Stockholms universitet

Is the European Central Bank Handling the stress?

Modi Lovenvall

Kandidatuppsats i matematisk statistik
Bachelor Thesis in Mathematical Statistics

Kandidatuppsats 2015:22
Matematisk statistik
September 2015

www.math.su.se

Matematisk statistik
Matematiska institutionen
Stockholms universitet
106 91 Stockholm

Stockholm University

Mathematical Statistics
Stockholm University
Bachelor Thesis **2015:22**
http://www.math.su.se

Is the European Bank Handling the Stress

Modi Lovenvall*

September 2015

Abstract

This rapports aim is to investiage the European Central Banks way of using mathematical statistical processes to affect the European economy. In order to accomplish this, we explain what methods are being used, how the are working, and what purpose they serve the European Central Bank. The models of concern are the Vector Autoregressive Model and other time series models. We use these models to simulate forecasts, using data from two time series; these being unemployment rate and house price index. We will further compare our results with official European Central Bank statistics, where we will see that the Vector Autoregressive Model could be a suitable method when fine-tuned. Additionally we will see that the Vector Autoregressive Model has good use for the European Central Bank when they are conducting their macroeconomic scenarios, we will attempt to explain why in this rapport.

*Postal address: Mathematical Statistics, Stockholm University, SE-106 91, Sweden. E-mail: modi.lovenvall@gmail.com. Supervisor: Mathias Lindholm.

ACKNOWLEDGEMENT

I would like to thank my supervisor for this thesis, Mathias Lindholm, for great support and valuable consultation.

CONTENT

1 Introduction

In recent times, as a response to tough historical economic crisis, the European Central Bank (ECB) has made an effort to enforce methodologies and macroeconomic scenarios to bank stress tests. The stress tests are conducted so banks can determine their ability to respond to financial crisis. The scenarios constitute an extreme percentile of a damaging future macroeconomic outcome.

A key element of the stress test is the statistical simulation of the macroeconomic scenarios, which are taking place during the stress test period, the standard stress test period being three years. These tests could include the increase in unemployment or a change upon the interest rates.

The European Systemic Risk Board (ESRB) has the task of overview and care of the macroeconomic risks in the financial sector. Additionally the ESRB plays a large part in the stress test preparation of the economic scenarios, which the banks ultimately use in their stress test evaluation.

One factor that makes the stress test interesting to evaluate is the poor performances across banks in Europe. One could argue that the stress test should be sufficient and prepare banks for any type of economic crisis. But the empirical evidence shows otherwise.

We will in this report examine how the ESRB are preparing the final result for the macroeconomic scenarios used in the stress tests. We are going to study the statistical models that are being used in the simulations, and explain how they work in theory. Then simulate scenarios with these models and compare our results with the official ECB scenarios.

We do this to get a better understanding of the way the ESRB are working, also to investigate if the models that are being used serve their purpose to create these macroeconomic scenarios.

Additionally, to restrict this report, we are focusing on the scenarios regarding unemployment and house price index that are made to accommodate the macroeconomic scenarios for the bank stress test in Sweden.

2 Definition of Terms/Models and Methodology

2.1 Methodology used by ECB

One should understand how the ECB are preparing the framework regarding the stress test scenarios. By knowing this, the conclusion and the final judgment for the results regarding the macroeconomic scenarios will seem just.

Every year the ECB is conducting a new stress test, the work begins with examining what scenarios are appropriate and should be forecasted during the stress test. The procedure of determining the scenarios contains a specific way of thinking regarding risk. The assessments reflect a plausible outcome in the economy that also warrants the probability of an impact on banks.

Naturally there are different types of risk during an economic crisis. So scenarios are designed to match specific risk. When doing this, the ECB explains that different models are used for different kind of scenarios, but they also conclude that some risks could have an effect on another risk, "*a joint shock calibration of the two risks makes more economic sense",* see p. 15 in [13].

The way of producing the scenarios at first hand dose not seem to have any real model or set procedure. Instead the model selection that is being used is "*tailored to the specific risk that the scenarios is supposed to reflect*", and "*the ECB uses an eclectic approach when selection models to produce scenarios*", see p. 15 in [13]

But we have a few models that are being mentioned by the ECB to produce the scenarios, one is the Stress test Elasticities tool (STEs Tool), another model being used for determining risk and scenarios outside the EU is the NiGEM model, see [14]. Further models that are being used are DSGE model, MSM model and also various types of VAR (vector autoregression) models (GVAR or BVAR), see p. 18 in [13].

Unfortunately the STE Tool and NiGEM model are not official to the public, so we would have to devote our attention on the latter mentioned models.

Since there are no clear or set rules for the for the final macroeconomic scenario result, it is hard to pinpoint where the calculations made by the ECB could be assumed too be deficient. Therefore it is a wise choice to focus on the scenario results that the ECB present for the stress test and make sure that they seem valid and that they have a statistical background that we could investigate, see e.g. [13][14].

2.2 Further detailed Explanations of Models

The focus of this report is the macroeconomic scenarios within the EU, more specific Sweden, which leaves us with the STE-tool and the VAR models that are being used along with the ad-hoc way of working at the ECB.

We know that the STE tool is a “multi-country EU-wide shock simulation tool based on impulse response functions of endogenous variables to pre-defined exogenous shocks”, see [14].

The information about the STE tool let us make an assumption about what kind of statistical analysis that is being used. We know that the use of historical data is essential when making future assumptions. Further one must have the knowledge of how this data is evolving over time and also how it is reacting to exogenous shocks to be able to create a simulation tool like the STE. This information gives us a hint of the use of time series analysis. To make this assumption more understandable, we will further ahead in the report talk about the theory of time series analysis.

However, time series analysis cannot be the only statistical tool that we should use. Since we are aware of that multiple variables can intertwine and affect each other, and that variables could have a joint effect on the macroeconomic scenarios. We must take this in consideration and also capture these affects.

One-way of doing this is the use of vector autoregression model (VAR), as is mentioned by the ECB. This is a tool that helps us understand the time series dependence on each other, and help with creating good forecast predictions, this is only mentioning a few things the VAR can help us with. Further in the report we will talk about VAR theory, with the focus being on how the ECB might use the VAR, how the procedure might look like, and the most likely way that the final ECB results are put together.

3 Methods and Theory

3.1 Methods

The use of univariate time series and multivariate vector autoregression will be the method of choice when conduction our work. Our reason to include univariate time series analysis is to get a better understanding of the time series process; also to have a second forecast result to compare with the vector autoregression forecast.

The two statistical analysis tools will help us produce our own scenarios, and compare results with the ECB final macroeconomic scenario. We will gather information from historical time series regarding unemployment rates and house prices index. Additionally we like to point out that making prefect forecast prediction with the limits of this report might to not be the easiest task, but since we are researching ECBs way of conducting their scenarios, and how the statistical models works, the end scenario result will serve its purpose.

With the support of the program R, we are going to create forecasts and produce scenarios from the times series models and the VAR models. For reference we are mainly using three packages in R, those being "tseries", "vars" and "forecast", these will suffice when conduction our work in R.

3.2 Theory

3.2.1 Time Series Analysis

With the help of the book "Time Series Analysis and Its Applications by Robert H. Shumway & David S. Stoffer (2006)", we can start to understand the process of time series and also explain and go through how we could use time series analysis for our purpose.

Time series analysis is the analysis of data that has been observed in different points in time. One could use this analysis tool to investigate trends over time or make forecast for the future. We can imagine that there is a clear correlation within the sampling of adjacent points of the same subject over time. This would mean that other conventional statistical tools would be restricted since they are dependent on the assumption of independent and identical distributed (i.i.d) observations, see p. 1 in [2].

When using time series analysis there are two approaches that one should be familiar with before analyzing the obtained data. Those approaches are Time-Domain and Frequency-Domain.

Time Domain is focused on analyzing data over a time period. This could be different macroeconomic data of different kinds over a time period. Or observations made on different occasions, whether it be the climate readings or the stock market. While frequency domain analysis refers to signals or mathematical functions based on frequency rather then time, see [1].

Which in our case, with the macroeconomic data we obtained over a time period, where the data is not connected to any apparent mathematical function or any type of frequency signal. Leads us to use the time series analysis with time domain approach.

3.2.1.1 Time Domain Approach

The time domain approach focuses on modeling future values of a time series as a parametric function of the current and past values. This approach is generally motivated by the presumption that correlation between adjacent points in time is best explained in terms of a dependence of the current value on past values, see p. 2 in [2]. Thus the results of the time domain approach can be used as a forecasting tool.

The approach with time domain is to begin with a linear regression of the present value of a time series on its own past value, and on the past values of other series, see p. 2 in [2]. With the linear regression we are able to estimate parameters with maximum likelihood methods and investigate the significance for them in the model.

Further when one realizes and understands how essential and important correlation is as a feature to time series analysis. It will lead us to make further analysis of our data with the aid of the auto-correlation function.

3.2.1.2 Auto-Correlation function

The auto-correlation function helps us measure the linear predictability of the time series. The auto-correlation function is defined as

$$\rho(s,t) = \frac{\gamma(s,t)}{\sqrt{\gamma(s,s)\gamma(t,t)}}$$

Where *s,t* are particular time values in the time series, and where $\gamma(s,t) = \gamma_y(s,t) = \mathrm{E}[(y_s - u_s)(y_t - u_t)]$ is the covariance, with $u_t = \mathrm{E}[Y_t]$ and $u_s = \mathrm{E}[Y_s]$.

This is a tool for clarifying relations that may occur within and between time series at various lags, see p. 84 in [2], in other words, the time period between two observations.

If we can predict y_t perfectly from y_s through a linear relationship, $y_t = \beta_0 + \beta_1 y_s$. Then the correlation will be 1 if $\beta_0 > 0$, i.e. showing a direct increasing linear relationship. And -1 if $\beta_1 < 0$, i.e. showing a decreasing linear relationship.

Hence with the correlation, giving us an estimate of how strong the linear relationship is, we have a rough estimate of the ability to forecast the series at time *t* from the value at time *s,* see p. 84 in [2]. Since we are trying to simulate a good forecast on unemployment rates and house price index to measure against the ECB stress test scenarios, the ACF is an important piece in our modeling.

Further analysis of the time series data is to make sure that all patterns in the time series are accounted for. The next step in modeling the series is estimation of the number of terms in the model that describes the dependency among successive observations best (autoregressive terms). Also looking at the number of terms in the model that describes the dependency among successive observations best (moving average terms), see p. 2 and 4 in [3].

To figure out the number of terms, one could use the moving average model and the autoregressive model, these is the two models will be helpful for this this purpose.

3.3 Autoregressive Model

Autoregressive models are based on the idea that the current value of the series y_t, can be explained as a function of p past values, $y_{t-1}, y_{t-2}, \dots,\ y_{t-p}$, where p determines the number of lags into the past needed to forecast the current value.

The definition of an autoregressive model of order p, or AR (*p),* has the following form

$$y_t = c + \phi_1 y_{t-1} + \phi_2 y_{t-2} + \cdots + \phi_p y_{t-p} + \varepsilon_t, \text{ or}$$

$$y_t = c + \Sigma_{i=1}^{p} \phi_i y_{t-i} + \varepsilon_t,$$

where y_t is stationary, meaning that the mean of the function is constant and does not depend on time, and that the covariance function $\gamma_y(s,t)$ depends on s and t only through their difference |s - t|. Further c is a constant, $\phi_1, \dots, \phi_p$ ($\phi_p \neq 0$) are the parameters of the model and ε_t is the error term, also called the white noise of the model, we assume that ε_t has a mean of zero and variance σ_ε^2, see p. 85 in [2] and see p. 24 in [2].

Under the assumption of a known order of p we have a possibility of estimating the parameters in the model, see p. 56 in [7].

3.3.1 Estimation of the Autoregressive Processes

If we know the distribution of the white noise that generates the autoregressive p process, AR (p) process, the parameters can be estimated using the maximum likelihood (ML) method. We are going to restrict this report to ML methods, since this is what we are using as an estimator method in our analysis when using our programing tool in R.

We would also like to point out that the ML of an autoregressive process might be a bit dense and complex, so to get an idea of how the ML method is working we are going to focus on an AR (1) process. We use the equation

$$y_t = c + \phi(y_{t-1}) + \varepsilon_t,$$

where we replace $c = \mu(1-\phi)$ and get $y_t = \mu + \phi(y_{t-1} - \mu) + \varepsilon_t$,

here $|\phi| < 1$ and $\varepsilon_t \sim iid\ N(0, \sigma_\varepsilon^2)$. Given data $y_1, y_2, \dots, y_n$ we seek the likelihood

$$L(\mu, \phi, \sigma_\varepsilon^2) = f_{\mu,\pi,\sigma_\varepsilon^2}(y_1, y_2, \dots, y_n).$$

Where $f_{\mu,\pi,\sigma_\varepsilon^2}(\cdot)$ is the density for $\mu, \pi, \sigma_\varepsilon^2$, and for future reference will be dropped to $f(\cdot)$, to ease the notation.

In the case of AR(1), we may write the likelihood as

$$L(\mu, \phi, \sigma_\varepsilon^2) = f(y_1) f(y_2|y_1) \dots f(y_n|y_{n-1}),$$

Because $y_n|y_{n-1} \sim N(\mu + \phi(y_{t-1} - \mu), \sigma_\varepsilon^2)$ we have

$$f(y_n|y_{n-1}) = f_\varepsilon[(y_t - \mu) - \phi(y_{t-1} - \mu)],$$

where $f_\varepsilon(\cdot)$ is the density of ε_t, that is, the normal density with mean zero and variance σ_ε^2. We may write the likelihood as

$$L(\mu, \phi, \sigma_\varepsilon^2) = f(y_1)\Pi_{t=2}^n f_\varepsilon[(y_t - \mu) - \phi(y_{t-1} - \mu)].$$

To find $f\ (y_1)$, we can use the representation

$$y_1 = \mu + \Sigma_{i=0}^\infty \phi^i \varepsilon_{1-i}$$

to see that y_1 is normal, with the mean μ and variance $\sigma_\varepsilon^2/(1-\phi^2)$. Finally for the AR(1), the likelihood is

$$L(\mu, \phi, \sigma_\varepsilon^2) = (2\pi\sigma_\varepsilon^2)^{-\frac{n}{2}}(1-\phi^2)^{\frac{1}{2}} \exp\left[-\frac{S(\mu,\phi)}{2\sigma_\varepsilon^2}\right],$$

where typically $S(\mu, \phi)$ is called the unconditional sum of squares,

$$S(\mu,\phi) = (1-\phi^2)(y_1-\mu)^2 + \Sigma_{t=2}^n[(y_t - \mu) - \phi(y_{t-1} - \mu)]^2,$$

see p. 126 in [2].

To maximize the likelihood model above would be hard and require iterative or numerical procedures. An alternative, and in the case of AR models, we have the advantage to conditioning on initial values, and set linear models. That is we can drop the term in the likelihood that causes the nonlinearity. Conditioning on y_1, the first observation, we can maximize the likelihood on the first observation. Regarding y_1 as deterministic and setting $f(y_1) = 1$, the conditional likelihood becomes

$$\begin{aligned} L(\mu, \phi, \sigma_\varepsilon^2|y_1) &= \Pi_{t=2}^n f_\varepsilon[(y_t - \mu) - \phi(y_{t-1} - \mu)] \\ &= (2\pi\sigma_\varepsilon^2)^{-(n-1)/2} \exp\left[-\frac{S_c(\mu,\phi)}{2\sigma_\varepsilon^2}\right] \end{aligned}$$

where the conditional sum of squares is

$$S_c(\mu, \phi) = \Sigma_{t=2}^{n}[(y_t - \mu) - \phi(y_{t-1} - \mu)]^2$$

, see p. 126 in [2] and [16]

Taking the partial derivative of the log of the conditional likelihood function with respect to σ_ε^2 and setting the result equal to zero, we see that for any given values of μ and ϕ in the parameter space, $\hat{\sigma}_\varepsilon^2 = S_c(\hat{\mu}, \hat{\phi})/(n-1)$ maximizes the likelihood. Thus the conditional maximum likelihood estimate of σ_ε^2 is

$$\hat{\sigma}_\varepsilon^2 = S_c(\hat{\mu}, \hat{\phi})/(n-1),$$

so $\hat{\mu}$ and $\hat{\phi}$ are the values that minimize the conditional sum of squares, $S_c(\hat{\mu}, \hat{\phi})$. Letting $\alpha = \mu(1 - \phi)$, the conditional sum of squares can be written as

$$S_c(\mu, \phi) = \Sigma_{t=2}^{n}[y_t - \alpha + \phi(y_{t-1})]^2.$$

From a least square estimation, we have $\hat{\alpha} = \bar{y}_{(2)} - \hat{\phi}\bar{y}_{(1)}$, where $\bar{y}_{(1)} = (n-1)^{-1}\sum_{t=1}^{n-1} y_t$, and $\bar{y}_{(2)} = (n-1)^{-1}\sum_{t=2}^{n-1} y_t$, and the conditional estimates are then

$$\hat{\mu} = \frac{\bar{y}_{(2)} - \hat{\phi}\bar{y}_{(1)}}{1 - \hat{\phi}}$$

$$\hat{\phi} = \frac{\sum_{t=2}^{n}(y_t - \bar{y}_{(2)})(y_{t-1} - \bar{y}_{(1)})}{\sum_{t=2}^{n}(y_{t-1} - \bar{y}_{(1)})^2},$$

see p. 127 in [2].

We are now going to restrict the report and leave the MLE of parameters of the data in its entirety to the program R.

With our now improved understanding for the autoregressive process. We will apply this in our analysis, and find a good model for our data to be able to create forecasts. A statistical model that will help us with this is the Autoregressive-Integrated-Moving-Average model (ARIMA-model).

This model combines the autoregressive model and the moving average model. We are not going to further investigate the moving average model, since our main focus is going to be on multivariate autoregression technique further in the report.

3.4 Autoregressive-Integrated-Moving-Average model (ARIMA)

There are a few components to the ARIMA *(p, d, q)* model that we should be aware of. First we have the auto-regressive element, *p*, which shows the number of terms in the model that describes the dependency among successive observations. Then we have the differencing element, *d*, it stands for how many times the data has been differenced, that is computing the differences between consecutive observations. Lastly the moving average element, q, it describes the persistence of a random shock from one observation to the next, see p. 1 in [3].

The middle element, d, is investigated before p and q. The goal is to determine if the process is stationary and, if not, to make it stationary before determining the values of p and q. We do this by calculating the differences between pairs of observations at some lag. If $d > 0$, then we would have to difference the data to make it stationary, and remove any seasonality, that is removing recurring patterns in data. An example being if $d = 1$, we would have a differencing process of $c_t = y_t - y_{t-1}$, where c_t is the first differenced series of y_t, see [17]. Now if $d = 0$, it would mean that the data is already stationary and no number of differencing is needed. Recall that a stationary process has a constant mean and variance over the time period, see p. 8 in [3].

The autoregressive component p represents the memory of the process for preceding observations. If the value of $p = 0$, there is no relationship between the closest observations. If the value is 1 there is a relationship and a correlation at the time periods between two following observations (lag 1), likewise $p = 2$ indicates a lag of 2, see p. 12 in [3].

Lastly the moving average component q represents the memory of the process for preceding random shocks. When $q = 0$, there are no moving average components. And similar to the autoregressive competent when the value is 1 there is a relationship between the current score and the random shock and a correlation in lag 1, and likewise $q = 2$ also has a relationship and correlation on lag 2, see p. 12 in [3].

To define the ARIMA model we can start by defining the ARMA part. A time series is an ARMA (p, q) if it is stationary and

$$y_t = \phi_1 y_{t-1} + \cdots + \phi_p y_{t-p} + w_t + \theta_1 w_{t-1} + \cdots + \theta_q w_{t-q},$$

with $\phi_p \neq 0, \theta_q \neq 0$ and $\sigma_w^2 > 0$. Where w_t is assumed to be white noise, and $\phi_1, \dots, \phi_p$ are the parameters of the autoregressive components and where $\theta_1, \dots, \theta_q$ are the parameters of the moving average component. So if the model would only contain one order of autoregressive component and one component from the moving average for example, we would have an ARMA (1,1) model, see p. 93 in [2].

Further the integrated ARMA model, or ARIMA model, is a broadening of the ARMA model to include differencing. This is needed, as we previously discussed, when the time series is not stationary and needs to be differenced. An example of a notation of an ARIMA model if a time series where to be differenced two times, with a model containing one order of autoregressive component and one component from the moving average, we would have an ARIMA (1,2,1), see p. 141 in [2].

One of the first steps to estimate the parameters of the ARIMA with MLE would be to difference the data, thus making it stationary. We have seen that the MLE is very dense, and that having stationary data is a benefit when maximizing the likelihood. It is difficult to write the likelihood as an explicit function of the parameters in an ARIMA model, instead, it could be advantageous to write the likelihood in terms of one-step prediction errors, $x_t - x_t^{t-1}$, see p. 128 in [2]. Luckily R will handle the MLE for the ARIMA, so we will not go into details, as this matter is very complex.

Furthermore, if our time series would show signs of seasonality in the data, that is, signs of recurring events. We would need to handle this in a form of a broadened ARIMA model, for example the seasonal autoregressive intergraded moving average model, the SARIMA model.

3.4.1 Seasonal-ARIMA models (SARIMA)

If our time series show sings of seasonal data, the best-fitted model would be a seasonal autoregressive integrated moving average model. This is to take account for seasonal and non-stationary behavior in our data. The dependence on the past could occur when we have underlying seasonal lag S, i.e. a recurring seasonal trend, this might be a quarterly or a yearly trend for example, see p. 157 in [2].

The SARIMA model adds seasonal autoregressive and seasonal moving average components to the ARIMA model. A Seasonal ARMA model would be denoted as; ARMA (*p, q*)(*P, Q)(S)*, and is given by

$$y_t = \phi_1 y_{t-1} + \cdots + \phi_p y_{t-p} + w_t + \theta_1 w_{t-1} + \cdots + \theta_q w_{t-q} + \Phi_1 y_{t-s} + \cdots + \Phi_p y_{t-s} \\ + \Theta_1 w_{t-s} + \cdots + \Theta_q w_{t-s} \,,$$

where $\phi_p \neq 0, \theta_q \neq 0, |\Theta_p| < 1, |\Phi_q| < 1$ and σ_w^2. The notation Θ_p is the seasonal autoregressive component and Φ_q is the seasonal moving average component with seasonal period of *s*.

Identical to the ARMA model, the SARMA model can be broadened to the integrated SARMA model, SARIMA model, i.e. including differencing of the time series if needed, see p. 155 in [2].

As an example we denote an ARIMA (*p, d, q*)(*P, D, Q)(S)*, which would mean that we would have an SARIMA model. Where the ordinary autoregression, integrated and moving average components are represented by *p, d* and *q*. Additionally, the seasonal- autoregression, integrated and moving average components, are represented by *P, D* and *Q.* Lastly *S*, which is representing the seasonal fluctuations, see p. 159 in [2].

Additionally drift could occur and can be added to the SARIMA model, this drift appears when the data is non-stationary and is differenced, it is the mean that is set for the specific produced SARIMA model, see p. 127 in [2].

When we want to select the appropriate model for our data, we will use the help of R. The process would be to first find the difference operators *d*, then continue and produce a roughly stationary series, set the autoregressive terms *p*, the moving average terms q, and if necessary a seasonal ARIMA model, see p. 159 in [2].

Estimation of the parameters within the SARIMA is hard to do directly. If we take a look at the previously denoted Seasonal ARMA model, it contains two expressions with each of their own lag. This would have to be telescoped out before starting any ML estimation, meaning we would have to expand the expressions to be able to handle them, see [19]. When this is done, we face the same difficulties as we do with the ARIMA model as we motioned in section *3.4*. We will again leave further theory of this complex matter for another paper.

The ARIMA and SARIMA analysis will be preformed on our time series data with the help of the program R. This will help us to fit the best model for the data over the investigated time period. Also we will leave it to R for the calculation of the Maximum Likelihood Estimation (MLE) of parameters in the fitted model.

Once this is done we will make sure that the model is good by focusing on the residuals of the model.

3.4.2 ARIMA/SARIMA Modell Checking

To investigate how well the fit is of the produced model, we will look at different plots to look for general assumptions regarding regression. We are going to investigate that our residuals follow the normality assumption, which would be normally distributed residuals in the model. Also we will investigate that our residuals are identical with mean zero and equal variance, i.e. that the residual are homoscedastic, oppose to not following those assumptions being heteroscedastic residuals. And lastly that we previously discussed, have a look at the ACF to investigate the correlation of residuals between lags, see p. 77 in [9].

Remember that our aim is to investigate how good we are able to forecast future values from the time series we have with the univariate time series method, so we can later expand to multivariate time series modeling and forecasting.

The above theories will help us find a model that best fit our data and that will allow us to make a fairly good forecast prediction. The forecast will be done by the aid of the program R.

3.5 Time Series Forecasting

When we are doing the forecasts, we are grounding the future on the historical data and on univariate time series. For the sole purpose of predicting the future this seems to be a good way. But to analyze if our forecasts are good predictions, before we move on to multivariate time series, we can consider looking at the predictions errors for the produced forecast.

There are a few way of analyzing the prediction errors, but we choose to look at the Mean Absolute Error (MAE). The MAE could be explained as a measure of the average scale of the errors in the forecast, since it is an absolute value it does not consider the direction of the error. It is given by $MAE = \frac{\sum_{i=1}^{n}|f_i - y_i|}{n}$, where f_i is the prediction and y_i true value. Also we want to note that the value of the MAE is not normalized, i.e. the size depends on the scale of data. Further the MAE is easy to understand and compute, but it can be a chore to understand the relative size of the error with MAE. We would have to understand what kind of scales we have in the data we are producing when forecasting, see p. 87 in [7].

Now since one of our ambitions is to compare our forecast with the ECB produced scenarios, the univariate time series forecast will not be sufficient. The methodology given to us by the ECB tells us that we have to take into account that different macroeconomic events can affect one another. For example when there is a crisis in employment the values of house prices could be affected.

To investigate multiple time series simultaneously, we have a look for linear independences among the different time series; to do this we are going to use a multivariate technique. A popular technique is the vector autoregressive model (VAR model), *remember that VAR is used and mentioned in the ECB methodology.*

3.6 Vector Autoregressive Model

VAR models are used for multivariate time series. It is essentially a system of equations where the dependent variable is regressed on lagged observations of all the variables in the procedure. The structure of each variable is a linear function of the time period of past observations from the individual variable, also called lags, that also includes past lags of the other variables concerning the individual variable in the VAR analysis, see [4].

Some benefits of this multivariate method, comparing to the univariate method, is that we can forecast on a collection of relatable variables, testing whether one variable if useful to forecast another and see the proportion of how much one variables variance is attributed with another variable, this is only to name a few benefits, see [18].

The VAR can be explained as an n-equation, n-variable model in which each variable is in turn explained by its own lagged values, plus current and past values of the remaining *n-1* variable, see [5]. The VAR model resembles the autoregressive model, the main difference is that we are dealing with multiple regressions and gathering them into vectors.

To make a clear definition of the VAR model we can have a look at how a VAR model is built, let $\boldsymbol{Y}_t = (y_{1t}, y_{2t}, \dots, y_{nt})'$ denote an $(n \times 1)$ vector of time series variables. Where y_{nt} is the dependent variable and "n" stands for which variable it represents and "t" stands for the observation in time "t". The p-lag VAR (p) has the form

$$\boldsymbol{Y}_t = \boldsymbol{c} + \boldsymbol{\Pi_1}\boldsymbol{Y}_{t-1} + \boldsymbol{\Pi_2}\boldsymbol{Y}_{t-2} + \cdots + \boldsymbol{\Pi_p}\boldsymbol{Y}_{t-p} + \boldsymbol{\varepsilon_t}\,, t = 1, \dots, T$$

where $\boldsymbol{\Pi_i}$ are $(n \times n)$ coefficient matrixes and $\boldsymbol{\varepsilon_t}$ is the error terms for the vector process with time invariant covariance matrix $\boldsymbol{\Sigma}$,

$$E(\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}'_t) = \boldsymbol{\Sigma}_\varepsilon = \begin{bmatrix} \sigma^2_{\varepsilon_1} & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \sigma^2_{\varepsilon_n} \end{bmatrix}$$

This error term should also satisfy $\boldsymbol{E}(\boldsymbol{\varepsilon}_t) = \mathbf{0}$, that every error term should have the mean 0 and $\boldsymbol{E}(\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}'_t) = \boldsymbol{\Sigma}_\varepsilon$ and 0 otherwise. That is, no correlation across time. Thus, a VAR is a system in which each variable is expressed as a function of own lags as well as lags of each of the other variables; see [5][6].

Thus, in a VAR of order p each component of the vector Y depends linearly on its own lagged values up to p periods as well as on the lagged values of all other variables up to order p. With this concept, the VAR has a process that allows the analyst to identify and interpret economic shocks and to assess their influence on macroeconomic variables, see p. 128 in [7]. Which is what we are looking to do, as it is this that we have detected that the ECBs is trying to achieve.

When creating the VAR from our time series, we need to have a process of identifying the best model, with the appropriate lags and best estimation of parameters to the VAR model. We will do this with the aid of R programming, where it will determine the amount of lags, and with that the amount of parameters in the model to best fit our time series data.

The way R handles the data from the time series to fit a VAR model is through a process that resembles stepwise selection, which means that R will try multiple was of fitting the data, i.e. trying different lags to get the optimal fit, and cancelling out the models with least significances. It will also help us estimate the parameters in the model that is produced. That will be conducted in a similar way of estimation parameters in the autoregressive process.

3.6.1 Estimation of parameters in the VAR process

We have covered the simple AR(1) process of estimating parameters, we are now going to show in a more simplistic approach, how the estimation is being done in the VAR(p) case.

First we are going to simplify the pervious VAR (p) notation in a compact form

$$\boldsymbol{Y}_t = \mathbf{B}'\mathbf{Z}_t + \boldsymbol{\varepsilon}_t, t = 1, \dots, \mathrm{T}$$

$$\boldsymbol{\varepsilon}_t \sim i.i.d\ \mathrm{N}(0, \boldsymbol{\Omega})$$

where $\mathbf{B}' = [c_0, \boldsymbol{\Pi}_1, \boldsymbol{\Pi}_2, \dots, \boldsymbol{\Pi}_k]$, $\mathbf{Z}'_t = [1, y'_{t-1}, y'_{t-2}, \dots, y'_{t-k}]$ and the initial values $\mathrm{Y}^0 = [y'_0, y'_{-1}, \dots, y'_{-k+1}]$ are given, also we have the notation $\boldsymbol{\Omega}$ for $\boldsymbol{E}(\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}'_t)$, the covariance matrix of error terms. We need to derive the equations for estimating $\mathbf{B}$ and $\boldsymbol{\Omega}$, which can be done by finding expressions for $\mathbf{B}$ and $\boldsymbol{\Omega}$ for which the first order derivatives of the likelihood function are equal to zero.

Consider the log likelihood function

$$\ln L(\mathbf{B}, \mathbf{\Omega}; \mathbf{Y}) = -\frac{TN}{2}\ln(2\pi) - \frac{T}{2}\ln|\mathbf{\Omega}| - \frac{1}{2}\sum_{t=1}^{T}(y_t - \mathbf{B}'\mathbf{Z}_t)'\Omega^{-1}(y_t - \mathbf{B}'\mathbf{Z}_t),$$

then calculate $\partial \ln L / \partial \mathbf{B} = 0$, which gives

$$\sum_{t=1}^{T} y_t \mathbf{Z}_t' = \widehat{\mathbf{B}}' \sum_{t=1}^{T} \mathbf{Z}_t \mathbf{Z}_t',$$

so the MLE of $\mathbf{B}$ is:

$$\widehat{\mathbf{B}}' = (\sum_{t=1}^{T} y_t \mathbf{Z}_t')(\sum_{t=1}^{T} \mathbf{Z}_t \mathbf{Z}_t')^{-1}.$$

Next we calculate $\partial \ln L / \partial \mathbf{\Omega} = 0$, which gives us the MLE of $\mathbf{\Omega}$

$$\widehat{\mathbf{\Omega}} = T^{-1}\sum_{t=1}^{T}(y_t - \widehat{\mathbf{B}}'\mathbf{Z}_t)(y_t - \widehat{\mathbf{B}}'\mathbf{Z}_t)^{-1}$$

, see p. 56 in [9].

We restrict ourselves and will leave the actual calculation on the data and the ML estimations to R. The next step will be to understand how good of a model we are producing with our data.

3.6.2 Vector Autoregression Model Checking

Once a VAR *(p)* model has been estimated, we will be able to do further analysis. The multivariate normality assumption underlying the VAR model can be checked against the data using the residuals, see p. 66 in [9]. We as researchers should be interested in diagnostic tests, such as testing for the absence of autocorrelation, heteroscedasticity or non-normality in the error process, see [8].

The VAR model describes the variation in $\boldsymbol{Y}_t$ as a function of lagged values of the process, but not of current values. That is, that all information about current effects in the data is contained in the residual covariance matrix $\mathbf{\Sigma}_\varepsilon$, see p. 66 in [9].

3.6.2.1 Test for Residual autocorrelation

The VAR methodology is based on the idea of decomposing the variation in the data into a systematic part describing all the dynamics and an unsystematic random part. If the test suggests that there are significant autocorrelations left in the model, our forecasting results will deviate systematically from an actual realization. If the chi-squared-test and F-tests derived for the VAR model, are based on the assumption of independent errors. If this would not be satisfied, the distribution of the tests will deviate from chi-squared and F in unknown ways.

Further the properties of the estimators may be sensitive to significant autocorrelations. In particular the ordinary least square (OLS) estimator is inconsistent when there are residual autocorrelations, see p. 74 in [9].

3.6.2.2 Test for Residual heteroscedasticity

Heteroscedasticity is often a concern for models based on data monitored with monthly or higher frequency, heteroscedastic residuals can indicate structural changes, see [10].

The presence of heteroscedasticity affects a number of standard inference procedures, such that the application of these methods may lead to conclusions that are not in line with the true underlying dynamics. In the time series context there is literature that makes some suggestions for valid inference with heteroscedasticity, see [11]. Meaning that much of the analysis can be done even if there is heteroscedasticity. That is because the VAR model represents the mean of the variables, which is often of primary interest.

Still, it may be useful to check for heteroscedasticity to better understand the properties of the underlying data, see [10]. We will do this in the results, by looking at residuals plots, looking for patterns.

3.6.2.3 Normality Test

Non-normality tests are often used for model checking, although normality is not a necessary condition for the validity of many of the statistical procedures related to VAR models. However, non-normality of the residuals may indicate other model deficiencies such as non-linearity or structural change. Multivariate normality tests are often applied to the residual vector of the VAR model, where, as a side note, the univariate versions are used to check normality of the errors of the individual equations, see [10]. We will use graphic methods to determine normality of the residuals.

Further the residual autocorrelation tests and the heteroscedasticity tests are derived under the assumption of normally distributed errors and the normality tests are derived under assumption of independent and homoscedastic errors.

This means that we do not know whether all the tests which already passed the residual examination can be trusted or not, and all residual examination tests in that case would need to be recalculated after the model has been re-specified, see p. 77 in [9]. This is nothing we will focus on, but it is worth mentioning.

With the aid of R, we will produce all the graphs and plots needed to investigate trends or irregularities on the residuals of the VAR model.

3.6.3 Vector Autoregressive Model Forecasting

Learning from the previous theory that we discussed, it indicates that VAR models are natural tools for forecasting, their setup as we gathered, is such that current values of a set of variables are partly explained by past values of the variables involved, see [10].

VAR model are said to often provide superior forecasts comparing to univariate time series models, since the forecasts of VAR models are quite flexible, and that they can be adjusted to one or multiple conditions regarding the potential future paths of specified variables in the model, see [6].

But one should be aware of one strain with VAR forecasting, that is, when we have big number of multiple dependent variables in a model, we will then get more parameters to estimate. An example might be if we would have 4 variables, and a lag of 4; this would mean that every equation would have 16 coefficients that would need to be estimated, making a total of 64 coefficients.

The more coefficients to be estimated the larger the chance of estimation error entering the forecast, i.e. it would be more likely getting a rather unreliable future scenario result. In practice it is usual to keep the number of variable small, and only include the ones that are correlated to each other so they are useful in the forecast, see [18]. Since we are only using two time series in the forecast to see if they are dependent on each other. We have a smaller chance of having big estimation errors.

Further we will not focus on the forecasting errors like we discussed with the univariate time series models. Instead we will focus on how much the variables actually contribute to one another in the forecast. We will do this with Variance Decomposition, which allows us to investigate and decompose the forecast error variance of a variable that is generated by the other variable in the system, see p. 146 in [7]. We will limit the rapport and not go into detail about how the actual calculations are made. We will rely on the help of R to make all the correct variance decomposition computations.

When preparing for the construction of our forecast. We keep in mind that when the ECB are producing their results for the stress test, they are forecasting 3 years in to the near future. We can argue that this is a long horizon forecast, but we have one way of conducting forecasts with this in mind, that is the *chain rule of forecasting.*

3.6.3.1 Chain rule of forecasting

When we are looking for longer future scenarios intervals, that we want to be able to compare with the ECB scenarios. The forecast loose accuracy the further we go into the future. Since we in common with the ECB have this problem we make sure that we understand how the forecasting works when we have fitted a model to our data. The question we have is how the expected value, or mean, in the future is produced for $\boldsymbol{Y}_{t+1}$. For our stochastic process $\boldsymbol{Y}_t$ and the VAR model with parameters $\boldsymbol{\Pi}_i$, we can produce a one step ahead forecast

$$E_t(\boldsymbol{Y}_{t+1}) = \boldsymbol{\Sigma}_{\boldsymbol{i=1}}^{\boldsymbol{n}}\ \boldsymbol{\Pi}_i \boldsymbol{Y}_{t+1-i}$$

If we would be interested in the *t+2* forecast, this method would not be used since $\boldsymbol{Y}_{t+1}$ is unknown. The chain of rule forecast allows $E(\boldsymbol{Y}_{t+1})$ to be substituted for $\boldsymbol{Y}_{t+1}$. So the *t+2* forecast would be

$$E_t(\boldsymbol{Y}_{t+2}) = \boldsymbol{\Pi}_{\boldsymbol{1}} E_t(\boldsymbol{Y}_{t+2}) + \boldsymbol{\Sigma}_{\boldsymbol{i=1}}^{\boldsymbol{n}}\ \boldsymbol{\Pi}_{\boldsymbol{i}} \boldsymbol{Y}_{t+2-i}\ .$$

Now we can derive the *t+k* forecast model

$$E_t(\boldsymbol{Y}_{t+k}) = \boldsymbol{\Sigma}_{\boldsymbol{i=1}}^{\boldsymbol{n}}\ \boldsymbol{\Pi}_{\boldsymbol{i}} E_t \boldsymbol{Y}_{t+k-i}\ , \mathrm{k} \geq 1$$

where

$$E_t(\boldsymbol{Y}_{t-i}) = \boldsymbol{Y}_{t-i}\ , i \geq 0$$

, see [4][12].

To get a result to compare with the macroeconomic scenarios produced by the ECB. We will be using programing in R with this approach, and R will handle the analysis and calculations when producing our forecasts.

4 Data

4.1 Description of Data

The macroeconomic scenarios produced by the ECB constitute an extreme percentile of a damaging macroeconomic scenario. Which also reflect the current economy and the near future. If we were to gather data from 50 years, we are going to catch trends that might not fully reflect our current macroeconomic status.

That is why we are going to gather historical data from the year 2004 up to 2014, as these will signify our current economic trends. To make sure we get enough of observations, we will be using quarterly data for our time series analysis. Which will be 44 observations within the different time series. Further the structure of the data is simple; each quarter gives us information about the percentage status of the current unemployment and change in house price index.

The data is gathered from the institution Statistics Sweden, SCB.

5 Results

5.1 Overview

Since we have talked about univariate and multivariate analysis, we are going to use our knowledge and conduct analysis on our data using both methods. We are going to fit models, look at the parameters significance in these models and test the residuals. This will end our analysis regarding the produced model; the models, created via stepwise selection, with the aid of R programing, will suffice and serve our aim regarding this paper. Since R is testing multiple different models and choosing the best one, we will limit the report here, and will not compare the models that might fit the data manually, just go through the main model that R will suggest.

Our next step will be to create the forecasts, look at possible errors in the time series model or error variances for the VAR model, and compare the two forecast results of each model. This is to see what kind of differences we might find in the end result. Our focus is going to be on the multivariate analysis, the VAR forecast, since this model resembles the needs of the ECB. And finally we will compare our end result with the ECBs constructed scenarios.

5.2 Time series

First lets take a look at how our time series look when plotting them over a time period.

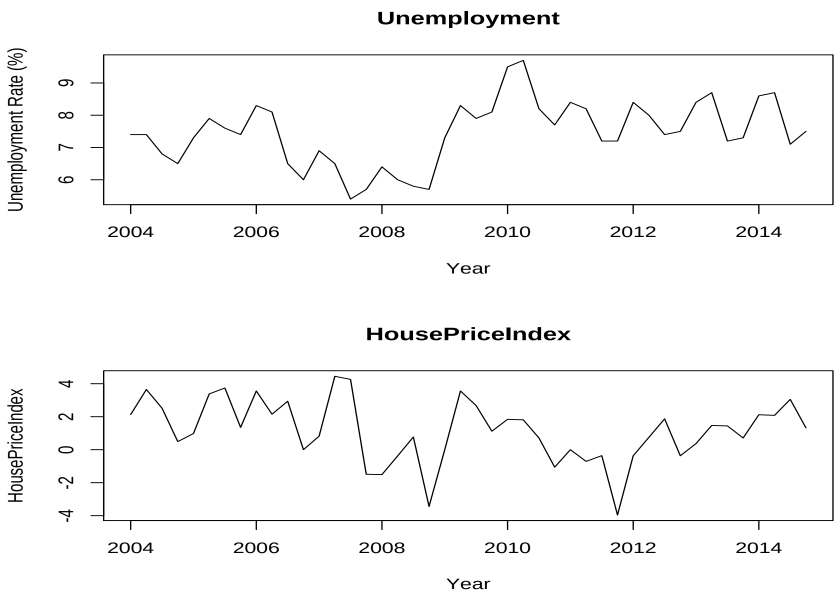

Figure 1. Time series plots over historical data of unemployment and house price index

By looking at figure 1 we are trying to distinguish what kind of assumptions can be made about our data. We can see that the time series shows some kind of trend near 2009, also some kind of seasonality cycle in half-year intervals. We cannot see that we have a constant mean and also the variation over the time periods is not constant. This gives us a hint of having non-stationary time series. If this would be the case, we would need to difference our data to make in stationary.

To make further assumption on the data we are going to look at the ACF plots, figure 2, here we might also see if our previous observation can be validated.

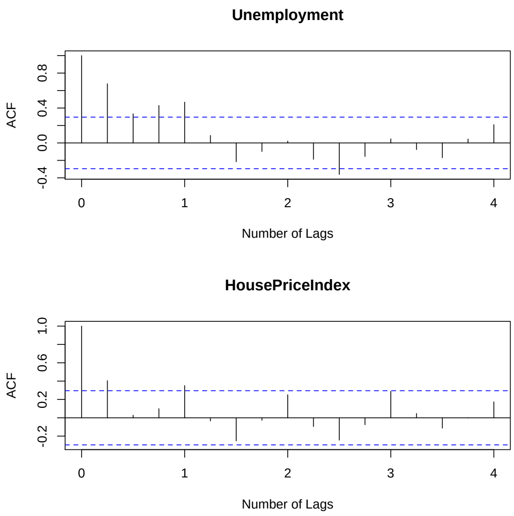

Figure 2. Plot of the autocorrelation function on unemployment and house price index.

When observing figure 2, we are looking at the lags over the different lag periods, if they cross the dotted line we can determine if they are significant. We are also looking for any type of periodic trend in the lags. That might tell us that we have some kind of seasonality in our data. Additionally we are looking at how fast the ACF will drop to zero correlation. If we can see that the drop is relatively fast, we might assume that the data is stationary. But if the decrease is slow we should assume a non-stationary data.

By examining figure 2, and focusing on the ACF for unemployment, we cannot find trends in our data, i.e. the spikes are random. Also we observe significant lag between 2 and 3. Which might hint of non-stationary data.

Investigating figure 2 with focus on the house price index, we can see a trend between the lags. Also there is a resemblance with the ACF of unemployment regarding the pace of the lags correlation dropping to zero. This tells us that we might have seasonal data of the house price index and that we are dealing with non-stationary data.

As a result of this, it is clear that we should be differencing our data to remove any seasonality and more importantly make our data stationary so the model fit will be good, as we discussed in section *3.4*.

5.2.1 ARIMA-model

The use of R and the function auto.arima, this will help us pick the best model for our data. It uses a stepwise selection method, which will try all different combinations of ARIMA models and ultimately give us the one model that fits our data best. It will difference the data if needed, also handle seasonality, and choose the proper components of auto-regressive, integrated- and moving average elements, to best fit our data. It will also estimate the parameters using MLE. When the model is fitted, we will go through what R produced, and test that they fulfill all the requirements we have on it.

5.2.1.1 Fitting data to ARIMA

The produced models, and the best fit for our data according to R, are for the time series of unemployment, an ARIMA (2,0,0)(2,1,1,)(4) with drift. And for the time series regarding house price index, the result is an ARIMA (1,0,0)(2,1,0,)(4) with drift.

It seems like a standard ARIMA model was not enough to fit the data from our time series. Instead R produced a Seasonal-ARIMA model with drift.

5.2.1.2 SARIMA Model for House Price Index

We would like to go through the models and understand which components where set to the data, figure 3 show us the model and the estimates of the coefficients with their respective standard errors.

```
Series: HousePriceIndex
ARIMA(1,0,0)(2,1,0)[4] with drift

Coefficients:
         ar1     sar1     sar2    drift
      0.6352  -0.6545  -0.3531  -0.0079
s.e.  0.1363   0.1531   0.1511   0.0812
```

Figure 3. House Price Index SARIMA model with parameters estimates and standard errors.

We see that from (1,0,0) that we have 1 auto-regressive-, 0 differenced and 0 moving average components for the non-seasonal part of the model. The seasonal part of the model (2,1,0)(4), tells us that model has 2 seasonal AR components, 1 seasonal differential component and 0 seasonal MA component. The (4) stands for the numbers of periods we have per season, and this would be quarterly periods in our case. Since our data is based on quarterly observations. Also we can see that drift was added to our SARIMA model.

Now we want to control that the coefficients are significant for the model, if we look at figure 4, we see the corresponding p-values. And we can tell that the p-values are very low, indicating significance to all parameters in the model.

```
         ar1         sar1         sar2        drift
3.144884e-06 1.910784e-05 1.946970e-02 9.222935e-01
```

Figure 4. P-values of the parameters in the house price index SARIMA model.

5.2.1.3 SARIMA Model for Unemployment

Proceeding to investigate the model of unemployment, we notice in figure 5 that we have a few more parameters to our model.

```
Series: Unemployment
ARIMA(2,0,0)(2,1,1)[4] with drift

Coefficients:
         ar1      ar2     sar1     sar2     sma1   drift
      1.1450  -0.2988  -0.1185  -0.4592  -0.7012  0.0297
s.e.  0.1693   0.1641   0.2014   0.1644   0.2188  0.0249
```

Figure 5. Unemployment SARIMA model with parameters estimates and standard errors.

The SARIMA model for unemployment, we see from (2,0,0) that we have 2 auto-regressive-, 0 differenced and 0 moving average components for the non-seasonal part of the model. The seasonal part (2,1,1)(4), tells us that model has 2 seasonal AR components, 1 seasonal differential component and 1 seasonal MA component. We have the same periods of (4) per season, and this is as we mentioned because we have quarterly periods. Also in this model we can see that drift was added to our SARIMA model.

We want to make the same control on the coefficients as we did to the model for house price index. Looking at figure 6, we again see very low p-values to the corresponding coefficients that indicate significance to our model.

```
         ar1          ar2         sar1         sar2         sma1        drift
1.358402e-11 6.861213e-02 5.563110e-01 5.228291e-03 1.351778e-03 2.335609e-01
```

Figure 6. P-values of the parameters in the unemployment SARIMA model.

5.2.1.4 Testing the SARIMA models residuals

The next step to investigate if R has done a good job of fitting these models will be to look for correlation, heteroscedasticity and normality in the residuals.

Figure 7. Scatterplot to investigate heteroscedasticity of residuals of SARIMA model with House Price Index

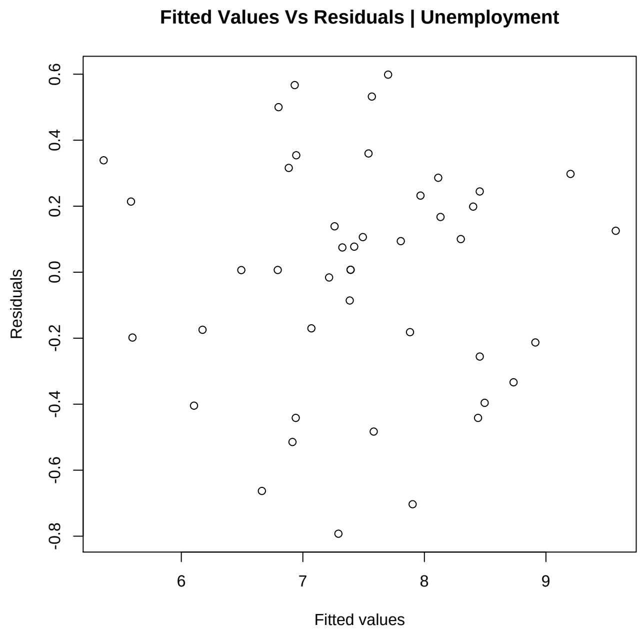

Figure 8. Scatterplot to investigate heteroscedasticity of residuals of SARIMA model with unemployment

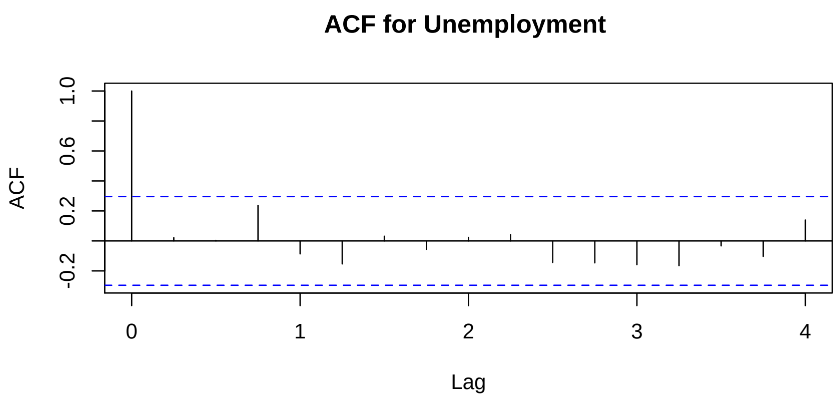

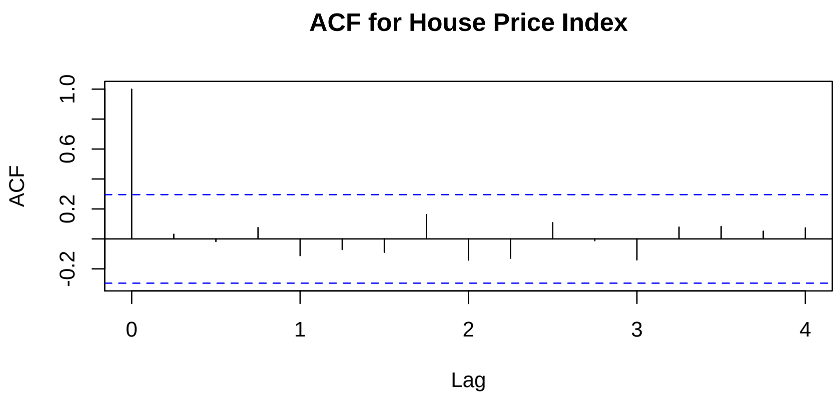

Figure 9. ACF of residuals from the SARIMA models of unemployment and house price index

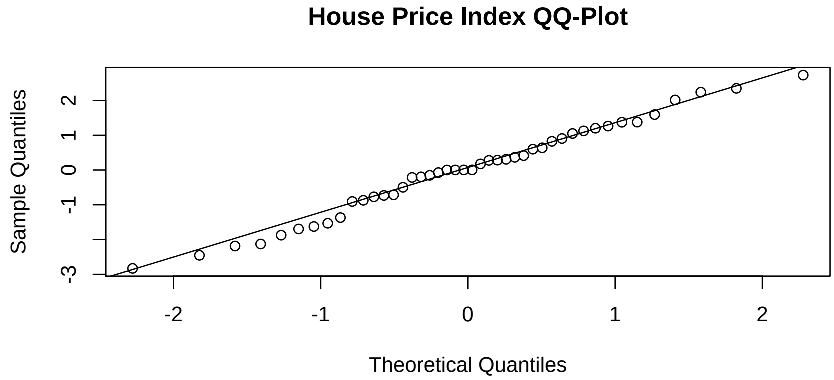

Figure 10. QQ-plots of residuals from the SARIMA models of unemployment and house price index

With the help of figure 7-10 we can make sure that the residuals of the created models are satisfying our assumptions, i.e. linearity, constant variance (homoscedasticity) and lack of auto-correlation.

From figure 7 and figure 8, we see that the cluster between the values has approximately the same width all over. We cannot see that the cluster of values has any tendency of a larger width when the values are getting larger. We choose to interpret this as a sign of homoscedasticity.

From figure 9 we can tell that the auto-correlation between our residuals are good. And from figure 10 we can see that our residuals are showing normally distributed pattern. Since all checks on residuals of the models satisfy our assumptions, we can go ahead and create future forecast on our data.

5.2.2 Forecast with the produced models

When creating the forecast, we need to be specific to what scenario we need to produce. In our case we want to be able to compare scenarios to the ECB, looking at section *4.1*, we mentioned that the ECB are producing their scenarios for a future outcome in the worst-case scenario, i.e. a damaging macroeconomic scenarios that constitute an extreme percentile.

When R helps us complete the forecast, it also creates prediction interval (P.I) of the expected value for the future, R uses the theoretical variance of the forecast distribution, created from the main model, and also two prediction levels of 80% and 95%, see p. 119 in [2].

We can see this in figure 11, where the external gray area represents a P.I of 80%, and the inner area a 95% P.I. Since we want the most "extreme" scenario, we will have to use the external results. As these represent the most extremt outcome in the future in our case. In table 1, we can see the 80 % P.I of the expected unemployment rate and the expected house price index for the stress test years.

Now to investigate what kind of error we might have, we take a look at the MAE that R has calculated for both forecast, we have a value of 1.043862 for the house price index forecast. And a MAE value of the unemployment rate forecast is 0.2823115. Meaning we have an approximal difference between forecast and observed house price index of 1.04, and for the observed unemployment rate, we have approximately difference of 0.28 between the observed value and the forecast. These errors are fairly low, and also quite good for the estimated forecast.

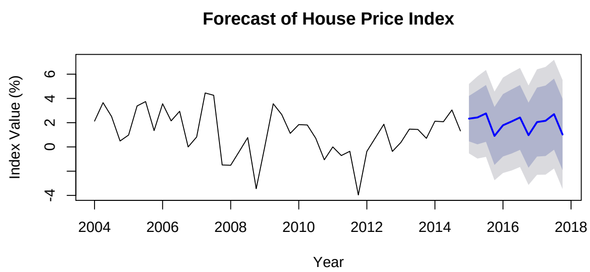

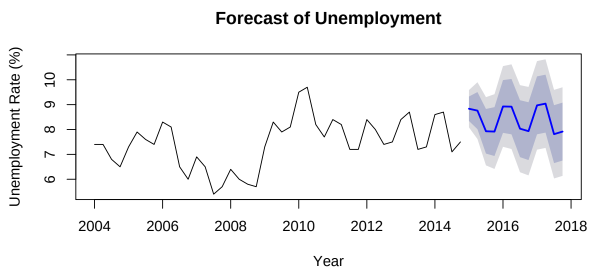

Figure 11. Future forecast of the unemployment and house price index SARIMA models

Year->	2015	2016	2017
Expected House Price Index (%)	-1,48 %	-1,71 %	- 1,92 %
Expected Unemployment Rate (%)	8,89 %	9,09 %	9,08 %

Table 1. Expected values at a 90% P.I of the future forecast of the unemployment- and house price index SARIMA models

5.3 Vector autoregression

Now that we have produced our scenarios with the univariate technique, we are proceeding with the use of the multivariate analysis of our data. We are again using aid from R to help us get the best-fitted VAR model to our data.

5.3.1 Fitting the data to the VAR model

Estimation of parameters and picking the best model for us is something R helps us with. The most important part of the process will be to determine the appropriate lags for our model. R will, with the use of a stepwise selection across the different lags selection, pick the optimal model. The result when preforming the stepwise selection ends with a lag of 1, which is a VAR (1) model. This means that the data is best represented when the subject variable has 1 lag of itself and 1 lags of the other variable.

That is the unemployment is best estimated when used 1 lag of its own data and 1 lag of house price index, and vice versa regarding house price index model. In other words we have at first glance some kind of dependency between our subjects, as we suspected. Further we need to control the model R fitted to our time series data.

5.3.1.1 VAR model for Unemployment

The VAR model for unemployment is best estimated when using 1 lag of its own data and 1 lag of house price index. Lets take a look at figure 12, and see how the estimations look and if they have significance to the model.

```
Estimation results for equation Unemployment:
=============================================
Unemployment = HousePriceIndex.l1 + Unemployment.l1

                    Estimate Std. Error t value Pr(>|t|)
HousePriceIndex.l1 -0.008247   0.009935  -0.830    0.411
Unemployment.l1     0.744471   0.106978   6.959 2.14e-08
```

Figure 12. Unemployment VAR model with parameters estimates and standard errors and p-values.

Further in figure 12, we see the VAR (1) for unemployment takes shape, but more important, we notice that the parameter of house price index is not significant with its high p-value of 0.411. The best action would be to reduce the model, but for the sake of restricting the report, we leave the VAR model as it is. We are more interested in how the model works, the focus is not to get a perfect result, but will take note of this issue when making a final assessment of the VAR model.

5.3.1.2 VAR model for House Price Index

Producing the VAR model for House Price Index with R, we get a similar result as for the unemployment VAR model, i.e. we have a best estimated when using 1 lag of its own data and 1 lag of unemployment. Taking a look at figure 13, we have a parameter that is not significant to the model, which is unemployment, with a p-value of 0.30672. Again the model should be reduced, but as we previously mentioned, for the sake of restricting the report leave the VAR model as it is, but take note of this issue.

```
Estimation results for equation HousePriceIndex:
================================================
HousePriceIndex = HousePriceIndex.l1 + Unemployment.l1

                   Estimate Std. Error t value Pr(>|t|)
HousePriceIndex.l1   0.3970     0.1423   2.789  0.00805
Unemployment.l1      1.5868     1.5326   1.035  0.30672
```

Figure 13. House Price Index VAR model with parameters estimates and standard errors and p-values.

We will continue our investigation on the models by looking at their residuals. And suspect that they might not be optimal since, in both VAR models, we have one parameter that is not significant to the model.

5.3.2 Control of the residuals for the VAR

Controlling the VAR model residual we will as we did before with the univariate models, focus on the basic assumptions regarding the residuals. This handles the autocorrelation, normality and looking for Heteroscedastic in the residuals.

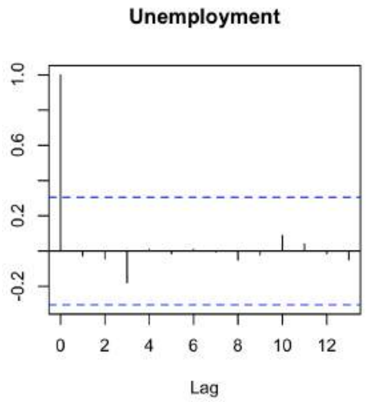

Figure 14. ACF of residuals from the VAR models of Unemployment

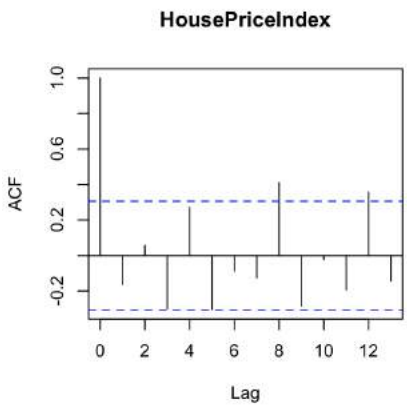

Figure 15. ACF of residuals from the VAR models of House Price Index

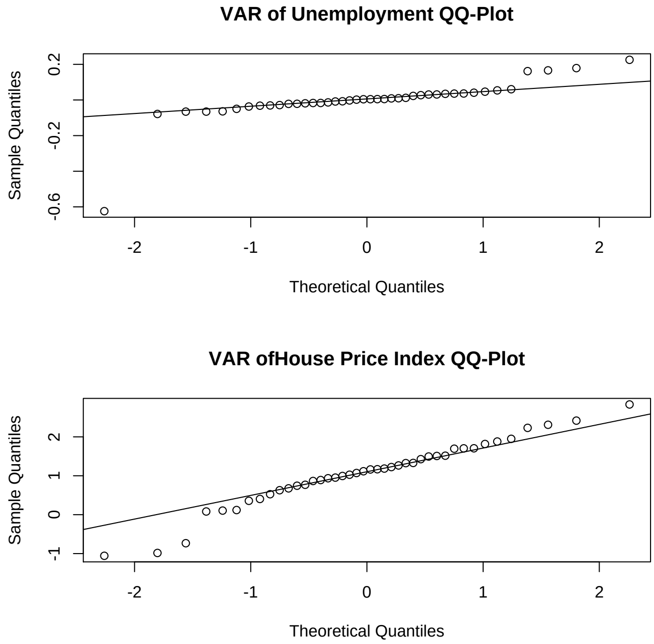

Figure 16. QQ-plots of residuals from the VAR models of unemployment and house price index

Figure 17. Scatterplot to investigate heteroscedasticity of residuals of VAR model of unemployment and house price index

We start by looking at the autocorrelation plots of the VAR model residuals, we can se that figure 14 is satisfying and not showing any evidence of late autocorrelation. But figure 15, ACF of house price index residuals, patterns has clear spikes on the significant blue doted line. This is an indicator of autocorrelation; this means that our forecasting results will deviate systematically from an actual realization.

An additional reflection on figure 11, is that we see little to none resemblance with figure 9, i.e. when comparing the ACF for the residuals with the univariate model. This could be because of the VAR feature where the data is fitted in a different way to the model, here the multiple variables are made to depend on each other, which is something the univariate model does not do. We will not go into detail or try to fix this issue, but we are taking this into account when forecasting and analyzing the end results.

Further investigation on the residuals shows that the normality seems to be fine when looking at figure 16. But we notice an extreme outlier on the QQ-plot for unemployment, this point should reflect the observation made near 2008, see figure 1. We also find outliers on the QQ-plot for house price index, these should be the observation made around 2008 and 2012. Luckily the outlier does not seem to influence the regression line, i.e. because the regression line runs through the majority of the observations and the regression line doesn't deviate from the observations to a point that we should further investigate the outliers, se we leave the outliers and conclude that figure 16 satisfies our requirements of normality on the residuals.

Lastly looking for heteroscedasticity of the residuals in figure 17, we see the same outliers that we see in figure 16, removing these outliers would not change our analysis of the scatterplots in figure 17. Additionally we make the same analysis as we did previously and leave them as they are.

Further looking at figure 17, we have a good pattern of homoscedasticity in the residuals of the house price index. The same cannot be said about the plot of residuals from the VAR model of unemployment. Here the patterns are really strong and suggests heteroscedasticity. Again we will only acknowledge the possibility of this problem, but wont take any action against it when we conduct our forecast.

We see that the assumptions of the residual in the regression are not fully fulfilled, and understand that this limits our ability to make conclusions. One alterative to adjusting the VAR model would be to transform the data and see if this would help with the residual assumptions. We will avoid transforming the data, and remind the reader that this paper is made to illustrate what different models the ECB might be using among its great arsenal of complicated analytical tools, and how they work in theory.

5.3.3 VAR model forecasting

We will now go ahead and produce a prediction into the future with the VAR models, we are again using R to help us conduct these forecast. The chain rule of forecasting, that we have discussed, now comes to use. This is since we again are producing predictions into a long time period. The results from the scenarios are expected values, with a prediction level of 95%, where the expected house prices index values are the lowest value of its P.I, and the expected unemployment rate value is the highest value of its 95% P.I, we pick these values for the same reason we picked the values for the univariate models, i.e. they would represent the very extremes of the future scenarios, the values are displayed in table 2.

Year->	2015	2016	2017
Expected House Price Index (%)	-2,47 %	-2,62 %	- 2,65 %
Expected Unemployment rate (%)	7,71 %	7,70 %	7,69 %

Table 2. Expected values at a 95% P.I of the future forecast, of the unemployment and house price index VAR models

Further we want to investigate the variance decomposition with the help of R, i.e. see how much of the error variance of the variables are generated by each other in the forecast models.

In table 3, we see the figures that R calculated for us, that is, we can with the results see how much of the variance in the forecast is dependent of its own data set and on the other variable in the model.

Further examining of table 3, we can see that in 2015, the forecast of the house price index VAR model impact on the ability to forecast the house price index, with unemployment, only accounts for 3,3 % of the variance. And for the forecast of the unemployment VAR model, we see that house price index only generates around 3,8 % of the variance of the forecast. These are rather low impacts, but we could expect this, since we previously in figure 12 and 13 noticed that the corresponded second parameters showed little significance to the models.

Further in table 3, we see that the variance of the forecasted variables in the where based on their own data set with 96 % in 2015, 95 % 2016 and 2017. The variables stand rather independent of each other.

Forecast Horizon (Year)		HousePriceIndex	Unemployment
2015	HousePriceIndex	0.9668209	0.03317915
	Unemployment	0.038003771	0.9619962
2016	HousePriceIndex	0.9549596	0.04504037
	Unemployment	0.045839592	0.9541604
2017	HousePriceIndex	0.9541140	0.04588596
	Unemployment	0.046381412	0.9536186

Table 3. Variance decomposition for house price index VAR model and unemployment VAR model

5.4 Comparing results of the chosen analysis models

We will compare the results from each model. If we start by looking at the different residuals plots. We could see that we had little to none hesitation regarding the analysis made over the plots using the SARIMA model. The same cannot be said about the VAR model, we had signs of autocorrelation and heteroscedasticity. This did not fit our assumption regarding the residuals of the VAR model, and it should be investigated. And preferably have a detailed explained to why we had issues, and ultimately correct the problems. This would probably make the models better and in our case, producing a more valid forecast.

Further we could see that the cross variable in our VAR model where not significant, and that they where not very useful to each other when doing the forecasts on the two VAR models. This points us in the direction of believing that the two time series are not very dependent on each other. Which would mean that the optimal way of dealing with these particular time series would be to use them in a univariate model.

Still the use of a VAR model is valid, we needed simulate a multivariate way of dealing with data, also analyzing if the two time series would add for a better forecast, and if they where dependent on each other when forecasting to get a indication of how the ECB are working with these questions.

The forecast results between the SARIMA models and the VAR models are different, that might cause an issue, but since we are dealing with future values nothing is an exact science. But the results are pointing in the same direction, and since we only need projection of what might happened in the future, the results do their purpose.

One way of controlling our scenario results and with that the models, would be to preform a backtesting on historical data, i.e. comparing scenario results with actual observed observation. This would indicate if our models would be able to create good forecasts, and if the models worked in a good way.

This is something we will leave out of the rapport, because our testing is against the ECBs results, and mainly getting an idea of how to use the models. Further we presume the ECB have made necessary back testing and are satisfied with their own results, so we will relay on this, and see the ECBs results as accurate predictions.

5.5 Comparing results with ECB macroeconomic scenarios

Before we start comparing our results with the ECB, we are going to present the ECBs macroeconomic scenarios, they represents a predicted result that strains the macroeconomic stability, i.e. these results have an extreme percentile that deviates from the un-strained predictions of the macroeconomic stability. The scenarios are produced and finalized by the ESRB regarding the Swedish macro economy. The times period for the scenarios are from the year 2015-2017, see [15].

Year->	2015	2016	2017
House Price Index (%)	-11,9 %	-19,1 %	- 19,1 %
Unemployment rate (%)	8,6 %	10,8 %	12,6 %

Table 4. Numerical values of future forecast, that constitute an extreme percentile, of the unemployment and house price shocks produced by the ESCB

We are now looking for similarities in our forecast, which we see in table 1 and 2, then comparing our values in relation to table 4. We can se a significant difference between our forecasted house price index and the predictions presented by the ECB. Does this result mean that we entirely made a wrong assessment of what models are used to produce these scenarios by the ECB?

The answer took some real detective work, and surely we found out why the predictions were so different from our results. By talking with Emil Hagström, who is working for the Swedish Financial Supervisory Authority. The explanation is given as follows, "*Before the ECB publishes the result of the macroeconomic scenarios. They send out a draft to each country equivalent of Financial Supervisory Authority, where there are banks that are taking part in the Stress-test exercise. This gives us a chance to comment and maybe enforce a change to the end result. In the case of house prices in Sweden we regarded the scenarios being very low. So we opted for a change with a higher stress level. And this is the change you know see in the rapport.*" With this knowledge we cannot compare this data to our result. But we can appreciate that our calculations where not horribly wrong.

The results of the unemployment are now of great interest. Are we able to see any signs of similar results? The answer is yes; yes we can actually see that the results are not far away. We see more similarities with the results of the SARIMA model, compared to the VAR model. The values of the SARIMA model are closer to ECB results and we have a similar up-going trend, with higher unemployment later in the forecast period. For our simplistic analysis, compared to the ECBs more refined way of working, we are happy with the similarities between the forecast.

A noticeable observation is that the predicted values from the ECB are very extreme. In our own conducted forecasts, with the lowest prediction interval being 90 %, where we picked the values from the very end of the interval range, we still could not get a wide enough range to get the extreme values comparing to the ECB scenarios. This would lead us to think that the results the ECB is presenting, has a very wide interval, this is probably to capture the very extremes of outcomes in the future macro economy, also it is very likely, that they are making some kind of an ad-hoc adjustment to the scenarios. This way of working might be a standard procedure within the ECBs output of the macroeconomic scenarios, to make them further severe, and make it tougher against the bank stress test.

6 Discussion and Conclusion

6.1 Discussion and Conclusion

The models we have analyzed have given us insight to what procedures are being used when the ECB are producing their macroeconomic scenarios for the bank stress tests. The theory behind the VAR model shows that the ECB has a powerful tool to aid them in the work of producing the scenarios. We have also seen that the VAR has many, if not all components that the ECB are looking for in a statistical tool.

Unfortunately our results from the VAR were not perfect. We have mentioned a few flaws in our results, and that these could possibly have been adjusted with different statistical methods. The flaws mainly came from the fitting of the time series data to the VAR model, which in turn, affected the end result.

Even without optimizing our model, we get a result that we can compare with the ECB and furthermore see similarities with. This is interesting as it shows that the model we chose to work with really could be used for the ECBs purpose, not only in theory but also in practice.

Additionally we found it very interesting that the SARIMA model showed more scenario similarities with the ECB data, comparing to the VAR model. The VAR model is used to capture covariance in multiple time series; this is something ARIMA and SARIMA models lack, where we do simple marginal adjustments. Also the VAR model is compressing multiple data sets, and as we have seen, if that process is not investigated, we might loose certainty in the end result. Further we would like to consider that this demonstrates that the univariate method could be useful in different forecasts, where you only focus on one time series.

We have also learnt that having a few models to be able validate the results is a clear advantage. As we mentioned in section *2.1*, we know that the ECB has several statistical tools, even more complicated tools, to aid them in their work. Further we have seen that the ECB is using ad-hoc methods, like the method of reaching out to the non-biased Swedish Financial Supervisory Authority, for a second opinion to optimize their end result.

This is a good thing, because making an exact science of an expected scenario is very hard. The ECB should not only rely on statistical methods, they should, as they are, use their knowledge of current and historical trends, data structure and choosing of the correct models to get the optimal result.

In conclusion, we recognize that the VAR model is a powerful tool, which can aid the ECB in their work for future scenarios. And also realize that the ECB seem to have a good portfolio of methods regarding the ways to produce macroeconomic scenarios for the bank stress tests.

7 Appendix

7.1 References

[1] S. Allen Broughton (2008) Discrete Fourier Analysis and Wavelets: Applications to Signal and Image Processing (Page 72).

[2] Robert H. Shumway & David S. Stoffer (2006) Time Series Analysis and Its Applications

[3] https://datajobs.com/data-science-repo/Time-Series-Analysis-Guide.pdf 22 April 2015

[4] https://onlinecourses.science.psu.edu/stat510/node/79 27 April 2015

[5] https://www2.bc.edu/~iacoviel/teach/0809/EC751_files/var.pdf 27 April 2015

[6] https://faculty.washington.edu/ezivot/econ584/notes/varModels.pdf 28 April 2015

[7] Gebhard Kirchgässner & Jürgen Wolters & Uwe Hassler (2013) Introduction to Modern Time Series Analysis

[8] http://cran.r-project.org/web/packages/vars/vignettes/vars.pdf 28 April 2015

[9] Katarina Juselius (2006) The Cointegrated VAR Model

[10] Helmut Luetkepohl (2011) Vector Autoregressive Models

[11] Ralf Brüggemann & Carsten Jentsch & Carsten Trenkler (2014) Inference in VARs with Conditional Heteroskedasticity of Unknown Form

[12] Craig S. Hakkio & Charles S.Morris (1984) Vector Autoregressions: A Users's Guide

[13] European Central Bank (2013) Occasional Paper Series; No 152

[14] European Central Bank (2013) Tools for Macro Stress Testing and Macro-Prudential Policy Assessment: The ECB Perspective

[15] European Systemic Risk Board (2014) EBA/SSM Stress Test: The Macroeconomic Adverse Scenario

[16] Chingnun Lee (2006) Maximum Liklihood Estimation (Page 7).

[17] Ruay S. Tsay (2005) Analysis of Financial Time Series (Page 68).

[18] Rob J Hyndman & George Athanasopoulos (2013) Forecasting: Principles And Practice (Page 273).

[19] http://sfb649.wiwi.hu-berlin.de/fedc_homepage/xplore/tutorials/xegbohtmlnode46.html#msarima_expand 20 Juni 2015