Stockholms universitet

Real currency exchange rate prediction.
- A time series analysis.

Malin Varenius

Kandidatuppsats i matematisk statistik
Bachelor Thesis in Mathematical Statistics

Kandidatuppsats 2017:13
Matematisk statistik
Juni 2017

www.math.su.se

Matematisk statistik
Matematiska institutionen
Stockholms universitet
106 91 Stockholm

Mathematical Statistics
Stockholm University
Bachelor Thesis **2017:13**
http://www.math.su.se

Real currency exchange rate prediction. - A time series analysis.

Malin Varenius*

June 2017

Abstract

The foreign exchange market is the largest financial market in the world and forecasting exchange rates are not solely an important task for investors, but also for policy makers. Since market participant do not have access to future information, they try to model the exchange rate by past information. In this thesis an ARIMA(1,1,0) and a VAR(1) model with the trade balance in the EU and the interest rate differential as additional variables are evaluated in a forecasting purpose. It is concluded that a VAR(1) generates the most accurate forecasts during a 1-month horizon, while the ARIMA(1,1,0) is the more suitable model during a 3-month horizon. Both model outperforms a random walk, which usually is considered to produce the most accurate forecasts.

*Postal address: Mathematical Statistics, Stockholm University, SE-106 91, Sweden. E-mail: malin-varenius@hotmail.com. Supervisor: Filip Lindskog and Mathias Lindholm.

Acknowledgements

I would like to thank my supervisors Mathias Lindholm and Filip Lindskog for their open minds and valuable suggestions during our discussions. I would also like to thank my fellow students for not only elaborating new ideas, but also for continuously remembering me of the importance of having fun while learning.

Contents

Abbreviations

ACF autocorrelation function

ADF augmented Dickey-Fuller

AIC Aikaike information criteria

AICc Aikaike information criteria with correction for finite samples

AR autoregressive

ARIMA autoregressive integrated moving average

ARMA autoregressive moving average

BIC Bayesian information criteria

ECB European central bank

FPE final prediction error

HQ Hannan-Quinn

i.i.d. independent and identically distributed

MA moving average

MLE maximum likelihood estimation

OLS ordinary least square

PACF partial autocorrelation function

RER real exchange rate

RIR real interest rate

RIRD real interest rate differential

SARIMA seasonal ARIMA

SC Schwarz criterion

TB trade balance

VAR vector autoregressive

VECM vector error correction model

1 Introduction

The foreign exchange market is the largest financial market in the world and forecasting exchange rates are not solely an important task for investors, but also for policy makers. The exchange rate has direct impact on nations' international trade, economic growth as well as on their interest rate. Thus, in a globalized world it is just as important for small open economies as for large economies to understand what causes exchange rate fluctuations. However, the international financial market is rapidly changing due to the constant access generated by electronic trading (King and Rime, 2010). The rapid changes cause the currency investments to entail an inevitable and uncontrollable risk. As a result, investors and policy makers constantly try to forecast the change of the exchange rate in an attempt to minimize the risk of holding currency.

The main purpose of this thesis is to present a validate model which is able to forecast the real EUR/USD exchange rate in a statistical satisfying way. We will in our pursuit of the best suitable model first try to explain the exchange rate by its historical values in a linear manner by determining a univariate ARIMA model in section 4. However, macroeconomic literature often suggest that the exchange rate is better modelled by other economic variables. For example, the International Fisher Effect theory states that the future spot exchange rate can be determined by the nominal interest rate differential. Although, since we in this thesis try to model the real exchange rate, we will instead make use of the related "real exchange rate - real interest rate differential" (RERI) relationship when making exchange rate predictions. This differential variable as well as the trade balance in Europe will be the supplementary endogenous variables in our vector autoregressive model, also referenced to as the economic model, in section 5. These two models' predicting capability will then be compared to a random walk in section 6 by computing forecast values for three succeeding months and then compare these to the observed ones.

First following this introduction is a brief review of the existing literature on exchange rate modelling as well as a synopsis of some recent papers on exchange rate forecasting. Thereafter, the theoretical framework in this thesis will be outlined in section 3 including theory of specific models and test before we begin to analyse the models above. Lastly, a discussion an conclusions will be made in section 6 and 7 respectively.

2 Literature review and previous work

There exist extensive literature on exchange rate modelling and forecasting. The numerous modelling approaches clearly emphasise the challenging nature of finding a representative model describing the fluctuations in the foreign exchange market. And as yet in literature, there is no specific model approach that fruitfully elucidate the changes of the EUR/USD exchange rate. The problem of forecasting was illustrated by Meese and Rogoff in 1983 when the authors compared out-of-sample forecasts from both structural and time series models. Meese and Rogoff (1983) find that although the models fit very well in-sample, none of the models make more accurate point forecasts than a random walk, when the forecast accuracy was compared by computing the root mean squared forecast error. Since the journal article was written, many authors have tried to refine the models used, particularly by incorporating the fact that the regressand in the study, the natural logarithm of the exchange rate, likely is non-stationary.

Akincilar, Temiz and Şahin (2011) fit several models to daily data in a forecasting purpose of the USD/TL, EURO/TL and POUND/TL and finds that the autoregressive integrated moving average (ARIMA) models gives comparable accurate forecasts. Additionally, Ayekple et al (2015) consider an ARIMA model for predicting the dynamics of the Ghana cedi to the US dollar. They find small differences between the out-of-sample forecasts for the ARIMA and the random walk. However, some literature emphasises the fact that fundamental macroeconomic variables may contain predictive power for exchange rate movements in the long-term. Weisang and Awasu (2014) presents three ARIMA models for the USD/EUR exchange rate using data of monthly macroeconomic variables and concludes that the exchange rate is best modelled by a linear relationship of its past three values and the past three values of the log-levels share price index differential.

Another traditionally-used linear time series model that incorporate multivariate systems is the vector autoregressive model (VAR) and the vector error correction model (VECM). Yu (2001) examines the monthly exchange rate for three North European countries by employing a VAR, restricted VAR, VECM and a Bayesian VAR with several macroeconomic variables such as domestic and foreign money supply, output, short-term interest rate and price level. The conclusions are that the random walk has better forecasting accuracy in the short term but that the models beat the random walk in the long term. Additionally, Mida (2013) compare 12 out-of-sample forecasts of the monthly USD/EUR exchange rate between a random walk and a VAR with inflation, interest rate, unemployment rate and industrial production index. Mida (2013) concludes that the VAR model outperforms the random walk in the short term, namely one to three months, but is heavily outperformed in the longer horizon of six, nine and twelve months. Furthermore, Sellin (2007) evaluates the forecast ability of the Swedish Krona's real and nominal effective exchange rate by estimating a VECM model. Sellin includes a cointegrating relationship between real exchange rate, relative output, net foreign assets and the trade balance and finds the model to make accurate forecast once the model has been augmented with an interest rate differential.

In this thesis, we aim to construct an adequate model for EUR/USD real exchange rate forecasting. Due to earlier research with varying outcomes, we first use past values to predict future values (our ARIMA model). Secondly, a VAR model with interest rate differential and trade balance in the Euro zone as additional variables is estimated. At last, the two models' predicting abilities are evaluated.

3 Theory

In the following section the theoretical framework used in this thesis will be outlined, including theory for the specific models and tests.

3.1 Macroeconomic variables

The theory in the following two subsections is from Blanchard, Amighini and Giavazzi (2013).

3.1.1 Real exchange rate

The real exchange rate (RER) compares the purchasing power of two currencies at the current nominal exchange rate and prices. Thus, the real exchange rate can be expressed as

$$\mathrm{RER} = \mathrm{e} \cdot \frac{\mathrm{P}^*}{\mathrm{P}} \tag{1}$$

where e is the nominal exchange rate expressed as the domestic currency price of a foreign currency, P^* the foreign price of a market basket and P the domestic price of a market basket.

3.1.2 Real interest rate

The real interest rate (RIR) is the rate of interest an investor receives after accounted for the inflation rate. The Fisher equation formally expresses the RIR as

$$\mathrm{RIR} \approx \mathrm{i} - \pi \tag{2}$$

where i is the nominal interest rate and π the inflation rate.

3.1.3 The RERI relationship

The real exhange rate - real interest rate differential (RERI) relationship is central to most open economy macroeconomic models and the reduced form of the equation is

$$\mathrm{RER} = \mu + \beta(\mathrm{RIR}_t - \mathrm{RIR}^*{}_t) + \mathrm{w}_t \tag{3}$$

where the RER and RIR variables follow the previous notations in (1) and (2) respectively and RIR* denotes foreign RIR and w_t is a disturbance term (Hoffman and MacDonald, 2003). The term $\mathrm{RIR}_t - \mathrm{RIR}^*{}_t$ is called the real interest rate differential (RIRD).

3.2 Time series

A time series $\{X_t\}$ is a set of observations $\boldsymbol{x}_t$ indexed in time order t. If the observations in a time series are recorded at successive equally spaced points in time it is called a *discrete-time time series.* (Brookwell and Davis, 2002). These kind of time series will be dealt with in this thesis as the data points are recorded once every month.

3.2.1 Stationarity

The theory in this subsection can be found in Tsay (2010, chapter 2).

When performing different time series techniques one often assumes that some of the data's properties do not change over time. The most fundamental assumption is that the data is stationary. A time series $\{X_t\}$ is said to be *strictly stationary* if the joint distribution does not change when shifted in time. A more commonly and weaker version of stationarity is often used and that is when both the mean of $\{X_t\}$ and the covariance between $\{X_t\}$ and $\{X_{t-l}\}$ is time invariant, l being an arbitrary integer. This leads us to the definition of a weakly stationary time series:

Definition 3.1. A time series is said to be weakly stationary if

- $E[\mathrm{X}_t] = \mu$ and
- $Cov(X_t, X_{t-l}) = \gamma_l$

where μ is a constant and γ_l only depends on the lag length l.

Hence, the two first moments of the distribution is of interest when examining a time series weak stationarity properties. This is shown in a time plot as the data points fluctuating with a constant variance around a fixed mean. A time series that are stationary in levels is denoted I(0), whereas if a first difference is needed for the series to fulfil the requirements is denoted I(1). Weak stationarity is of special interest when one wants to make inference about future observations.

3.2.2 AR

The theory in the following five subsections can be found in Cryer and Chan (2008).

The autoregressive (AR) model is used when the output variable depends linearly on its past values plus an innovation term e_t that incorporates everything new in the series at time t that the past values fail to explain. Specifically, a pth-order autoregressive process $\{X_t\}$ can be expressed as

$$X_t = \phi_1 X_{t-1} + \phi_2 X_{t-2} + ... + \phi_p X_{t-p} + e_t, \tag{4}$$

where we assume e_t is independent of X_{t-1},X_{t-2},X_{t-3},... .

3.2.3 MA

The moving average (MA) process can be expressed as a weighted linear combination of present and past white noise terms. The moving average process of order q satisfies the equation

$$X_t = e_t - \theta_1 e_{t-1} - \theta_2 e_{t-2} - ... - \theta_q e_{t-q}. \tag{5}$$

3.2.4 ARMA

If a series have traits from both an autoregressive - and a moving average process, we say that the series is a mixed autoregressive moving average (ARMA) process. In general, if the series $\{X_t\}$ can be expressed as

$$X_t = \phi_1 X_{t-1} + \phi_2 X_{t-2} + ... + \phi_p X_{t-p} + e_t - \theta_1 e_{t-1} - \theta_2 e_{t-2} - ... - \theta_q e_{t-q} \tag{6}$$

we say that $\{X_t\}$ is an ARMA(p,q) process.

3.2.5 ARIMA

If a time series does not exhibit the features connected to stationarity one looks for transformations of the data to generate a new series with the desired properties. If the data requires differencing to become stationary one talks about the class of autoregressive integrated moving average (ARIMA) models. These models are a generalization of the class of ARMA models discussed previously and with $\Delta X_t = X_t - X_{t-1}$ an ARIMA$(p, 1, q)$ takes the following form:

$$\begin{aligned}\Delta X_t = \phi_1 \Delta X_{t-1} + \phi_2 \Delta X_{t-2} + ... + \phi_p \Delta X_{t-p} + \\ e_t - \theta_1 e_{t-1} - \theta_2 e_{t-2} - ... - \theta_q e_{t-q}.\end{aligned} \tag{7}$$

3.2.6 SARIMA

If a time series is a non-stationary seasonal process one may use the important tool of seasonal differencing. The seasonal difference of period s for the series $\{X_t\}$ is denoted $\nabla_s X_t$ and is defined as

$$\nabla_s X_t = X_t - X_{t-s}$$

A process is said to be a multiplicative seasonal ARIMA (SARIMA) model with nonseasonal orders p, d and q, seasonal P, D and Q, and seasonal period s if the differenceed series ΔX_t satisfies

$$\Delta X_t = \nabla^d \nabla_s^D X_t \tag{8}$$

We say that $\{X_t\}$ is a SARIMA$(p, d, q)(P, D, Q)_s$ model with seasonal period s.

3.2.7 VAR

The theory in the succeeding two subsections can be found in Lütkepohl, Krätzig and Phillips (2004, chapter 3).

Ordinary models usually consider a unidirectional relationship where the variable of interest is influenced by the predictor variables, but not the opposite way. However, in many macroeconomic models the reversed is often also true - all the variables have an effect on each other. When studying a set of macroeconomic time series vector autoregressive (VAR) models are frequently used. The structure is that each variable is a linear function of past lags of itself and past lags of the other variables. With vector autoregressive models it is possible to approximate the actual process by arbitrarily choosing lagged variables. Thereby, one can form economic variables into a time series model without an explicit theoretical idea of the dynamic relations.

The basic model for a set of K time series variables of order p, a VAR(p) model, has the form

$$\mathbf{y}_t = \mathbf{A}_1\mathbf{y}_{t-1} + ... + \mathbf{A}_p\mathbf{y}_{t-p} + \boldsymbol{u}_t \tag{9}$$

where the $\mathbf{A}_i$'s are $(K\text{x}K)$ coefficient matrices and $\boldsymbol{u}_t$ is a vector of assumed zero-mean independent white noise processes. The covariance matrix of the error terms, E($\boldsymbol{u}_t\boldsymbol{u}_t$')=$\mathbf{\Sigma}_u$, then assumes to be time-invariant and positive definite. The error terms $u_{i,t}$ may be contemporaneously correlated, but are uncorrelated with any past or future disturbances and thus allowing for estimation following the ordinary least square (OLS) method. By introducing the notation $\boldsymbol{Y} = \left[y_1, ..., y_T\right]$, $\boldsymbol{A} = \left[A_1 : ... : A_p\right]$, $\boldsymbol{U} = \left[u_1, ..., u_T\right]$ and $\boldsymbol{Z} = \left[Z_0, ..., Z_{T-1}\right]$, where

$$\boldsymbol{Z}_{t-1} = \begin{bmatrix} y_{t-1} \\ \vdots \\ y_{t-p} \end{bmatrix}$$

the model can be expressed as

$$\boldsymbol{Y} = \boldsymbol{AZ} + \boldsymbol{U}.$$

and the OLS estimator of $\boldsymbol{A}$ is

$$\hat{\boldsymbol{A}} = \left[\hat{A}_1 : \ \dots : \hat{A}_p\right] = \boldsymbol{YZ}'\left(\boldsymbol{ZZ}'\right)^{-1}.$$

The covariance matrix $\mathbf{\Sigma}_u$ may be estimated in the usual way. By denoting the OLS residuals as $\hat{\boldsymbol{u}} = \boldsymbol{y_t} - \hat{\boldsymbol{A}}\boldsymbol{Z}_{t-1}$ the matrix

$$\hat{\mathbf{\Sigma}}_u = \frac{1}{T - Kp}\sum_{t=1}^{T}\hat{\boldsymbol{u}}_t\hat{\boldsymbol{u}}_t' \tag{10}$$

where T is the number of observations and $\hat{\boldsymbol{\Sigma}}_u$ is an estimator which is consistent and asymptotically normally distributed independent of $\hat{\boldsymbol{A}}$.

Furthermore, the process is defined as stable if the determinant of the autoregressive operator has no root in/on the complex unit circle. Otherwise, some or all of the time series variables are integrated.

3.2.8 VECM

If the variables in the time series vector $\mathbf{y}_t$ has a common stochastic trend, there is a possibility that there exist linear combinations of the variables that are I(0), even though the individual time series are I(1) . This phenomenon is called *cointegration* and two or more variables are cointegrated if there exists a long run equilibrium relationship between them. In that case a vector error correction model (VECM) is useful since the model supports the analysis of the cointegration structure by combining levels and differences. The VECM is obtained from the VAR(p) model by subtracting $\mathbf{y}_{t-1}$ from both sides and rearranging. The result is the following form

$$\Delta \mathbf{y}_t = \Pi \mathbf{y}_{t-1} + \Gamma_1 \Delta \mathbf{y}_{t-1} + ... + \Gamma_{p-1} \Delta \mathbf{y}_{t-p+1} + \boldsymbol{u}_t \tag{11}$$

where $\Pi = -(\boldsymbol{I_k} - \boldsymbol{A_1} - ... - \boldsymbol{A_p})$ contains the cointegrating relations and is called the long run part. Likewise, $\Gamma_i = -(\boldsymbol{A_{i+1}} + ... + \boldsymbol{A_p})$, (i=1,...,$p-1$), is referred to as the short run or the short term parameters. The same assumptions about the error terms, $\boldsymbol{u}_t$, as in the VAR model also holds here.

3.3 Statistics and tests

The theory in the succeeding three subsections can be found in Tsay (2010, chapter 2).

3.3.1 Augmented Dickey-Fuller test

If a time series appears non-stationary one may verify the existence of a unit root in a AR(p) series by performing an augmented Dickey-Fuller (ADF) test. The null hypothesis $\mathrm{H}_0 : \beta = 1$ is tested against the alternative $\mathrm{H}_a : \beta \leq 1$ using the regression

$$X_t = c_t + \beta X_{t-1} + \sum_{i=1}^{p-1} \Delta X_{t-i} + e_t$$

where c_t is a deterministic function of the time index t and $\Delta X_j = X_j - X_{j-1}$ is the differenced series of X_t. Thus, the ADF-test is the t-ratio of $\hat{\beta} - 1$ expressed as

$$\text{ADF-test} = \frac{\hat{\beta} - 1}{std(\hat{\beta})} \tag{12}$$

where $\hat{\beta}$ is the least-squares estimate of β. The interpretation of the ADF-test is if the null hypothesis is rejected, then the time series is stationary.

3.3.2 Autocorrelation and partial autocorrelation functions

The autocorrelation function (ACF) is considered when the linear dependence between X_t and its past values X_{t-i} is of interest. The autocorrelation coefficient between X_t and X_{t-l} is denoted ρ_l which under the weak assumption of stationarity is a function of l only:

$$\rho_l = \frac{\text{Cov}(X_t, X_{t-l})}{\text{Var}(X_t)} \tag{13}$$

where $\rho_0 = 1$, $\rho_l = \rho_{-l}$ and $-1 \leq \rho_l \leq 1$.

The partial autocorrelation function (PACF) is a function of ACF and is the amount of correlation between a variable and a lag of itself that is not explained by correlations at all lower-order-lags. Considering the AR models:

$$\begin{aligned}
X_t &= \phi_{0,1} + \phi_{1,1}X_{t-1} + e_{1t}, \\
X_t &= \phi_{0,2} + \phi_{1,2}X_{t-1} + \phi_{2,2}X_{t-2} + e_{2t}, \\
X_t &= \phi_{0,3} + \phi_{1,3}X_{t-1} + \phi_{2,3}X_{t-2} + \phi_{3,3}X_{t-3} + e_{3t}, \\
X_t &= \phi_{0,4} + \phi_{1,4}X_{t-1} + \phi_{2,4}X_{t-2} + \phi_{3,4}X_{t-3} + \phi_{4,4}X_{t-4} + e_{4t}, \\
&\vdots
\end{aligned}$$

where $\phi_{0,j}$, $\phi_{i,j}$ and e_{it} are, respectively, the constant term, the coefficient of X_{t-i} and the error term of an AR(j) model. Since the equations are in the form of a multiple linear regression we may estimate the coefficients using the ordinary least-square method. The estimates $\hat{\phi}_{1,1}$, $\hat{\phi}_{2,2}$ and $\hat{\phi}_{k,k}$ of respective equation are called the lag-1, lag-2 and lag-k *sample PACF* of X_t. Thus, the complete sample PACF describes the time series' serial correlation with its previous values of a specific lag controlling for the values of the time series at all shorter lags.

By looking at the ACF and PACF plots one can tentatively identify the number of MA and AR terms needed. If the PACF displays a sharp cutoff and/or the lag-1 autocorrelation is positive, then the series could be explained by adding AR terms to the model. The lag at which the PACF cuts off is the indicated number of AR terms. In a similar manner, the lag at which the ACF cuts off indicates the number of MA terms. However, if both the ACF and PACF cuts off at a low lag order, a mixed ARMA model could be considered.

3.3.3 Ljung-Box Portmanteau test

A Ljung-Box Portmanteau test is performed to jointly test if several autocorrelations of X_t are zero. The null hypothesis $\mathrm{H}_0 : \rho_1 = ... = \rho_m = 0$ is tested against the alternative hypothesis $\mathrm{H}_a : \rho_i \neq 0$ for some $i \in \{1, ..., m\}$ with the test statistics

$$Q(m) = T(T+2) \sum_{l=1}^{m} \frac{\hat{\rho}_l^2}{T_l} \tag{14}$$

where T denotes the sample size, $\hat{\rho}_l^2$ the sample autocovariance at lag l and m the number of autocovariances tested. $Q(m)$ is asympototically a $\chi^2(m)$ variable under the assumption that X_t is i.i.d.. The null hypothesis is thus rejected if $Q(m) > \chi^2_\alpha$, where χ^2_α denotes the $100(1-\alpha)$th percentile of a chi-squared distribution with m degrees of freedom.

3.3.4 Johansen cointegration test

The theory in this subsection can be found in Burke and Hunter (2005, chapter 4).

Johansen cointegration test uses two test statistics do determine the number of cointegration vectors. The first, the maximum eigenvalue statistic, tests the null hypothesis of $H_0 : r \leq j-1$ cointegrating relations against the alternative of $H_a : r = j$ cointegrating relations for $j \in \{1, 2, ..., n\}$. It is computed as:

$$\mathrm{LR}_{max}(j-1, j) = -T \cdot \log(1-\lambda_j) = \lambda_{max}(j-1) \tag{15}$$

where T is the sample size. Thus, the null hypothesis of no cointegrating relationship against the alternative of one cointegrating relationship is tested by $\mathrm{LR}_{max}(0, 1) = -T \cdot \log(1-\lambda_1)$ where λ_1 is the largest eigenvalue.

The second test statistic, the trace statistic, tests the null hypothesis $H_0 : r \leq j-1$ against the alternative $H_a : r \geq j$ for $j \in \{1, 2, ..., n\}$, and is computed as:

$$\mathrm{LR}_{trace}(j-1, n) = -T \left[\sum_{i=j}^{n} \log(1-\lambda_i) \right] = \lambda_{trace}(j-1) \tag{16}$$

Both tests rejects the null hypothesis for large values of the test statistic. Thus, if cv stands for the critical value of the test and $\lambda(j-1)$ the statistic, the form of the test is:

$$\text{Reject } H_0 \text{ if } \lambda(j-1) > cv$$

The critical values for the two tests are different in general (except when $j = n$) and come from non-standard null distributions that are dependent on the sample size T and the number of cointegrating vectors being tested for. The interested reader can further read on the distribution theory leading to the critical values of the test in appendix D in Burke and Hunter (2005).

3.3.5 Model specification methods

There is a number of approaches to choose the right ARMA(p,q) model. One of the most common is the Aikaike information criteria (AIC). This criterion chooses the best model as the model that minimizes

$$\text{AIC} = -2\log(\text{maximum likelihood}) + 2k,$$

where k is the total number of parameters; $k = p + q + 1$ if the model contains an intercept or constant term and $k = p + q$ otherwise. The last term operates as a "penalty function" where larger models are penalised due to many parameters. This helps to ensure the selection of parsimonious models (Cryer and Chan, 2008).

Another approach is to select the model that minimizes the Bayesian information criteria (BIC). This criterion is defined as

$$\text{BIC} = -2\log(\text{maximum likelihood}) + k\log(n) \tag{17}$$

and is known to return consistent p and q orders when the true process follow an finite ARMA(p,q) process. On the other hand, if the process is not of a finite order ARMA process, the AIC will return the best suitable model that reflects the true process.

Even though AIC is commonly used in model selection, it should be known that it is a biased estimator and that the bias can be noticeable for large parameter per data ratios. Though, the Aikaike information criteria with correction for finite samples (AICc) is an estimator which is shown to approximately eliminate the bias by adding one more penalty term to the AIC. It is defined as

$$\text{AICc} = \text{AIC} + \frac{2(k+1)(k+2)}{n-k-2}, \tag{18}$$

where n is the sample size. It is suggested that for cases where $k/n \geq 10\%$ AICc outperforms many selection criteria, including the AIC and BIC (Huruvich and Tsai, 1989).

Furthermore, when deciding on the appropriate lag order in a VAR model there are three other model selection criteria used, the Hannan-Quinn (HQ), Schwarz criterion (SC) and final prediction error (FPE) (Pfaff, 2008). They are defined as

$$\text{HQ} = \log\,\det(\sum_u(p)) + \frac{2\log(\log(T))}{T}pK^2, \tag{19a}$$

$$\text{SC} = \log\,\det(\sum_u(p)) + \frac{\log(T)}{T}pK^2, \tag{19b}$$

$$\text{FPE} = \Big(\frac{T+p^*}{T-p^*}\Big)^K \det(\sum_u(p)) \tag{19c}$$

where $\boldsymbol{\Sigma}_u(p) = T^{-1}\sum_{t=1}^{T}\hat{\boldsymbol{u}}_t\hat{\boldsymbol{u}}'_t$ and p^* is the total number of parameters in each equation and p the lag order.

3.3.6 Root mean square error

Since the aim of this thesis is to construct and compare models for real exchange rate forecasting we need a measure of the models' adequacy. The root mean square error (RMSE) measures the actual deviation from the predicted value to the observed value.

Let $\hat{X}_i$ be the predicted value of the corresponding observed values X_i at time i, then the root mean squared error is computed as

$$\text{RMSE}(\hat{X}) = \sqrt{\text{MSE}(\hat{X})} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(\hat{X}_i - X_i)^2} \tag{20}$$

where n is the number of observations.

3.4 Forecasting time series models

As the main purpose of this thesis is to compare different models' forecasting accuracy by computing the root mean square error we need to be able to make out-of-sample forecasts. An ARIMA(1,1,1) model will be illustrated as an example of the procedure of ARIMA forecasting . The example below is a modified example of the one illustrated in Hyndman and Athanasopoulos (2014, chapter 8.8).

An ARIMA(1,1,1) is on the form

$$\Delta X_t = \phi_1 \Delta X_{t-1} + e_t - \theta_1 e_{t-1}$$

and by utilizing that $\Delta X_t = X_t - X_{t-1}$, the equation above may be written as

$$X_t = (1 + \phi_1 X_{t-1}) - \phi_1 X_{t-2} + e_t - \theta_1 e_{t-1}. \tag{21}$$

Hence, by replacing t by $T+1$ we get

$$X_{T+1} = (1+\phi_1)X_T - \phi_1 X_{T-1} + e_{T+1} - \theta_1 e_T$$

where, by assuming observations up to time T, all values on the right hand side are known except for e_{T+1}, which is replaced by zero, and e_T which is replaced by the last observed residual $\hat{e}_T$. Thus, the forecasted value in time $T+1$ is

$$\hat{X}_{T+1|T} = (1+\phi_1)X_T - \phi_1 X_{T-1} - \theta_1 \hat{e_T}.$$

Furthermore, a forecast of X_{T+2} is obtained by instead replacing t with $T+2$ in equation (20) above. All values on the right hand side will be known at time T except X_{T+1}, which is replaced by $\hat{X}_{T+1|T}$, and e_{T+2} and e_{T+1}, which both are replaced by zero. A forecasted value of X_{T+2} is thus given by

$$\hat{X}_{T+2|T} = (1+\phi_1)\hat{X}_{T+1|T} - \phi_1 X_T.$$

The process continues in this manner to get point forecasts for all future time periods.

When plotting predicted values one usually depicts a shaded prediction interval in the plot. A prediction interval is an estimated interval in which future observations will fall, with a certain probability, given what has already been observed. The calculation of ARIMA prediction intervals is difficult and the derivation cumbersome, not providing a simple interpretation. Thus, the interested reader can find the details of multi-step forecast intervals in Brockwell and Davis (2002, chapter 6.4). However, the first prediction interval is easy to compute and is given by

$$\hat{X}_{T+1|T} \pm c_{1-\gamma/2}\hat{\sigma}, \tag{22}$$

where $c_{1-\gamma/2}$ is the $(1-\gamma/2)\cdot 100$ percentage point of the standard normal distribution and $\hat{\sigma}$ the standard deviation of the residuals.

In a similar manner, one may forecast a VAR model according to theory in Lütkepohl, Krätzig and Phillips (2004, chapter 3). Forecasts are generated in a recursive manner for *each* variable in the VAR model following the notation in equation (9). Assuming a fitted VAR model by OLS for all observations up to time T the h-step ahead forecast is generated by

$$\hat{\boldsymbol{X}}_{T+h|T} = \hat{\boldsymbol{A}}_1 X_{T+h-1} + \hat{\boldsymbol{A}}_p X_{T+h-1}.$$

The corresponding forecast error is

$$\boldsymbol{X}_{T+h} - \hat{\boldsymbol{X}}_{T+h|T} = \boldsymbol{u}_{T+h} + \boldsymbol{\phi}_1 \boldsymbol{u}_{T+h-1} + \ldots + \boldsymbol{\phi}_{h-1}\boldsymbol{u}_{T+1}$$

where it can be shown by successive substitution that

$$\boldsymbol{\phi}_s = \sum_{j=1}^{s} \boldsymbol{\phi}_{s-j}\boldsymbol{A}_j, \quad s = 1, 2, \ldots,$$

with $\boldsymbol{\phi}_0 = \boldsymbol{I}_k$ and $\boldsymbol{A}_j = 0$ for $j > p$. The u_t is the one-step forecast error in period $t-1$ and as the forecasts are unbiased, the forecast errors have expectation 0. The mean square error of an h-step forecast is

$$\boldsymbol{\Sigma}_y = \boldsymbol{E}\{(\boldsymbol{X}_{T+h} - \hat{\boldsymbol{X}}_{T+h|T})(\boldsymbol{X}_{T+h} - \hat{\boldsymbol{X}}_{T+h|T})'\} = \sum_{j=0}^{h-1} \boldsymbol{\phi}_j \boldsymbol{\Sigma}_u \boldsymbol{\phi}_j'.$$

If the process $\boldsymbol{X}_t$ is Gaussian, implying $\boldsymbol{u}_t \in$ i.i.d. $N(0, \Sigma_u)$, then the forecast errors follow a multivariate distribution. Thus the prediction interval is given by

$$X_{k,T+h|T} \pm c_{1-\gamma/2}\hat{\sigma}_k(h) \tag{23}$$

where $\hat{\sigma}_k(h)$ is the square root of the kth diagonal element of $\boldsymbol{\Sigma}_y(h)$

4 Univariate analysis

The most important step in building dynamic econometric models is to get an understanding of the characteristics of the individual time series variables involved. This section will contain a thorough analysis of the EUR/USD real exchange rate as well as a somewhat lighter analysis of the real interest rate differential and the trade balance in the Euro zone, all series covering the period from January 1999 to February 2016. The analysis is important since the properties of the individual series will increase our understanding when we later analyse them together in a system of series. The central part of this univariate time series analysis is to discover if the series are stationary, since this will play a role in our vector autoregression modelling later. However, we will perform a more extensive analysis of the real exchange rate series with the aim to find a suitable ARIMA model for this series alone, as this is one of the models whose predicting ability will be compared.

It is worth to mention that there are many other important determinants of exchange rate changes that involves political and economic stability and the demand for a country's goods and services. However, increasing the number of variables in a single time series study does not generally lead to a better model since it makes it more difficult to capture the dynamic, intertemporal relationships between them. Therefore, the multiple time series analysis that follows in section 5 will focus on three variables, namely the real interest rate differential, the trade balance in the Euro zone as well as the variable of interest; the real exchange rate.

4.1 EUR/USD real exchange rate

The data for the real exchange rate in the following analysis is downloaded from the Statistical Data Warehouse of the European central bank (ECB). It is measured as the European price level relative to the American price level expressed in US dollar.

We start our analysis of the EUR/USD real exchange rate by plotting the series in figure 1. As visible in the plot, the exchange rate fluctuates considerably and appears to exhibit non-stationarity. This appearance is confirmed by the slow decay of the upper autocorrelation function plot in figure 2 as well as the augmented Dickey-Fuller (ADF) test presented in table 1. To be able to preform different time series techniques on the data it is important to adjust for the non-stationarity. Thus, a first difference is applied to the series in an attempt to achieve a stationary time series. Table 1 contains the ADF test when a first difference is applied to the series. The lag order is determined by the `ar` function which uses the AIC criterion and the result is later compared to ACF and PACF plots if it seems reasonable. Since the plot of the real exchange rate in levels in figure 1 and the plot of the differenced series in figure 13 in appendix B do not reveal a linear deterministic trend a priori, both series is tested based on a model without a trend. Furthermore, the mean of the differenced series is not significantly different from zero when a t-test is preformed and is therefore tested without a constant. On the contrary, the mean of the levels series is significantly nonzero and is thus tested based on a model with a constant. The test results in table 1 indicate that the series is not stationary in levels but stationary in first difference. The appropriateness of a first difference is also demonstrated in the differenced ACF in figure 2 where there is now a drop to zero quickly. It is therefore concluded that the RER is integrated of order 1, I(1), and that we cannot reject a unit root in the levels series.

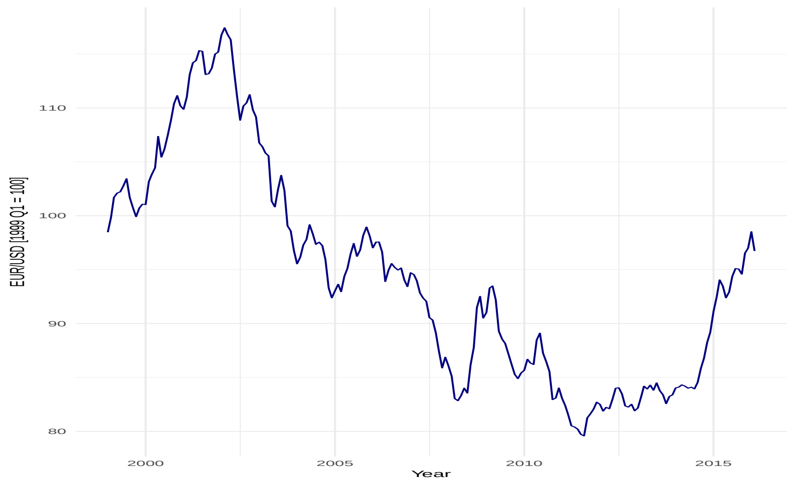

Figure 1: The monthly real Euro/US dollar exchange rate

Augmented Dickey Fuller Unit Root test			
Variable	Deterministics	Lag order	P-value
RER	constant	2	0.6063
Diff RER		1	<0.01

Table 1: Unit root test for the levels and first differenced Real Exchange Rate series

In an attempt to find a suitable ARIMA model the autocorrelation function (ACF) and partial autocorrelation function (PACF) are used. The correlograms are visible in figure 2 for the levels and the differenced series respectively. As stated previously, the slow decay of the positive sample autocorrelation function for the levels series suggests the appropriateness of an ARIMA model.

Since the first lag is significant in both the ACF and PACF plot of the differenced RER series potential p and q values are 1. Therefore, a possible candidate is an ARIMA(1,1,1) model. To investigate the matter further we employ the `auto.arima` function in `R` which uses a variation of the Hyndman and Khandakar algorithm [5] to obtain a suitable ARIMA model. This algorithm combine unit root tests, minimization of the AICc and maximum likelihood estimation (MLE). The reader is referred to appendix A for further information of the steps.

As a result, the function returns an ARIMA(2,1,2) model with a seasonal AR at the 12th lag of order 1, a SARIMA$(2,1,2)(1,0,0)_{12}$ model. Since it is not advantageous from a forecasting point of view to choose large p and q we restrict the `auto.arima` function to not look for seasonal components. This decision is based on the observation that we do not see a clear seasonal pattern in the series or a single significant spike at lag 12 in the PACF. With the imposed restriction, the `auto.arima` function return an ARIMA(1,1,0) model. Consequently, this initial analysis suggests three potential models for the real exchange rate series. The models are fitted and the supplied information criterion values computed. The results can be found in table 2 below.

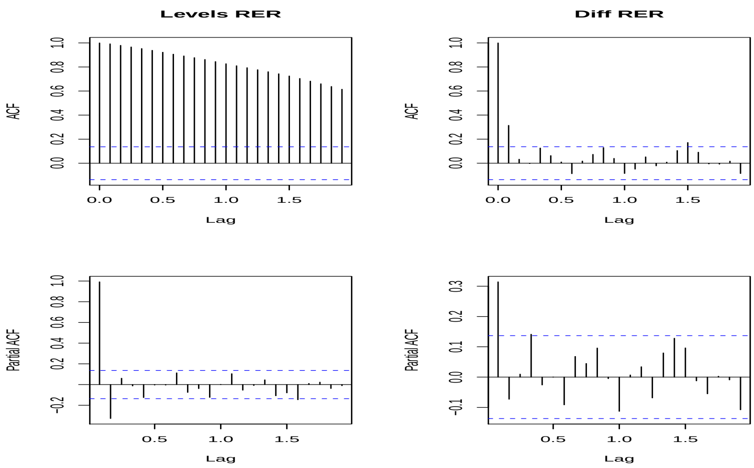

Figure 2: The acf and pacf for the levels series to the left and for the differenced series to the right. The lags on the y-axis are scaled by a 10^{-1}-multiple.

The obtained AIC and AICc values in table 2 advocate the SARIMA$(2,1,2)(1,0,0)_{12}$ as the most suitable model out of the three considered. However, the BIC implies that the ARIMA(1,1,0) is the appropriate model and that the SARIMA model is the least appropriate model. Though, as one can see, the values obtained do not differ greatly between the models and since there were not much evidence of a seasonal effect in our data we

Information criterion values			
Model	AIC	AICc	BIC
ARIMA(1,1,1)	641.16	641.28	651.13
SARIMA(2,1,2)(1,0,0)$_{12}$	639.63	640.05	659.56
ARIMA(1,1,0)	640.2	640.26	646.84

Table 2: Information criterion values for the different suitable ARIMA models

choose to proceed our analysis with the ARIMA(1,1,1) model as well as the ARIMA(1,1,0).

Thereon, we verify that the residuals of the fitted models possess the wanted properties by first establishing that they are normally distributed by normal quantile-quantile plots in figure 3 below.

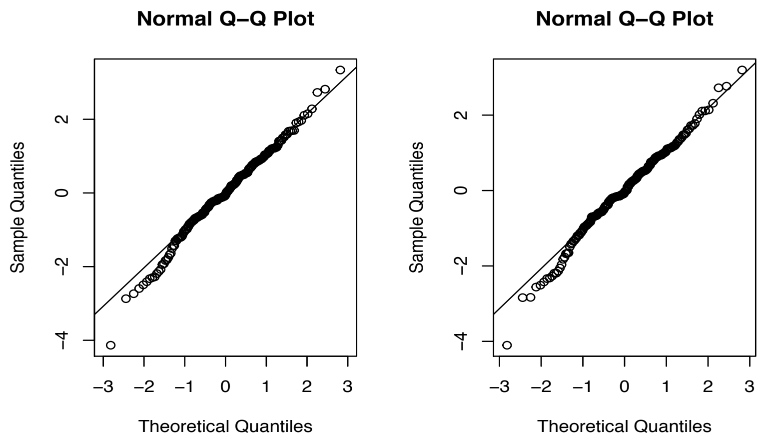

Figure 3: Normal Q-Q plots of the residuals from the ARIMA(1,1,1) model to the left and the ARIMA(1,1,0) model to the right.

The linearity of the points in figure 3 suggests that the data are normally distributed. There is one data point that deviates from the rest and we detect it as the one for November 2008. This is due to the financial crisis and a possible model for step response intervention could be employed but as the data point do not seem to interfere with the normality assumption more than appearing as an outlier we simply leave it and proceed with the possibility of changing the model specification with a dummy variable later.

Thereon, we continue by plotting the standardized residuals together with the sample ACF. These can be found in figure 14 and 15 in appendix B. The plots of the standardized residuals obtained from both models gives no indication of a nonzero mean, trend or changing variance and thus resembles white noise. The sample ACF of the residuals further indicates that they are independent and identically distributed (i.i.d.) since no significant autocorrelation is present. Lastly, a Ljung-Box Portmanteau test of the residuals is performed. The test statistics leads us not to reject the null hypothesis of independence and the associated p-values for all lags up to 10 can as well be viewed in figure 14 and 15 in appendix B. As a result, there is no cause to reject the fitted models. We continue by plotting the one-step in-sample forecast generated from the fitted models in figure 4 and 5.

Figure 4: Time series plot of the in-sample forecast generated by the ARIMA(1,1,1) in pink and the EUR/USD real exchang rate in blue

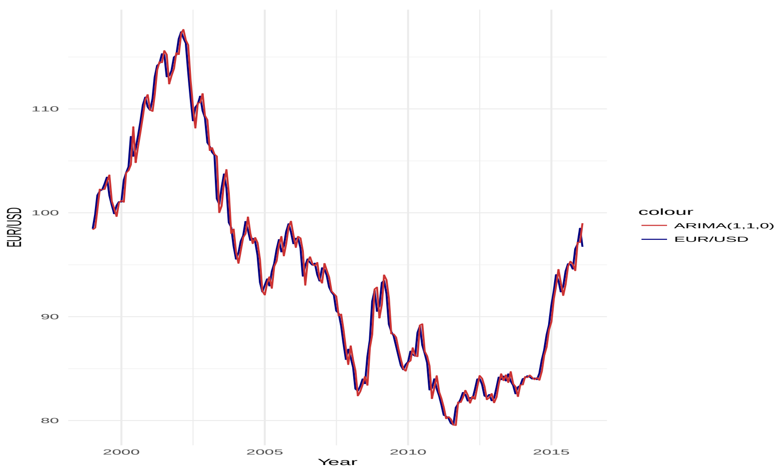

Figure 5: Time series plot of the in-sample forecast generated by the ARIMA(1,1,0) in red and the EUR/USD real exchang rate in blue

The lines of the fitted ARIMA models look almost the same which is also strengthen by the similar values of the squared correlation between the observed and in-sample forecast values, approximately 0.9874 and 0.9873 respectively (a difference of 5.66e-5). This states that both models make a good description of the past. However, a good fit does not necessarily lead to a good forecast. For example, overfit models will usually have very small in-sample errors, but not lead to favourable out-of-sample forecasts. Hence, we will return to the fitted ARIMA model equations and find a bootstrap confidence interval of the coefficients using model-based residual resampling conditioned on the first $(p + d)$ initial values and test for these being zero. The fitted ARIMA models are given by the following equations

$$x_t = \underset{(0.1768)}{0.1496} x_{t-1} + e_t + \underset{(0.1715)}{0.1913} e_{t-1} \tag{24a}$$

$$x_t = \underset{(0.0666)}{0.3185} x_{t-1} + e_t \tag{24b}$$

where $x_t = X_t - X_{t-1}$, $x_{t-1} = X_{t-1} - X_{t-2}$ and e_t is the random shock noise occurring at time t. The corresponding standard errors are presented under respective parameter estimation. First, the associated parametric 95% confidence interval of the coefficients are presented in table 3 below.

Parametric 95% Confidence Intervals		
Model	AR1	MA1
ARIMA(1,1,1)	(-0.197,0.496)	(-0.145,0.527)
ARIMA(1,1,0)	(0.188,0.449)	

Table 3: Parametric 95% confidence intervals of the parameters in the fitted models

Table 3 reveals that both of the parameters in the fitted ARIMA(1,1,1) model has 0.0 inside the interval and thus indicating insignificant parameters. For comparison, the bootstrap confidence intervals can be found in table 4.

Bootstrap 95% Confidence Intervals		
Model	AR1	MA1
ARIMA(1,1,1)	(-0.688,0.868)	(-0.762,0.823)
ARIMA(1,1,0)	(0.183,0.448)	

Table 4: Boostrap 95% confidence intervals of the parameters in the fitted models

The confidence interval based on the bootstrap method for the ARIMA(1,1,0) yield a similar interval as the parametric whereas the bootstrap confidence interval for the ARIMA(1,1,1) is wider, which is generally the case, and still includes 0.0. Therefore, we conclude that the parameter estimates in the ARIMA(1,1,1) model is insignificant at the 5% level. To give the reader a more intuitive sense of the confidence intervals, histograms of respective bootstrap distribution for the parameter estimates can be found in figure 16 and 17 in appendix B.

Since the estimates of the coefficients in the ARIMA(1,1,1) are insignificant we discard this model as suitable for the real exchange rate series. Thus, we continue our analyze with solely the ARIMA(1,1,0) model. By utilizing that $x_t = X_t - X_{t-1}$ and $x_{t-1} = X_{t-1} - X_{t-2}$ equation (24b) can be written on the equivalent form

$$X_t = 1.3185X_{t-1} - 0.3185X_{t-2} + e_t \tag{25}$$

Its predicting ability will later be compared to the one of the economic model in section 7.

4.2 Macroeconomic variables

The real interest rate differential is computed as the 3-month real Euribor rate minus the 3-month Libor rate in USD adjusted for the inflation in the US. Moreover, the terms of trade is computed as the ratio between export prices and import price. The data for the 3-month Euribor rate is downloaded from the Statistical Data Warehouse of the ECB whereas the data for the export price index in the Euro zone is downloaded from Eurostat. Lastly the data for the import price index in the Euro zone, the inflation rate in the US as well as the 3 month Libor rate in USD is downloaded from the Federal Reserve Bank of St. Louis.

In this section a lighter analysis of the two macroeconomic variables; real interest rate differential (RIRD) and European trade balance (TB) will be made. We start by plotting the series in figure 6 and 7 respectively and we notice that both series appear non-stationary with high volatility around the immediate time and in the reverberations of the financial crisis. As a consequence of the non-stationary characteristics, a first difference is applied to both series and as observable in figure 18 and 19 in appendix B, the transformation seems to be satisfying in a stationary purpose. Though, the time series of the differenced RIRD appear to exhibit an outlier and we identify it as the one for December 2008. The outlier is not adjusted for now, but we may have to adjust the vector autoregressive model with a dummy variable if the multivariate normality assumption is violated.

Next, an ADF test is used to detect whether the differenced series possess a unit root. Both the series are modelled without a constant since the respective t-tests do not reject a zero mean. The resulting p-values are presented in table 5.

Augmented Dickey Fuller Unit Root test			
Variable	Deterministics	Lag order	P-value
Diff RIRD		1	<0.01
Diff trade balance EU		1	<0.01

Table 5: Unit root test for the differenced macroeconomical variables

The ADF tests conclude both series to be I(1). This discovery will play a significant role when we define an economic model in the next section since there is a possibility that the time series exhibit a common stochastic trend due to all being I(1). However, before we proceed with finding a valid vector model some descriptive statistics of the differenced series are displayed in table 6 and 7 below.

The informative descriptive statistics in table 6 reveal that both of the macroeconomic series, but principally the differenced TB, is closely oscillating around zero. Table 7 display the covariance of the three differenced series with respective sample variance on the diagonal.

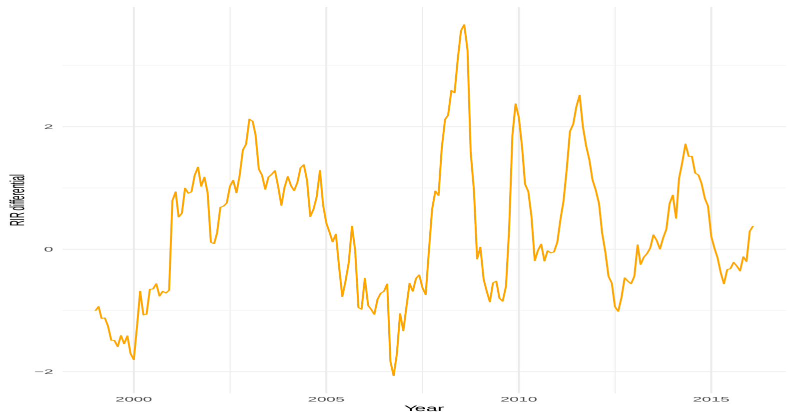

Figure 6: Time series plot of the real interest rate differential

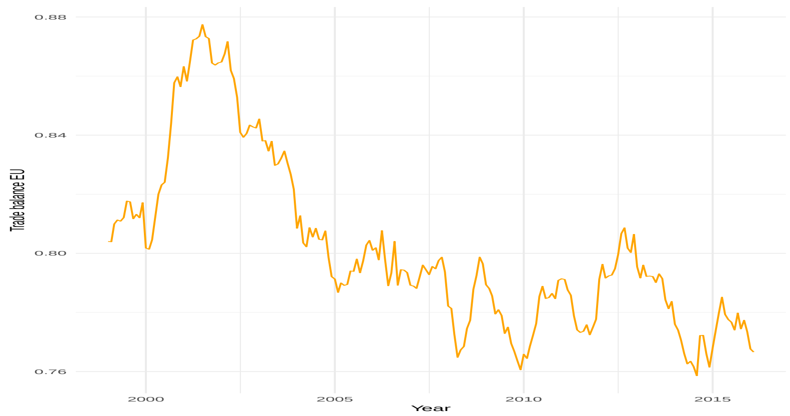

Figure 7: Time series plot of the trade balance in the Euro zone

Descriptive statistics			
Variable	Mean	Highest value	Lowest value
Diff RIRD	0.0068	1.5322	-1.6653
Diff trade balance EU	-0.0002	0.01372	-0.0154

Table 6: Descriptive statistics for the differenced real interest rate differential as well as for the differenced trade balance.

Covariance matrix			
	Diff RER	Diff RIR	Diff TB
Diff RER	1.4564	-0.1285	0.0017
Diff RIR	-0.1285	0.1492	-0.0003
Diff TB	0.0017	-0.0003	2.697e-05

Table 7: Covariance matrix of the three differenced series. The figures are rounded to 4 decimals.

5 Economic model

In this section, we will define the model which includes endogenous macroeconomic variables. Since the results of the stationarity test according to ADF indicate that all series are stationary in their first difference, we initiate this section by testing for cointegration to detect if there exist a linear combination of the variables that is stationary. The results of the Johansen test for cointegration using the maximum eigenvalue statistic can be found in table 8 whereas the test with trace statistic can be found in table 9.

	Statistic	10%	5%	1%
r $\leq$ 2	1.27	7.52	9.24	12.97
r $\leq$ 1	2.44	13.75	15.67	20.20
r = 0	14.86	19.77	22.00	26.81

Table 8: Johansen test for cointegration rank: Max eigenvalue statistic

	Statistic	10%	5%	1%
r $\leq$ 2	1.27	7.52	9.24	12.97
r $\leq$ 1	3.71	17.85	19.96	24.60
r = 0	18.57	32.00	34.91	41.07

Table 9: Johansen test for cointegration rank: Trace statistic

Both test statistics reject the hypothesis of a cointegrating vector at all significance levels. Hence, we continue by estimating a VAR model with the differenced series. The lag order is first specified by the model selection criteria described in section 3.3.7. All the criteria selects a VAR(1) as the most suitable model. Values of the different criteria for up to six lags can be found in table 11 in appendix C.

Thereafter we proceed with estimating a VAR(1) model by utilizing ordinary least square per equation. The model is estimated without a constant term. This do not seem contradictory since we tested the sample means of the first-differenced series and none of the means were significantly different from zero. We therefore restrict the vector autoregressive model to exclude deterministic drifts in the individual series. As in the case of the ARIMA models in section 4, we initialize by performing a multivariate Ljung-Box Portmanteau test to test for autocorrelation in the residuals. The test is preformed to detect if the choice of one lag is to restrictive and if the data instead should be modelled with a higher lag order. Autocorrelation in the residuals is an unwanted trait, since autocorrelation biases the estimators and makes them less efficient. However, the test statistics leads us not to reject the null hypothesis of independent residuals (associated p-values for up to ten lags can be viewed in figure 20 in appendix B). Additionally, we compare the distribution of the residuals received from the RER equation to a normal distribution using a normal qunantile-quantile plot in figure 8.

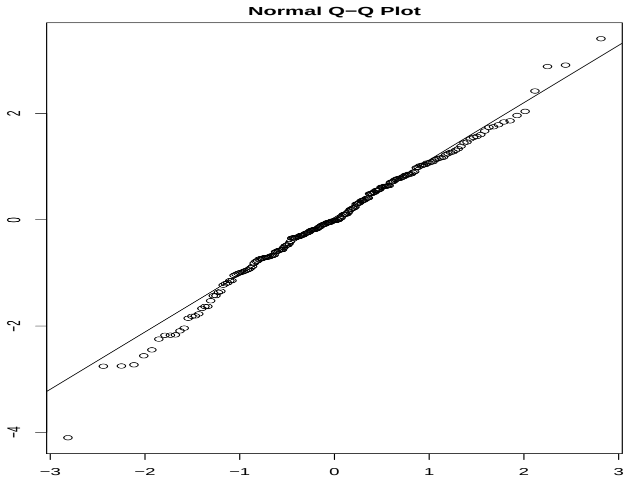

Figure 8: Normal Q-Q plot of the residuals from the RER equation in the VAR(1) model

The Q-Q plot in figure 8 do exhibit some kurtosis but not enough to violate the normal distribution assumption. Again, we observe an outlier and identify this as the one for November 2008. A modification of the VAR model can possibly be made where a dummy variable for this specific point could be added to adjust for the outlier. However, we forego doing so as this would require an additional parameter to estimate in the VAR(1) model and that the data point does not appear to interfere with the normality assumption. Further diagnostic plots can be found in figure 21 in appendix B. None of the plots give a reason to question a normal distribution assumption of the residuals and we may proceed with the VAR(1) model.

Usually, the estimated VAR coefficients are not reported since they are poorly estimated and, except for the first lag, often insignificant (Canova, 2007). However, since we will use the model in a forecasting comparison purpose, depicting the equation could result in a intuitive sense of the model. Consequently, the equation for the differenced RER takes the form

$$\mathrm{RER}_t = \underset{0.07148}{0.34768}\,\mathrm{RER}_{t-1} + \underset{0.21660}{0.07769}\,\mathrm{RIRD}_{t-1} \underset{16.06442}{-19.66126}\ \mathrm{TB}_{t-1} + u_{1t} \quad (26)$$

where the standard errors are stated below the corresponding parameter estimate and u_{1t} is the error term. All time series X_t in the above equation should be considered as $X_t - X_{t-1}$ since they are all differenced. The table with all of the estimated parameters can be found in table 12 in appendix C. The corresponding covariance matrix of all error terms, u_{1t}, u_{2t} and u_{3t}, in the VAR(1) model is given by

$$\boldsymbol{\Sigma}_u = \begin{bmatrix} 1.3094 & -0.1056 & 0.0008 \\ -0.1056 & 0.1327 & -9.252\text{e}-05 \\ 0.0008 & -9.252\text{e}-0.5 & 0.2170\text{e}-05 \end{bmatrix}$$

As a model validation check for a VAR model one can make use of the formula for covariance of the causal series in Brockwell and Davis (2004, section 8.4). The authors derive the covariance of the series for a VAR(1) with one lag as

$$\boldsymbol{\Sigma}_y = \boldsymbol{\Sigma}_u + \boldsymbol{A}\boldsymbol{\Sigma}_u\boldsymbol{A}'$$

By computing the matrix $\boldsymbol{A}$ with the parameter estimate table 12 in appendix C we get the variance for differenced RER series to 1.4707, which is similar to the value in the covariance matrix in section 4 which was 1.4564.

An initial observation of the RER equation is that the parameter estimate for the RER_{t-1} is quite similar to the estimate in the ARIMA(1,1,0) in section 3, with an additional value of approximately 0.03. Furthermore, it is worth to remark that the absolute contribution from the two macroeconomic variables are quite small considering that both of these differenced series are closely oscillating around zero. Thus, the resemblance of this model with the ARIMA(1,1,0) is palpable. Furthermore, a 95% bootstrap confidence interval of the parameter estimates indicates, as outlined by Canova (2007), only the lagged RER to be significant.

Instead of parameter estimates one usually report functions of the VAR coefficients which tend to summarize information better and have some economic meaning whilst they are generally more precisely estimated (Canova, 2007). Thus, before we move on to the main purpose of this thesis; forecasting, we analyze the responsiveness of the dependent variable RER to exogenous impulses, which in macroeconomic modelling often is referred to as shocks, to the other macroeconomic variables. This analysis will be carried out by impulse response functions and will describe what sign the effect has on the response variable as well as how long the effect will last. In this way we are able to detect dynamic relationships over time.

The idea is to look at the adjustment of the variable after a hypothetical shock at time t. This adjustment is then compared with the time series process without a shock, that is, the actual process. The impulse response sequences plot the difference between the two time paths.

- **Shocks to the real interest rate differential -** When the RIRD increase, the exchange rate is supposed to appreciate. This is because higher interest rate in a country make the currency more valuable relative to the country offering lower interest rate. We can see that this also is the case as depicted in figure 8 where the exchange rate responded positively on a shock to the RIRD.

- **Shocks to the trade balance -** As depicted in figure 9, an increase in the trade balance has a negative effect on the exchange rate. This is in line with theory; if the export price index of European goods increase, American importers will lower their demand for European currency and thus making the currency less valuable.

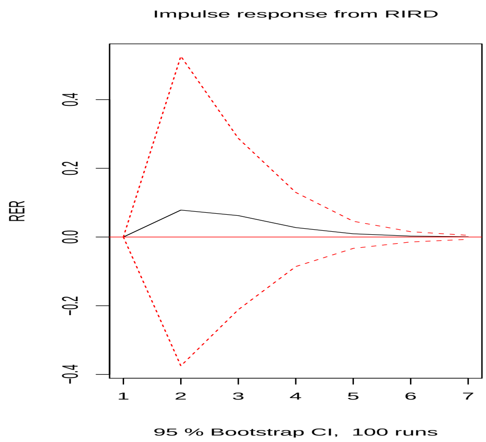

Figure 9: Impulse responses for the shocks to the real interest rate differential with a 95% bootstrap confidence interval

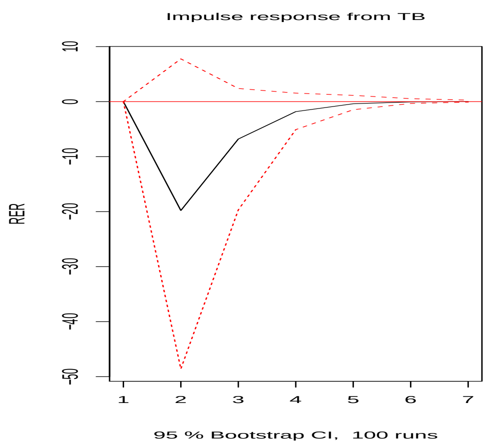

Figure 10: Impulse responses for the shocks to the trade balance with a 95% bootstrap confidence interval

Since the real exchange rate responds according to theory after shocks to the macroeconomic variables we will continue analyzing the estimated VAR(1) model in a forecasting purpose in the next section.

6 Forecasting

The main purpose of this thesis is to forecast the real EUR/USD exchange rate for the four months succeeding February 2016 and compare these predicted values to the observed values. We will compare the forecast accuracy of the ARIMA(1,1,0) and the VAR(1) with a random walk, which is often considered to produce the best forecasts. The predicted values from the random walk is thus the value of the latest observation, $\hat{X}_t = X_{t-1}$, which corresponds to 96.7301. The choice of test period is set to three months. This is due to the fact that an ARIMA(1,1,0) would converge to a random walk when the amount of forecast steps increase.

Below, three out of sample forecasts from the ARIMA model are depicted together with a lighter shaded 80 % and a darker 95 % prediction interval in figure 11 whereas the three out of sample forecast from the VAR model are depicted with a corresponding 95 % prediction interval in figure 12. The intervals are calculated by using the formulas for prediction interval in equation (22) and (23) respectively. The variance of the ARIMA(1,1,0) model was estimated to 1.31 and the variance of the RER equation in the VAR(1) model was derived and estimated to 1.4707 in section 5.

It is difficult to notice any obvious differences between the predicted values based on the plots and this is not surprising since the two fitted models are similar, as remarked in section 4. After examining the point forecasts more thoroughly we conclude that both series of the predicted exchange rates are steadily depreciating for the three out of sample forecasts, where the first value of the forecast for March 2016 from the VAR(1) model is slightly lower than the corresponding forecast obtained from the ARIMA(1,1,0) model. However, the two point forecast obtained for April and May 2016 are lower for the ARIMA model. The predicted values can be viewed in table 13 in appendix C.

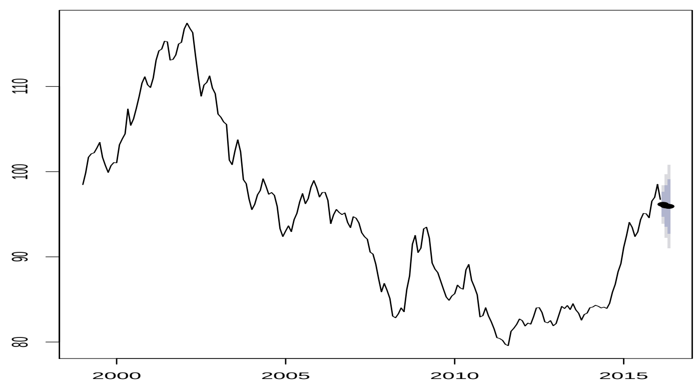

Figure 11: Three out of sample forecasts from ARIMA(1,1,0)

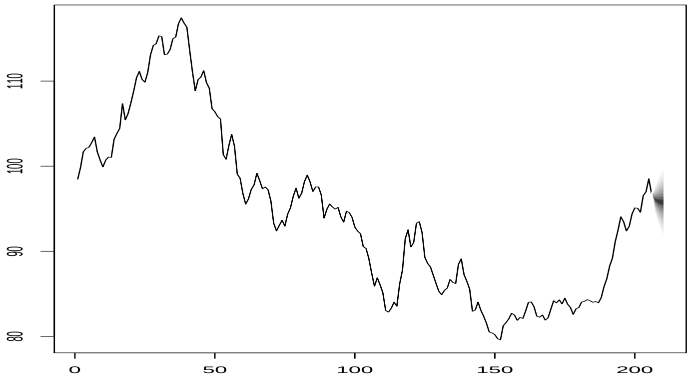

Figure 12: Three out of sample forecasts from VAR(1)

To compare the forecast accuracy, the root mean square error (RMSE) of the out of sample forecasts are computed. The values can be depicted in table 10 where the values for the random walk (RW) model also are computed.

Time horizon	ARIMA(1,1,0)	VAR(1)	RW
1 month	0.8302	0.8053	1.3987
3 months	1.5722	1.6001	2.2714

Table 10: Root mean square errors for the ARIMA, VAR and RW models.

According to the RMSE in table 10, the VAR(1) outperforms both the ARIMA(1,1,0) and the RW during a 1-month horizon. However, the ARIMA model produce more accurate predictions during the 3-month horizon. This contradicts the finding by Meese and Rogoff (1983) and a possible explanation is that we have incorporated that the real exchange rate time series is non-stationary in levels, a feature that was not considered in their paper. Another explanation is that EUR/USD real exchange rate experiences a period of successive deprecation causing the forecasts from the random walk to continuously overshoot.

Nonetheless, in our setting, with access to the actual values for the three succeeding months, we may conclude that the random walk is outperformed by both models and that the VAR model seem to be better used in a one-step forecast purpose. However, the rather small difference in RMSE values during the 1-month horizon between the two models implies that a simple ARIMA(1,1,0) captures dependencies in the data in a sufficient manner.

7 Discussion

The purpose of this thesis was not to discover new results in the field of exchange rate modelling and forecasting. However, it is still interesting to compare the findings in this thesis to conclusions made in the literature on econometric models and ARIMA models fitted to exchange rate time series.

In section 3 and 5, an ARIMA(1,1,0) model and a VAR(1) model with additional macroeconomic variables were fitted to the EUR/USD real exchange rate data. Both of these models were concluded to make better predictions than a random walk after the comparison of root mean square error values as a measure of forecast accuracy in section 7. This finding contradicts the conclusion in the paper of Meese and Rogoff (1983). The authors found that the random walk model was superior to an ARMA model. However, a possible explanation of the diverging conclusion is that Meese and Rogoff (1983) did not incorporate the fact that the exchange rate series probably was non-stationary in levels. Though, after examining equation (25) we realise that the sum of the AR coefficients are 1 and thus, unless the past values have changed remarkably during the last two periods, we see that our ARIMA(1,1,0) has similarities to a random walk.

The VAR(1) model in section 5 was derived with additional macroeconomic variables for the real interest rate differential and the trade balance in the EU. The parameter estimates for the lagged additional variables were proven to be insignificant. However, economic theory suggests that there are several important exchange rate determinants and other variables may explain the exchange rate fluctuations in a more sufficient manner. The findings of Yu (2001) and Mida (2013) diverges in the short- and long run. Yu (2001) concludes that neither a VAR, restricted VAR, VECM or a Bayesian VAR generates better forecasts than a random walk in the short run but that the models have better forecast accuracy in the long run. Meanwhile, Mida (2013) find the VAR model to outperform the random walk in the short run but not in the long run. Sellin (2007) also finds a VECM model for the Swedish Krona's real and nominal effective exchange rate to make accurate forecasts once the model has been augmented with an interest rate differential. The diverging results are most likely due to differing exchange rates and the usage of different macroeconomic variables. Thus, finding significant explanatory variables to an exchange rate is troublesome since insignificant variables for one exchange rate may contain valuable information in explaining another.

Lastly, we have used monthly time series data for the period from January 1999 to February 2016 in this thesis. Since the models are based on explaining the past, they will be biased toward the past in the sense that they will weigh historical information more heavily than newer information. This will usually lead to poor prediction performance and is a problem when deriving models that aim to forecast. Thus, using fewer variables and lags in a model are usually beneficial in a forecasting perspective since overfit models often have small in-sample errors, but not lead to favourable out-of-sample forecasts. However, our ARIMA only have one AR-term and our VAR is made up by solely three first-lagged variables. So the problem of overfitting is not applied here. Though the choice of other macroeconomic variables in our VAR may have yielded more accurate out-of-sample forecasts.

Additionally, new information is incorporated quickly on foreign exchange markets due to its easy access for market participants. Market forces tend to adjust the market to a new equilibrium within a short time frame, often faster than a monthly or even a weekly frequency. Therefore, using monthly data do not allow one to quantify how the foreign exchange market react to some new information, e.g. a change in interest rates or a change in price level, since the change has already occurred and been consumed by the time you predict it. Thus, daily data is probably more reliable when predicting exchange rate movements.

8 Conclusions

The main purpose of this thesis was to find a model that makes accurate predictions of the real EUR/USD exchange rate for the three months succeeding February 2016. It was first concluded that the real exchange rate was non-stationary in levels but stationary in first difference after examining the ADF-test presented in table 1 and the ACF and PACF plots in figure 2. Thereafter, we presented three different ARIMA models based on a built-in algorithm for automatic ARIMA modelling in `R` as well as from examinations of the ACF and PACF plots in section 4.1. The candidates were an ARIMA(1,1,1), SARIMA$(2,1,2)(1,0,0)_{12}$ and an ARIMA(1,1,0) model. Based on values of the Bayesian information criterion (BIC) as well as no significant 12th lag in the PACF plot in figure 2, the SARIMA$(2,1,2)(1,0,0)_{12}$ was considered the least appropriate model and the two remaining models were further analysed. The ARIMA(1,1,1) and ARIMA(1,1,0) displayed similar one-step-in-sample forecasts in figure 4 and both models' residuals were presumably from a normal distribution based on quantile-quantile plots displayed in figure 3. As the confidence interval for the parameters in the ARIMA(1,1,1) model included 0.0, and thus were insignificant, we reached to the conclusion that the ARIMA(1,1,0) was the most appropriate one of the three models we had begun our univariate analysis with.

Thereafter, in section 5, an economic model with trade balance in the EU as well as the interest rate differential as additional variables was estimated. Since all the variables were stationary in their first difference according to the ADF-test in table 5, there was a possibility that there existed a linear combination of the levels series that were stationary, and thus a long run relationship between the variables. This kind of cointegrating relationship was tested for by two different Johansen test statistics, both rejecting such hypothesis as observed in table 8 and 9. Thus, a VAR model with the first differenced series was estimated. It was concluded that a VAR model with variables of first lag was selected by all model selection criteria and that the residuals with the real exchange rate as dependent variable in the VAR(1) model demonstrated residuals with desired properties; no autocorrelation and white noise resemblance. However, an outlier in the quantile-quantile plot in figure 8 was observed, but since the normality assumption was not violated, a dummy variable was not added to the model.

When examining equation (26) of the real exchange rate, it was remarked that the equation resembled the one from the estimated ARIMA(1,1,0) in equation (25). As a last part of the analyze of the VAR(1) model, the real exchange rate impulse response to exogenous shocks to the other macroeconomic variables were presented in figure 9 and 10. The conclusions were that the response of the RER variable was in line with theory; an increase in the RIRD caused an appreciation, whereas an increase in the TB caused a depreciation.

Although economic theory suggests that other macroeconomic variables improve the explanation of exchange rate fluctuations, it is observed in section 7 that a simple ARIMA(1,1,0) model gives comparatively good one-step-forecast and even outperforms the VAR model during a 3-month time horizon, when comparing RMSE with actual observed forecast values from ECB in table 10. Furthermore, the random walk is outperformed by both models during both 1-month and 3-month horizon. This suggests that the most important variable to explain the real exchange rate is the lagged variable itself, which also is in line with the parameter estimations in the VAR model in section 5 where only the first lagged RER variable was significant. Thus, this thesis concludes the ARIMA(1,1,0) with its simple interpretation captures dependencies in the data in a relatively sufficient manner.

9 Further research

Since the models are derived to explain the past, new information is not incorporated in the model. Thus, one possible extension to this thesis is to use rolling window forecasts where the parameters are re-estimated after each step in which we includes a new true observation. Furthermore, it would be interesting to include a dummy variable in the model since we observed an outlier in our quantile-quantile plots for the residuals in both the ARIMA(1,1,0) model and the VAR(1) model. Lastly, the inclusion of other macroeconomic variables may yield a different result where the parameters estimates are significant.

References

[1] Brockwell, Peter J. and Davis, Richard J. 2002. Second edition. *Introduction to time series and forecasting.* New York: Springer-Verlag New York Inc.

[2] Hyndman, Rob J. and Athanasopoulos, George. 2014. *Forecasting: principles and practice.* https://www.otexts.org/fpp/ (Downloaded 2017-03-28)

[3] Cryer, Jonathan D. and Chan, Kung-Sik. 2008. *Time series analysis with applications in R.* Second edition. New York: Springer Science+Business Media, LCC.

[4] Tsay, Ruey S. 2010. *Analysis of financial time series.* Third edition. Hoboken, New Jersey: John Wiley Sons, Inc.

[5] Hyndman, Robert J. and Yeasmin Khandakar 2008. Automatic time series forecasting: the forecast package for R. *Journal of statistical software* 27 (3). http://robjhyndman.com/papers/automatic-forecasting/

[6] Blanchard, Olivier Amighini, Alessia and Giavazzi, Fransesco. 2013. Second edition. *Macroeconomics a European perspective.* Harlow: Pearson Education Limited.

[7] Hurvich, Clifford M. and Tsai, Chih-Ling. 1989. Regression and time series model selection in small samples. *Biometrika* 76, 297-307.

[8] Canova, Fabio 2007. *Methods for Applied Macroeconomic Research.* Princeton University Press

[9] Lütkepohl, Helmut Krätzig, Markus and Phillips, Peter C.B. 2004. *Applied time series econometrics.* Cambridge University Press.

[10] Pfaff, Bernhard. 2008. *Analysis of integrated and cointegrated time series with R.* New York: Springer Science+Business Media, LLC.

[11] Hoffmann, Mathias and MacDonald, Ronald 2003. A reeximination of the link between real exchange rates and real interest rate differentials. *CESifo* 894.

[12] King, Michael R. and Rime, Dagfinn. 2010. The 4$ trillion question: what explains FX growth since the 2007 survey?. *Bank for International Settlements* December 2010.

[13] Meese, Richard A. and Rogoff, Kenneth. 1983. Empirical exchange rate models of the seventies - Do they fit out of sample? *Journal of International Economics* 14, 3-24.

[14] Ayekple, Yao E. Harris, Emmanuel Frempong, Nana K. and Amevialor, Joshua.. 2015. Time series analysis of the exchange rate of the Ghanaian Cedi to the American dollar. *Journal of Mathematics Research* 7(3), 46-53.

[15] Akincilar, Aykan Temiz, Izzettin and Şahin Erol. 2011. An application of exchange rate forecasting in Turkey. *Gazi University Journal of Science* 24(4), 817-828.

[16] Weisang, Guillaume and Awazu, Yukika. 2014. Vagaries of the Euro: an Introduction to ARIMA Modeling. *Case Studies In Business, Industry And Government Statistics* 2(1), 45-55.

[17] Yu, Yongtao. 2011. *Exchange rate forecasting model comparison: A case study in North Europe*. Master thesis. Uppsala University. http://www.diva-portal.se/smash/get/diva2:422759/FULLTEXT01.pdf. (Downloaded 2017-04-22)

[18] Mida, Jaroslav. 2011. *Forecasting exchange rates: a VAR analysis.* Bachelor thesis. Charles University in Prague.

[19] Burke, Simon P. and Hunter, John. 2005. *Modelling nonstationary time series: a multivariate approach.* Hampshire and New York: Palgrave MacMillan.

Appendices

A Functions in `R`

A.1 The `auto.arima` function

The Hyndman-Khandakar algorithm for automatic ARIMA modelling follows these steps (Hyndman and Athanasopoulos, 2014, section 8.7):

1. The number of differences d is determined using repeated KPSS tests.
2. After differencing the data d times, the values of p and q are chosen by minimizing the AICc in a stepwise search to traverse the model space

   (a) The model with the smallest AICc is selected from the following four:

   - ARIMA(2,d,2)
   - ARIMA(0,d,0)
   - ARIMA(1,d,0)
   - ARIMA(0,d,1)

   If d=0 then the constant c is included. For $d \geq 1$ c is set to zero. This is called the current model.

   (b) Then variations of the current model are considered:

   - Vary p and/or q from the current model with ± 1;
   - Include/exclude c from the current model.

   The best model after that is either the current model or one of the variations. The best model then becomes the new current model.

   (c) Repeat Step 2(b) until no lower AICc can be found.

The algorithm works in a similar manner when seasonal components are allowed.

B Figures

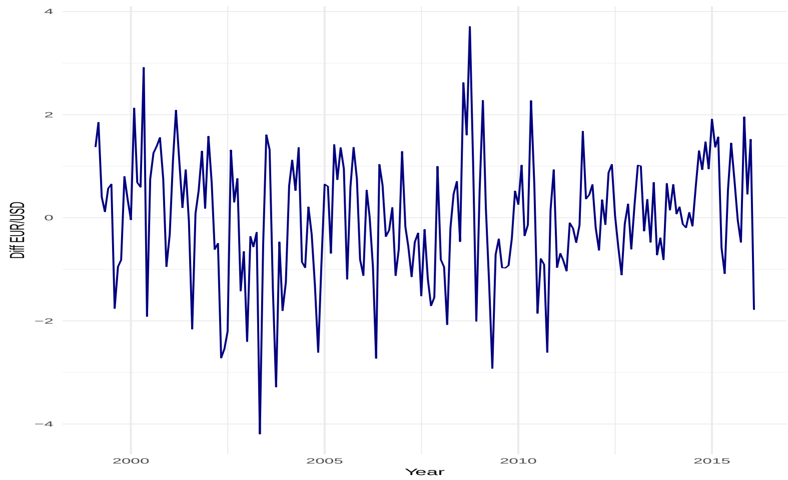

Figure 13: The monthly differenced real euro/us dollar exchange rate

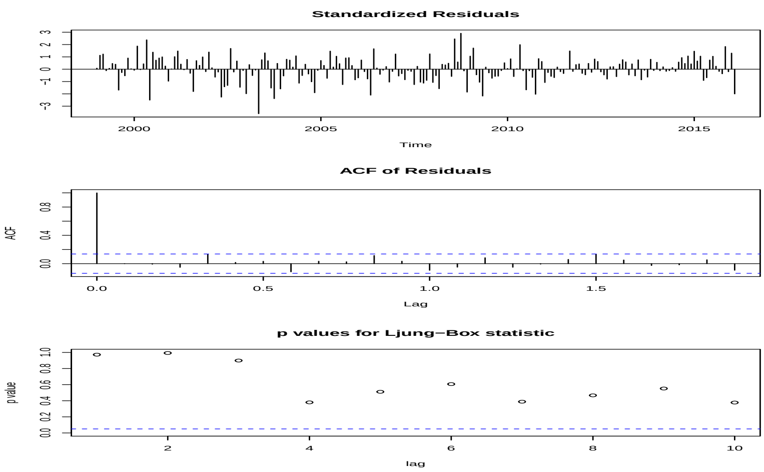

Figure 14: Diagnostic plots of the ARIMA(1,1,1) model

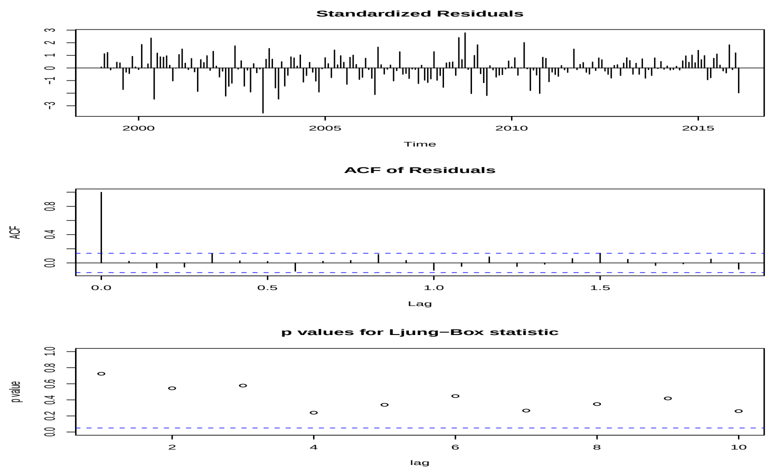

Figure 15: Diagnostic plots of the ARIMA(1,1,0) model

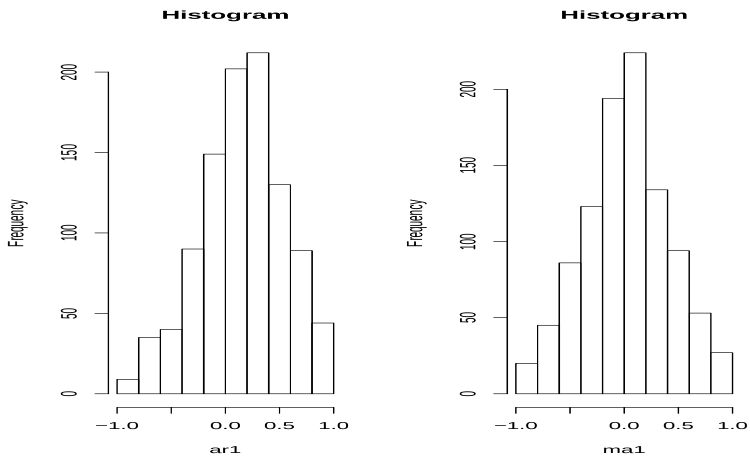

Figure 16: Histograms of bootstrap distribution for the parameter estimates of the ARIMA(1,1,1) model

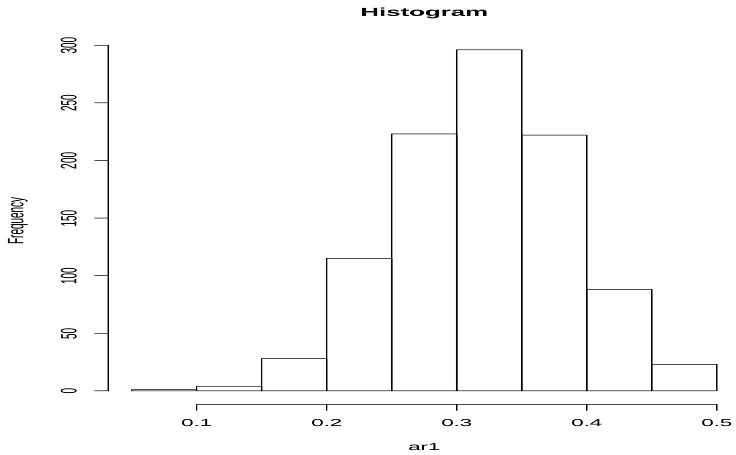

Figure 17: Histograms of bootstrap distribution for the parameter estimates of the ARIMA(1,1,0) model

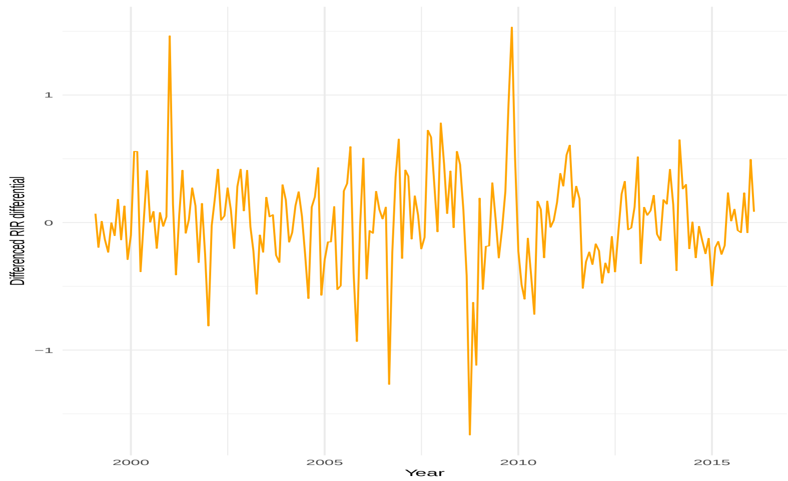

Figure 18: Time series plot of the differenced real interest rate differential

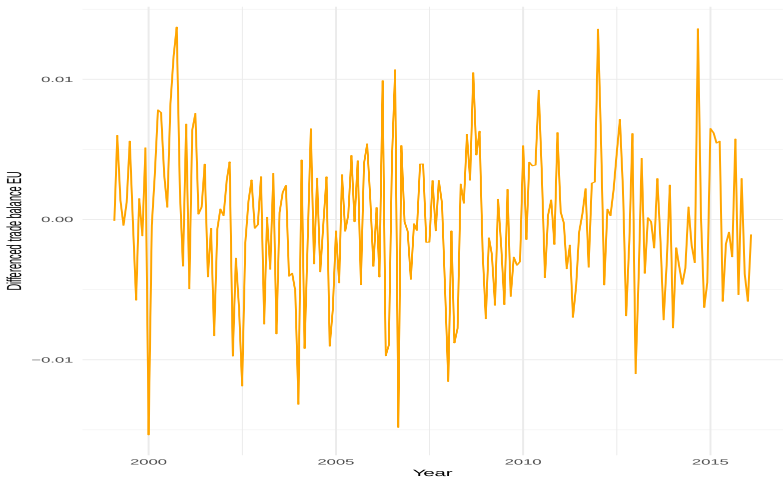

Figure 19: Time series plot of the differenced trade balance in the Euro zone

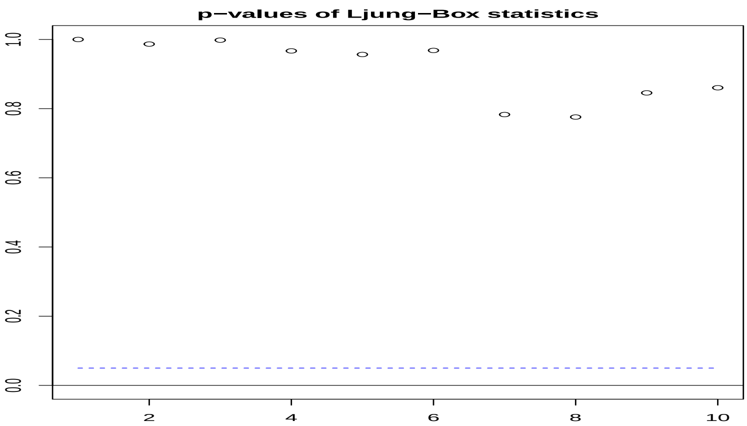

Figure 20: P-values of the multivariate Ljung-Box Portmanteau test

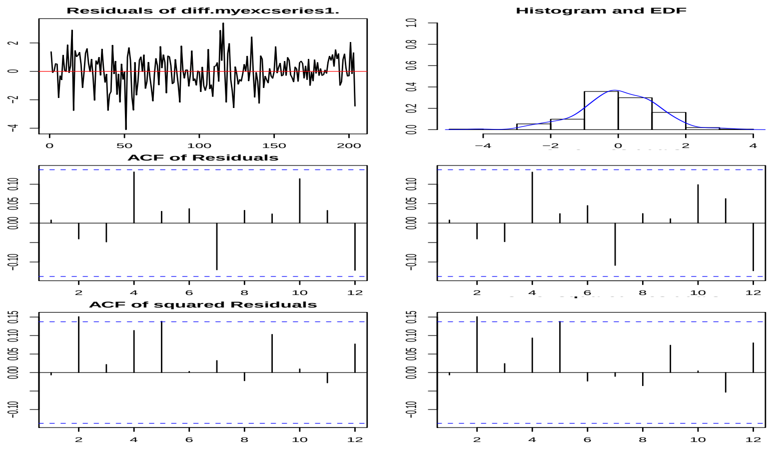

Figure 21: Diagnostics test of the RER equation in the VAR(1) model representing the residuals, estimated distribution function, acf and pacf of the original as well as the squared residuals

C Tables

VAR model selection criteria						
Criteria	1	2	3	4	5	6
AIC(n)	-12.42	-12.36	-12.28	-12.25	-12.22	-12.26
HQ(n)	-12.33	-12.22	-12.08	-11.98	-11.89	-11.77
SC(n)	-12.21	-12.01	-11.78	-11.59	-11.41	-11.20
FPE(n)	4.05 e-06	4.28 e-06	4.66 e-0.6	4.79 e-06	4.94 e-06	5.26 e-06

Table 11: Approximated model selection criteria for the VAR with differenced time series variables.

Parameter estimation VAR(1)			
Equation number	**1**	**2**	**3**
Variable	**RER**	**RIRD**	**TB**
$\mathbf{RER}_{t-1}$	0.34768	-0.03878	0.0019242
Standard error	0.07148	0.02275	0.0002912
P-value	2.32e-06***	0.089836 .	3.42e-10***
t-value	4.864	-1.704	6.608
$\mathbf{RIRD}_{t-1}$	0.07769	0.25927	-0.0007509
Standard error	0.21660	0.06895	0.0008824
P-value	0.720	0.000222***	0.396
t-value	0.359	3.760	-0.851
$\mathbf{TB}_{t-1}$	-19.66126	-7.52898	-0.0328862
Standard error	16.06442	5.11362	0.0654421
P-value	0.222	0.142493	0.616
t-value	-1.224	-1.472	-0.503

Table 12: Estimated coefficients for the first lag with corresponding values of standard error, P-value and t-value

Out of sample forecasts		
Model	ARIMA(1,1,0)	VAR(1)
1M	96.1616	96.1367
2M	95.9806	95.0066
3M	95.9229	95.9887

Observed values	
1M	95.3314
2M	93.8515
3M	94.4420

Table 13: Three out of sample forecasts for the real exchange rate from VAR(1) and ARIMA(1,1,0) in the left tabular. The right tabular displays the real observed valus for March, April and May 2016 with data from ECB.