

Scenario Generation for Financial Market Indices

Peter Nystrup
s082634

Kongens Lyngby
January 2012

Technical University of Denmark
Informatics and Mathematical Modelling
Building 305, DK-2800 Kongens Lyngby, Denmark
Phone +45 4525 3351, Fax +45 4588 2673
reception@imm.dtu.dk
www.imm.dtu.dk

IMM-BSc-2012-2

Abstract

The purpose of the project is to analyse the given index data with the aim of generating scenarios that can form the basis of decisions regarding strategic asset allocation. The data available is observations of the daily closing value of eleven financial market indices; six stock market indices, four bond market indices, and one money market index.

The distribution of the log returns is a mixture of a normal distribution and a few extreme observations from a different distribution. Traditional measures reject any resemblance with the normal distribution due to the presence of extreme values. Yet, more robust measures are able to find certain similarities.

The index data has a growing mean trend, a time dependent variance, auto- and cross-correlation. A log return transformation is suitable for making the mean stationary, sampling on weekly basis instead of daily eliminates most of the autocorrelation, a regime model, where upturns and downturns are modelled independently, handles the changing cross-correlation, and finally ARCH-models are employed to model the changing variance.

The approach taken is to model the principal components of the log return series. The advantage of modelling the independent principal components is the possibility of reducing the dimensionality based on the strong cross-correlation exhibited, as well as easy applicability for scenario generation.

The simulated scenarios were compared to bootstrapped scenarios, and were found to be better at reproducing the crises observed in the dataset. The abil-

ity to reproduce realistic crises appeared to be connected with the ability to reproduce volatility clustering correctly.

The bootstrapped scenarios were also more optimistic with median returns closer to the average returns from the dataset, where the scenarios simulated from the regime model were more influenced by the recession at the end of the data period/the beginning of the scenarios.

Keywords: Financial Returns, Leptokurtic Distributions, Principal Component Analysis, Conditional Heteroscedasticity, Non-linear Time Series Modelling, Regime Models, Leading Indicators, Scenario Generation, Bootstrapping.

Resumé

Formålet med projektet er at analysere det givne indeksdata med henblik på at generere scenarier, der kan danne grundlag for en beslutning om strategisk aktivallokering. Det tilgængelige data er observationer af den daglige lukkepris for elleve finansielle indeks: Seks aktieindeks, fire obligationsindeks og et penge-markedsindeks.

Fordelingen af de logaritmiske afkast er en blanding af en normalfordeling og nogle få ekstreme observationer fra en anden fordeling. Traditionelle test afviser enhver lighed med normalfordelingen grundet de ekstreme værdier. Mere robuste teststørrelser er dog i stand til at finde visse ligheder med normalfordelingen.

Indeksdataet har en voksende middelværdi, en tidsvarierende varians, auto- og krydskorrelation. En transformation til logaritmiske afkast gør middelværdien stationær, ugentlig sampling i stedet for daglig eliminerer det meste af autokorrelationen, en regimemodell, hvor opsving og nedgang modelleres uafhængigt, håndterer den varierende krydskorrelation og endelig anvendes ARCH-modeller til at modellere den skiftende varians.

Fremgangsmåden, der benyttes, er at modellere de principale komponenter for de logaritmiske afkast. Ved at modellere de uafhængige principalkomponenter kan dimensionen af problemet reduceres, idet den udviste stærke krydskorrelation udnyttes, hvilket har yderligere den fordel, at modellen er let anvendelig til scenariegenerering.

Ved en sammenligning med de bootstrappede scenarier blev de simulerede scenarier fundet at være bedre til at genskabe de observerede kriser i datasættet. Evnen til at genskabe realistiske kriser syntes relateret til evnen til at genskabe volatilitetsklumpning korrekt.

De bootstrappede scenarier var også mere optimistiske med medianafkast tættere på de gennemsnitlige afkast fra dataperioden, hvor scenarierne simuleret ud fra regimemodellen var mere præget af recessionen i slutningen af dataperioden/begyndelsen af scenarierne.

Nøgleord: Finansielle afkast, leptokurtiske fordelinger, principal komponent analyse, betinget heteroskedasticitet, ikke-lineær tidsrækkemodellering, regime-modeller, ledende indikatorer, scenariegenerering, bootstrapping.

Preface

This thesis was prepared at the Department of Informatics and Mathematical Modelling at the Technical University of Denmark in partial fulfilment of the requirements for acquiring the BSc degree in engineering.

The thesis deals with different aspects of mathematical modeling of financial index data with the aim of generating scenarios useful for an asset allocation decision. The main focus is on data analysis and non-linear time series modelling.

J.P. Morgan, Morgan Stanley Capital International, Nasdaq OMX, Nordea, as well as the Tokyo Stock Exchange should be mentioned as sources of the index data that form the foundation of the thesis.

I would like to thank my supervisors Lasse Engbo Christiansen and Kourosh Marjani Rasmussen for the many instructive conversations and discussions regarding the project and its context.

Copenhagen, January 2012

Peter Nystrup

Contents

Abstract	i
Resumé	iii
Preface	v
1 Introduction	1
1.1 Asset Allocation	2
1.2 Scenario Generation	3
1.3 Problem Statement	4
2 Market Indices	7
2.1 Stock Market Indices	8
2.2 Bond Market Indices	11
2.3 Money Market Indices	13
3 Data Analysis	15
3.1 Data Reconstruction	16
3.2 Considering Stationarity	17
3.3 The Distribution of Data	22
3.4 Correlations	29
3.5 Principal Component Analysis	37

4	Modelling	45
4.1	Choosing a Regime Approach	46
4.2	Choosing a Model	51
4.3	Estimating Principal Components	54
4.4	Estimating Model Parameters	58
4.5	Assessing Model Performance	61
5	Scenario Generation	63
5.1	The Choise of Regime	64
5.2	Generating Scenarios	66
5.3	Assessing Scenarios	68
5.4	Bootstrapping	72
5.5	Bootstrapping Performance	75
6	Summary and Conclusions	81
6.1	Discussion	83
6.2	Possibilities for Future Work	84
	Bibliography	85
A	Additional Graphs and Tables	89
B	R-code	97
B.1	R-code for Chapter 3	98
B.2	R-code for Chapter 4	109
B.3	R-code for Chapter 5	113

CHAPTER 1

Introduction

In a pension fund like ATP¹, the board decide on the overall risk budget based on their expectation for the future Sharpe ratio, that is excess return per unit volatility. The larger the expected Sharpe ratio, the more risk the board will be willing to take on. Based on their expectations, the board also decide on certain benchmarks such as the ratio of stocks to bonds as part of the risk management.

The portfolio managers research the investment opportunities and come up with proposals that the chief investment officer then either rejects or allocates risk to. This is where the actual asset allocation happens, even though the benchmarks set by the board put certain constraints on the allocation.

When searching for investment opportunities, the portfolio managers apply a top-down approach. Based on scenarios and analyses they decide on a region and country of interest, and then the asset class. Then they narrow it down to a sector, and then finally a specific asset within that sector.

The first chapter gives an introduction to the concepts of asset allocation and scenario generation, which are the framework of this thesis.

¹The investment decision in ATP, Arbejdsmarkedets Tillægspension, as sketched by a portfolio manager from ATP during a guest lecture at DTU.

1.1 Asset Allocation

Investors, both institutional and private, decide upon investment policies or strategies through considerations of their objectives and constraints. Objectives being return requirements and risk tolerance. Constraints being liquidity, time horizon, regulations, taxes, and unique needs. Portfolio policymaking is about asset allocation, diversification, risk positioning, tax positioning, and income generation, where the most important part is the allocation of assets. Asset allocation is the decision on how to divide the portfolio onto the different major asset categories:[2]

- a. Money market instruments (usually called cash).
- b. Fixed-income securities (usually called bonds).
- c. Equities (usually called stocks).
- d. Derivatives such as forwards, futures, options, and swaps.
- e. Real estate.
- f. Commodities such as oil and precious metals.
- g. Foreign currency.
- h. Other.

Often, investors only consider including some of the above mentioned asset categories in their portfolio. The process of asset allocation can be divided into four steps:

- i. Specifying the asset classes to be considered for inclusion in the portfolio.
- ii. Identifying the expectations of future rates of return.
- iii. Finding the portfolios that deliver the maximum expected return for any given degree of risk.
- iv. Determining the allocation of assets that best meets the risk and return objectives while satisfying the constraints.

The expectations of future rates of return are normally determined from historical data and economic analyses. This includes generating scenarios.

The fundamental idea behind asset allocation is to exploit that different asset classes offer returns that are not perfectly correlated, for which reason diversification reduces the overall risk in terms of tail risk and the volatility of returns for a given level of expected return.² There are different asset allocation strategies depending on investment objectives, time frame, risk tolerance, and demands for diversification. In this thesis, the focus will be on strategic asset allocation. The primary goal of strategic asset allocation is to find the optimal balance between expected risk and return for a long-term investment horizon.[27]

1.2 Scenario Generation

A general way to describe risk is by using scenarios. A scenario being a realisation of the future value of all parameters that influence the portfolio. A collection of scenarios should capture the range of variations that is likely to occur in these parameters including the impact of the shocks that are likely to come. These representations of uncertainty are the cornerstone of risk management. It is important to emphasise that the purpose of generating scenarios is not to forecast, what will happen.

There are three overall approaches to generating scenarios that should be mentioned. The first is to generate scenarios by sampling historical data through bootstrapping. The second approach is to generate scenarios through random sampling and then accept each scenario if its statistical moments match those of the observed data. The third approach, which is also the approach that will be emphasised in this project, is to develop a theoretical model with parameters calibrated to historical data and then simulate this model to generate scenarios. In order to obtain a standard of comparison there will also be generated scenarios using the bootstrapping approach.

Before scenarios are used as basis of an asset allocation decision, the quality of the scenarios has to be checked. This can be divided into three stages; checking accuracy, correctness, and consistency. The process of generating scenarios typically involves a discretisation of a continuous process. Approximations like this of course give rise to a number of errors. The sum of the accumulated errors has to be controlled in order to secure accuracy.

Correctness is whether the generated scenarios contain the properties known

²Tail risk arises, when the probability of an investment moving more than three standard deviations from the mean is greater than what is assumed by the normal distribution.

from theory or empirical evidence. This also implies that the scenarios have to satisfy the no-arbitrage principle. An arbitrage opportunity arises, if it is possible to make a riskless profit. In a perfect financial market, in which all investors have access to the same information and can react instantaneously, there should not be any arbitrage opportunity. Otherwise each investor would immediately try to realise the riskless profit. The resulting transactions would of course eradicate the arbitrage opportunity. Thus, the estimated future values of asset prices should not allow arbitrage opportunities.[6]

When generating scenarios for more than one stochastic variable, the scenarios also have to be internally consistent. For instance an increase in the interest rate together with an increase in the bond prices is inconsistent.

1.3 Problem Statement

This thesis is concerned with the second part of the asset allocation decision only, that is scenario generation to be exact. As outlined in the introduction, the generated scenarios are of utmost importance to the investment decision process and the risk management in for instance a pension fund. Generating sufficient scenarios is therefore a practical problem of high relevance.

The asset classes considered are limited to money market instruments, bonds, and stocks, with the goal of establishing the correlation between these asset classes. Inclusion of the other major asset classes remains a possibility for future work.

The data available is twelve years and seven months of daily values of eleven different indices covering the period from 1st of January 1999 to 12th of August 2011. Six of them are stock indices, four are bond indices, and finally a Danish LIBOR-index, which will serve as the link to the money market. For the LIBOR-index, the data is only available from 16th of June 2003 and onwards. The indices will be explored in more detail in the following chapter.

The purpose of the project is to analyse the index data with the aim of generating scenarios that can form the basis of decisions regarding strategic asset allocation. A scenario in this sense is the future values of the indices. The time horizon of the generated scenarios will be five years, which is a reasonable horizon for a short term, strategic asset allocation decision. With ten years' index data

available, it would not be meaningful to look at a longer horizon than five years. There will be generated a number of scenarios, and the quality of these scenarios will be assessed.

The analysis will proceed according to the following steps:

1. The raw data is analysed for outliers, distribution, trends, autocorrelation, and cross-correlation.
2. A time series model is chosen and calibrated to the index series.
3. The model performance is tested on the data.
4. There will be generated scenarios using two different methods and the quality of the scenarios will be assessed.

The analysis will be conducted using the statistical software R.³ All the R-code used can be found in appendix B on page 97. There will be no prejudices as to what class of models that will be the better choice. The approach that will be used is therefore, through thorough data analysis to determine the necessary properties of a time series model that is able to describe the observed main features of the index data.

Part of the project work has been done in collaboration with Emil Ahlmann Østergaard, but the model chosen by him in connection with point two on the above list is different from the model that will be presented in this thesis. As a consequence, also the work done in connection with point three and four will differ. Apart from this subsection presenting the problem statement, the two theses have been written independently. In the concluding chapter, a comparison to the results from Emil Ahlmann Østergaards work[25] will be part of the discussion.

³www.r-project.org

CHAPTER 2

Market Indices

In the context of finance, an index is a statistical measure usually expressed in terms of the change from a base value. A financial market index tracks the performance of an imaginary portfolio of securities representing a particular market or a portion of it. Companies like Standard and Poor, Morgan Stanley Capital International (MSCI), and J.P. Morgan specialise in the provision of such indices.

Stock and bond market indices are used to construct index mutual funds and exchange-traded funds (ETFs), whose portfolios mirror the components of an index. These funds allow investors to invest in securities representing market indices. Indices are important benchmarks for fundmanagers, who use them to measure the performance of portfolios such as mutual funds.

Each index has its own calculation methodology. This chapter gives a presentation of the eleven financial market indices included in this thesis and a description of how they are calculated.

2.1 Stock Market Indices

2.1.1 KAXGI: CSE All Share GI Index

The Copenhagen Stock Exchange (CSE) All-Share Index includes all the shares listed on the Copenhagen Stock Exchange. The index is measured in Danish kroner. The aim of the index is to reflect the current status and changes in the market. KAXGI is the Bloomberg short name for the index. The Bloomberg short names will be used throughout the thesis as references to the different indices.

KAXGI is a free float-adjusted market capitalisation weighted index, implying that the index is computed by summing the capitalisation of the free float of all the firms in the index and dividing by some constant base number chosen in order to obtain a desirable scale (initially).¹

$$I(t) = \frac{1}{n_{base}} \sum_{\text{all shares}} n_i S_i(t). \quad (2.1)$$

The free float market capitalisation is the market value $n_i S_i(t)$ (number of shares \times stock price) of all the shares that are freely tradable among investors.

It is a total return index with gross dividends, meaning that returns from dividends are reinvested without inclusion of any tax effects.² The dividends are reinvested at the close of trading on the day the security is quoted ex dividend. The reinvestment is carried out by adjusting the base number to make the index level unaffected by the dividend payout:

$$I(t_{cd}) = \frac{1}{n_{base}} \sum_{\text{all shares}} n_i S_i(t_{cd}) = \frac{1}{\tilde{n}_{base}} \sum_{\text{all shares}} n_i S_i(t_{xd}) = I(t_{xd}), \quad (2.2)$$

for $t_{xd} = t_{cd} + \delta t$, where δt is an infinitesimal time span immediately after the payment of dividend. cd is short for “cum dividend” and xd is short for “ex dividend”.

If the adjustment wasn’t made, then the stock index would decline following a dividend payment by one of its constituents, since the share price itself drops

¹The base date for the CSE All-Share Index is December 31, 1995, with a base value of 100.

²GI in KAXGI is short for gross index.

by an amount equal to the dividend that is paid:

$$S(t_{cd}) = S(t_{xd}) + (1 - \gamma) \textit{Dividend}, \quad (2.3)$$

where γ is the tax rate that makes the gross dividend amount into the net dividend after tax. Hence, $\gamma = 0$ when looking at gross dividend indices.

Since the total value of the portfolio does not decline, when a dividend is paid, the index needs to be adjusted to reflect this. The adjustment corresponds to reinvesting the dividend in all index constituents in proportion to their respective weights. Consequently, the index measures the total return from a portfolio of shares with the same weights as the index disregarding any tax effects.[23]

2.1.2 NDDUE15: MSCI Daily TR Net Europe

An index designed by Morgan Stanley Capital International (MSCI) to measure the equity market performance of the developed markets in Europe. The index is measured in US dollars. The index consists of the market indices of the following 16 countries: Austria, Belgium, Denmark, Finland, France, Germany, Greece, Ireland, Italy, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, and United Kingdom.³

It is a free float-adjusted market capitalisation weighted index just like the KAXGI. It is a total return index with net dividends, implying that returns from dividends are reinvested after subtraction of withholding taxes.⁴ Thus, $\gamma \neq 0$ in equation (2.3) contrary to the gross index. In this way, the index is a measure of the total return from a portfolio of shares with the same weights as the index including tax effects.[18]

2.1.3 NDDUJN: MSCI Daily TR Net Japan

The same as NDDUE15 except for the region considered being Japan instead of Europe. For all the technical specifications see subsection 2.1.2.

³As of May 30, 2011.

⁴The tax rate used is applicable to non-resident institutional investors, who do not benefit from double taxation treaties.

2.1.4 NDDUNA: MSCI Daily TR Net North America

The same as NDDUE15 except for the region considered being Canada and USA. For all the technical specifications see subsection 2.1.2.[19]

2.1.5 NDUEEGF: MSCI Daily TR Net Emerging Markets

The MSCI Emerging Markets Index measures equity market performance of emerging markets. The index is measured in US dollars. The index consists of the following 21 emerging market country indices: Brazil, Chile, China, Colombia, Czech Republic, Egypt, Hungary, India, Indonesia, Korea, Malaysia, Mexico, Morocco, Peru, Philippines, Poland, Russia, South Africa, Taiwan, Thailand, and Turkey.⁵ See subsection 2.1.2 for all the technical specifications.[18]

2.1.6 TPXDDVD: TOPIX Total Return

TOPIX is an abbreviation for Tokyo Stock Price Index. The index tracks the performance of all domestic companies of the Tokyo Stock Exchange's (TSE) First Section. The TSE comprises three sections: The First section is the largest companies, the Second section is for mid-sized firms, and the "Mothers" section is for emerging and high-growth stocks.⁶

TOPIX is measured in Japanese yen. It is a free float-adjusted market capitalisation weighted index like all the other stock indices included in this thesis. Furthermore, it is a total return index with gross dividends just like the KAXGI-index. Thus, the index measures the total return from a portfolio of shares with the same weights as the index disregarding any tax effects. Review subsection 2.1.1 for all the technical details.[2]

⁵As of May 30, 2011.

⁶The First Section includes all the largest companies according to a number of criteria regarding the company's profit, the number of shareholders, the number of free-float shares, the market capitalisation, the ratio between assets and liabilities, the trading volume etc. The exact criteria can be found at <http://www.tse.or.jp/english/rules/transfers/index.html>.

2.2 Bond Market Indices

2.2.1 CSIHYYI: JPMorgan High Yield Bond Index Global

J.P. Morgan High Yield Bond Index Global is an index tracking the performance of a fund investing in fixed income instruments, primarily in bonds with a maximum credit rating of BB+ of issuers in developed countries, primarily corporations and banks. Although it is a global index, more than 90 % of the bonds are American. The index is measured in US dollars.

High yield is a reference to the credit ranking of the bonds included in the index. A high yield bond is classified as having a credit ranking of BB+ or lower. High yield bonds carry a higher risk of default and therefore also offer a higher yield.

It is a capitalisation weighted total return index. Hence, the index measures the total return from an imaginary portfolio of bonds with the same weights as the index. Capitalisation weighting in the context of bonds means that weight is assigned based on the total debt outstanding. This implies that most weight is appointed to the less creditworthy issuers with a larger outstanding debt. An ETF tracking a specific bond market index will therefore have to put most money into the most indebted companies, which of course induces a higher risk.[17]

2.2.2 JPGCCOMP: JPMorgan EMBI Global Diversified

J.P. Morgan Emerging Markets Bond Index Global Diversified measures the total return from US dollar-denominated sovereign bonds issued by a selection of emerging market countries. The index is measured in US dollars. It includes US dollar-denominated Brady bonds⁷, Eurobonds⁸, and traded loans issued by sovereign and quasi-sovereign entities of emerging market countries.⁹ The debt instruments included must have an outstanding face value of at least \$500 million and must meet certain criteria for secondary market trading liquidity.

⁷Bonds issued by the governments of developing countries.

⁸Bonds issued in a currency other than the currency of the country or market in which they are issued.

⁹A sovereign entity is a government. A quasi-sovereign entity is an agency that has government backing.

The emerging market countries included are selected from a combination of World Bank-defined per capita income brackets and each country's debt restructuring history. These two criteria allow the inclusion of a number of countries with a higher credit rating compared to those included in the High Yield Index discussed in the previous section.

The country weights are evenly distributed in the index. This is achieved by limiting the weights of those index countries with larger debt stocks by only including a portion of these countries eligible current face amounts of debt outstanding.[12] It is not possible to be more specific regarding the criteria for countries included, the weighting etc., since J.P. Morgan only shares this information with their clients.

2.2.3 NDEAMO: Nordea Mortgage Index

Nordea Mortgage Index is a benchmark index for Danish mortgage bonds used mainly by portfolio managers to measure their investment performance. It is a total return index measured in Danish kroner. It comprises fixed rate, callable, DKK-denominated mortgage bonds with a term to maturity of more than one year that have been issued by one of the six leading mortgage institutions in Denmark.¹⁰ Fixed rate means that the bonds carry a predetermined coupon rate. On the call date(s), the issuer of the bond has the option to buy back the bond from the bond holder at the predefined call price. The value of this option gives rise to a higher coupon rate and a lower price compared to "straight" bonds.

The weights of the index are determined by minimising the sum of the products of the absolute and the relative deviations from the market weights, subject to a constraint specifying a target for the modified duration.¹¹[1] The modified duration is the weighted average, of the times until the future cashflows are received, calculated using annual compounding. In other words, the weighting should reflect the liquidity of the bonds while maintaining a specified duration.

¹⁰Realkredit Danmark A/S, Nykredit Realkredit A/S, DLR Kredit A/S, BRFKredit a/s, Totalkredit A/S, and Nordea Kredit Realkreditaktieselskab.

¹¹ $\min_{x_i} \sum x_i \frac{(x_i - \bar{x}_i)^2}{\bar{x}_i}$, s.t. $\sum_i x_i \text{ModDur}_i = \text{TargetModDur}$, $\sum_i x_i = 1$, and $x_i \geq 0$, where \bar{x}_i is the market weight and x_i is the index weight.

2.2.4 NDEAGVT: Nordea Government Index

Nordea Government Index measures the total return from treasury bonds issued by the Danish government. The index is measured in Danish kroner. Only bonds with a term to maturity of more than one year are included in the index. Government bonds are often referred to as risk-free, since the risk of default is minimal (this is of course not true for all countries). For the same reason, they are characterised as the safest type of bonds and therefore pay the lowest interest rate. The index is often used as a benchmark for the risk-free rate. The weights of the index constituents are found in the same way as for the mortgage index.[1]

2.3 Money Market Indices

2.3.1 DK00S/N: LIBOR DKK Spot/Next

The London Interbank Offered Rate (LIBOR) is a daily reference rate based on the interest rates at which banks borrow unsecured funds from other banks in the London interbank lending market.[28] LIBOR is computed for ten different currencies, in the present case Danish kroner. It is a short-term interest rate and is calculated for maturities between overnight and twelve months. Spot/next (S/N) refers to the purchase of an asset for delivery the day after the spot delivery date. In most markets the spot delivery date is two business days after the day the transaction was made. So the spot/next rate matures in three days.

The LIBOR is released every business day just after 11 am London time by the British Bankers' Association in conjunction with Reuters. There are ten LIBOR panels, one for each of the ten currencies for which the rate is determined. Each panel is composed of at least eight banks, chosen for their reputations and their perceived expertise in a given currency. The LIBOR is calculated for each currency as the interquartile mean of the daily deposit rates reported by the banks in the panel for the given currency.¹² The average rates at which these banks say they would lend to one another is taken as an indication of the health of the banking system.[26]

The core of the money market consists of interbank lending. The instruments

¹²The interquartile mean is the mean of the middle 50%.

used in the lending process, like commercial papers, repurchase agreements, and similar, are often benchmarked to the LIBOR for the appropriate term and currency.[29] LIBOR has become the premier short-term interest rate in the European money market, and it serves as a reference rate for a wide range of transactions.[2, p. 60] The Danish LIBOR spot/next will serve as the link to the money market in this project.

CHAPTER 3

Data Analysis

The objective of this chapter is to explore the main features of the dataset and through that form the basis of the following chapter, where the data will be modelled. The main features being stationarity properties, distribution, auto- and cross-correlation. The data is displayed and there are drawn general conclusions regarding the distribution and the correlations within and between the series through application of suitable transformations and statistical tests.

The dataset contains 3291 observations of the value of eleven different financial market indices recorded on a daily basis (weekdays only) in the period between the 1st of January 1999 and the 12th of August 2011. On bank holidays the closing price of the previous weekday is quoted. As noted in the introduction, observations are only available from the 16th of June 2003 and onwards for the LIBOR-index.

The chapter begins with a reconstruction of the first 24 months' observations of the NDUEEGF-index.

3.1 Data Reconstruction




Figure 3.1: The original and the reconstructed index data for the first 24 months for the NDUEEGF-index.

Values for the first 24 months have only been recorded monthly for the NDUEEGF-index. All the values in-between the first of each month are simply repetitions. This appears from the black curve in figure 3.1. These data are not applicable for the analysis, since they would lead to a lot of zero returns and a few very large returns.

The data for the first 24 months are instead reconstructed using linear interpolation between the first in each month adding some normally distributed noise. It is evident from figure 3.1, that doing linear interpolation between more than two successive months at a time, would lead to increasing trends in months, where the overall index level actually has gone down and vice versa.

The standard deviation of the noise is decided from the standard deviation of the dataset containing the difference between the first value in each of the 24 months. This quantity is estimated as 8.0. This standard deviation of the monthly variation can in some sense be interpreted as the standard deviation of the sum of the daily variations. Hence, the standard deviation of the daily variations can be estimated as $8/\sqrt{22} \approx 1.7$, assuming there are roughly 22 bankdays in each month. This estimate is based on the assumption that the daily returns are uncorrelated, which is in fact not true. The contribution from the covariance is ignored for the time being, since the estimate is very rough anyway.

For $T \in [T_i, T_{i+1}]$ X_T is interpolated as

$$X_T = \frac{T - T_i}{T_{i+1} - T_i} X_{T_{i+1}} + \frac{T_{i+1} - T}{T_{i+1} - T_i} X_{T_i} + \varepsilon_T, \quad \varepsilon_T \sim \mathcal{N}(0, 1.7^2). \quad (3.1)$$

The application of linear interpolation results in constant returns within each month, and consequently the standard deviation of the returns will be nil. Adding normally distributed noise with the estimated standard deviation of the daily variations leads to returns that have the correct overall magnitude of variation. The shortfall of the used approach is that the standard deviation of the

daily returns is constant for the first 24 months, which is not consistent with reality. But the simple reconstruction is sufficient for the use of the data in this work. The reconstructed data is displayed together with the original data in figure 3.1.

3.2 Considering Stationarity

The daily values of the eleven indices are shown in figure 3.2 as a function of time (black lines). The red lines are recursively estimated mean values, meaning that the value plotted at time t is the mean of all the values prior to time t weighted equally.

None of the series are stationary, since the mean is either growing as a function of time (e.g. NDEAGVT), has strong local trends (e.g. NDDUJN), or both. The four bond indices near the bottom of figure 3.2 are seen to be much less volatile than the stock indices.

Figure 3.2 also reveals a remarkable resemblance between most of the indices. The stock indices and the first two of the bond indices tend to move up and down at almost the same points in time. The two Nordea indices do not show the same distinct fluctuations. Also, the LIBOR-index, at the bottom of figure 3.2, seems to move a bit out of sync with the rest of the indices.

The magnitude of the fluctuations are not the same for all the indices, and for instance the TOPIX-index has not recovered from the financial crisis in 2008 in the same way the other stock indices have. Moreover, the CSYHYI- and JGCCOMP-indices do not have a distinct peak around year 2000, like the stock indices have. There are after all certain differences, but the resemblance is still remarkable.

3.2.1 Log Returns

In order to obtain stationarity a transformation is needed. A difference is suitable to get rid of the growing mean value that most of the series possess. But instead of differencing the series directly, a log-transformation is applied prior

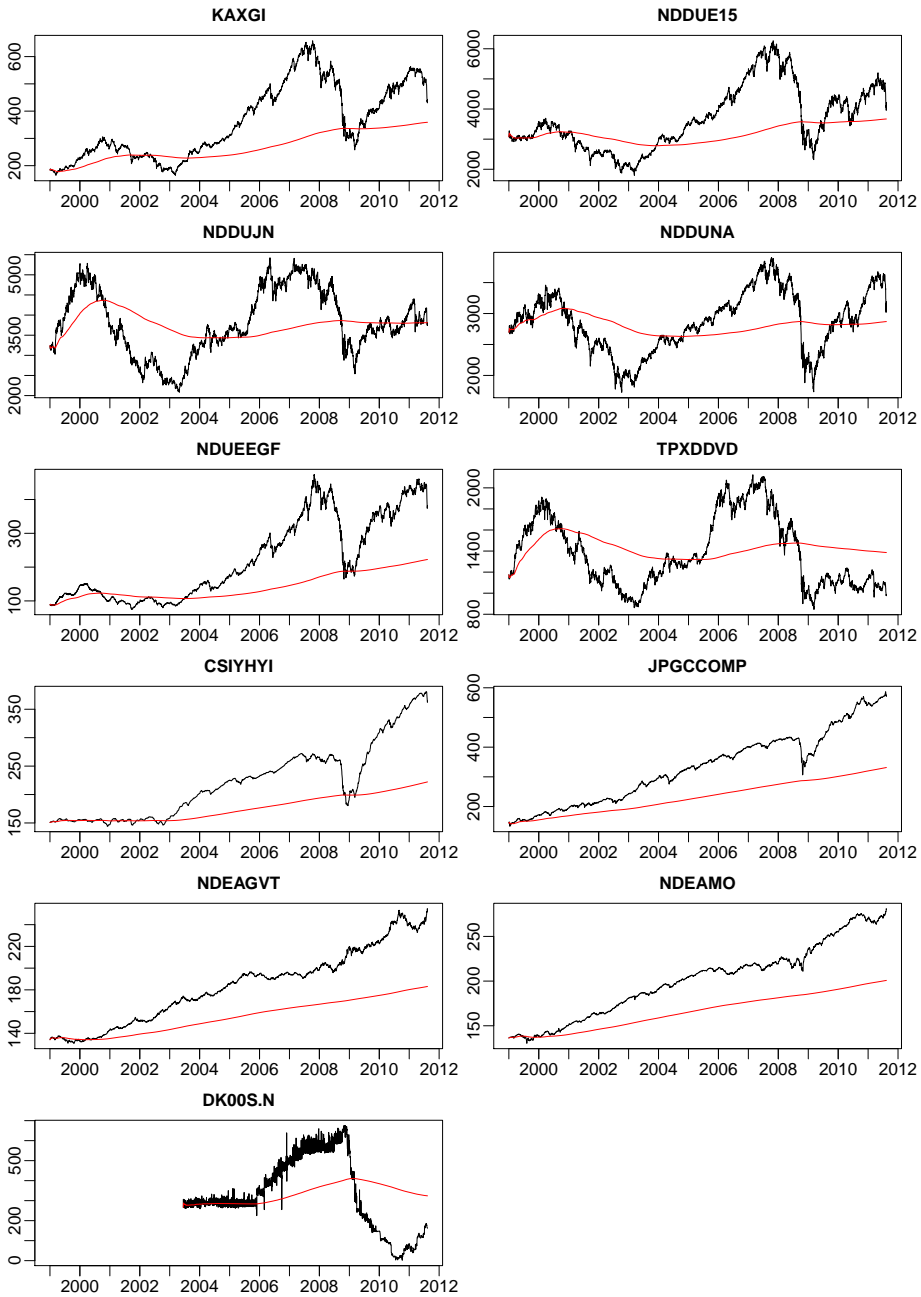


Figure 3.2: The daily value of the eleven indices as a function of time (black lines) together with the recursively estimated mean values (red lines).

to taking the difference. This leads to the log return defined as

$$r_t = \ln P_t - \ln P_{t-1} = \ln \frac{P_t}{P_{t-1}}. \quad (3.2)$$

Taking the natural logarithm before the difference results in a key figure of approximately the same scale for all the indices. In the raw index data some of the stock indices encountered values of up to 6000, while the bond indices lay around 200. Taking the logarithm ensures a convenient rescaling that narrows the gap between the indices. This appears from figure 3.3, where the daily log return values are plotted as a function of time.

Performing a log-transformation is always a good choice when interested in small changes in large numbers. For returns less than 10%, the log return actually approximates the discrete return very well anyway, since it is the first order Taylor approximation.¹[6] The log return values are seen to be mean stationary, since they fluctuate around a constant mean level close to zero for all the indices.

It appears from figure 3.3 that the log returns are much more volatile in some periods than others. This effect, as noted by Mandelbrot in 1963, is referred to as volatility clustering. That is, after a large price change (positive or negative) a large price change tends to occur, while small changes tend to be followed by small changes.[6]

Figure 3.4 shows the standard deviation of the eleven log returns series estimated recursively as a function of time using exponential weights (black lines). The weight assigned to the t 'th observation at time T is $\lambda_{t,T} = c \cdot 0.95^{T-t}$, where $t \leq T$ and $c = \frac{1-0.95}{1-0.95^T}$ is a normalising factor. The red line is a smoothing of the recursively estimated standard deviation showing the local mean level of volatility. The smoothing is done using locally weighted polynomial regression in R with a span of $1/5$, which is the proportion of points that influence the smooth at each value.

The standard deviation is far from constant for any of the log return series as noted in relation to figure 3.3. Furthermore, the figure shows that large peaks in the standard deviation tend to occur in periods, where the local mean level of volatility is above the overall mean level, whereas there are few large peaks in the periods, where the local mean level of variation is low. The standard deviation

¹ $r_t = \ln \frac{P_t}{P_{t-1}} = \ln(1 + R_t) = \ln(1) + R_t - \frac{R_t^2}{2!} + \frac{R_t^3}{3!} + \mathcal{O}(R_t^4) \approx R_t$ for discrete returns R_t close to zero.

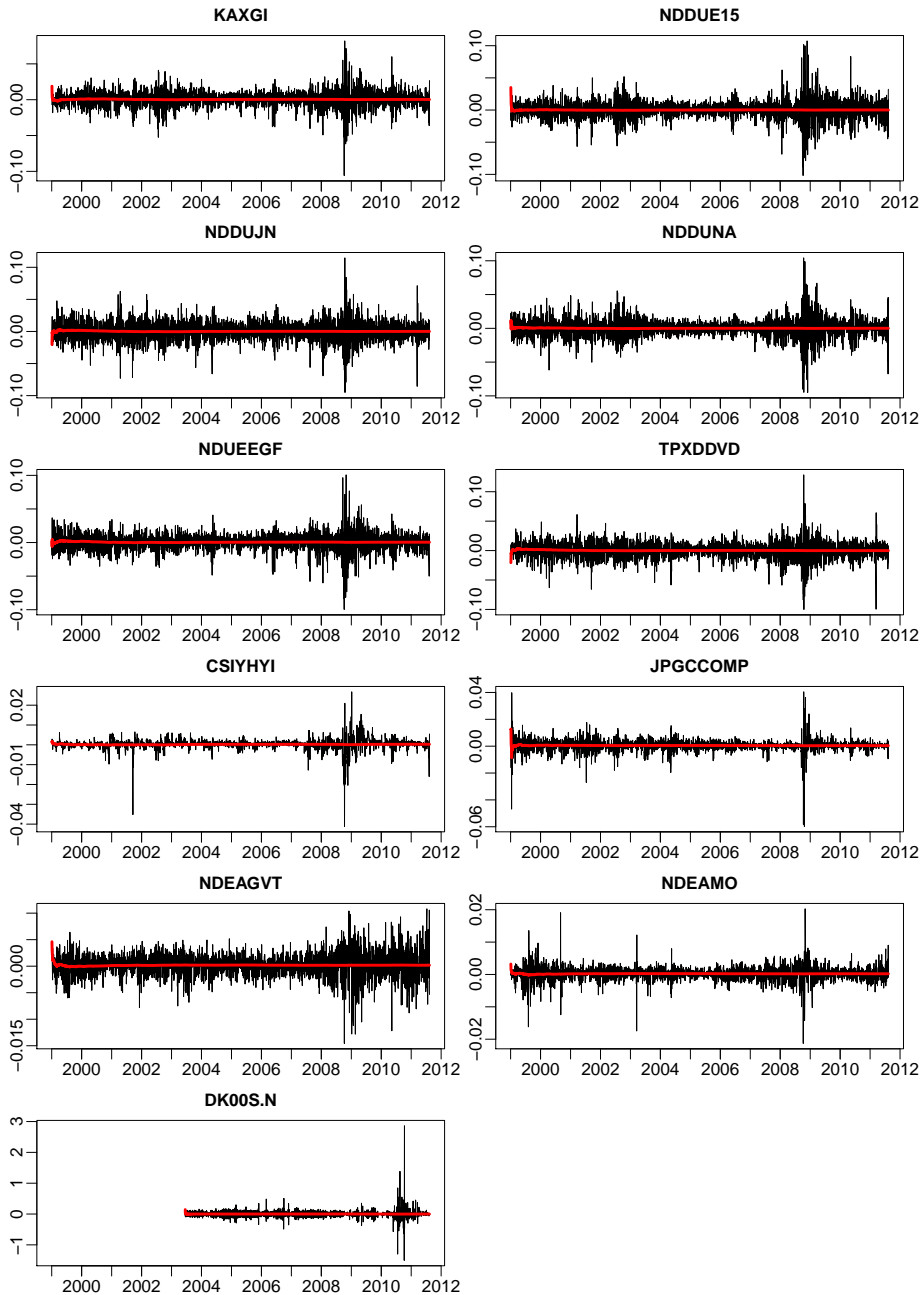


Figure 3.3: The daily log return values of the eleven indices as a function of time. The red lines are recursively estimated mean values applying uniform weights to all past observations.

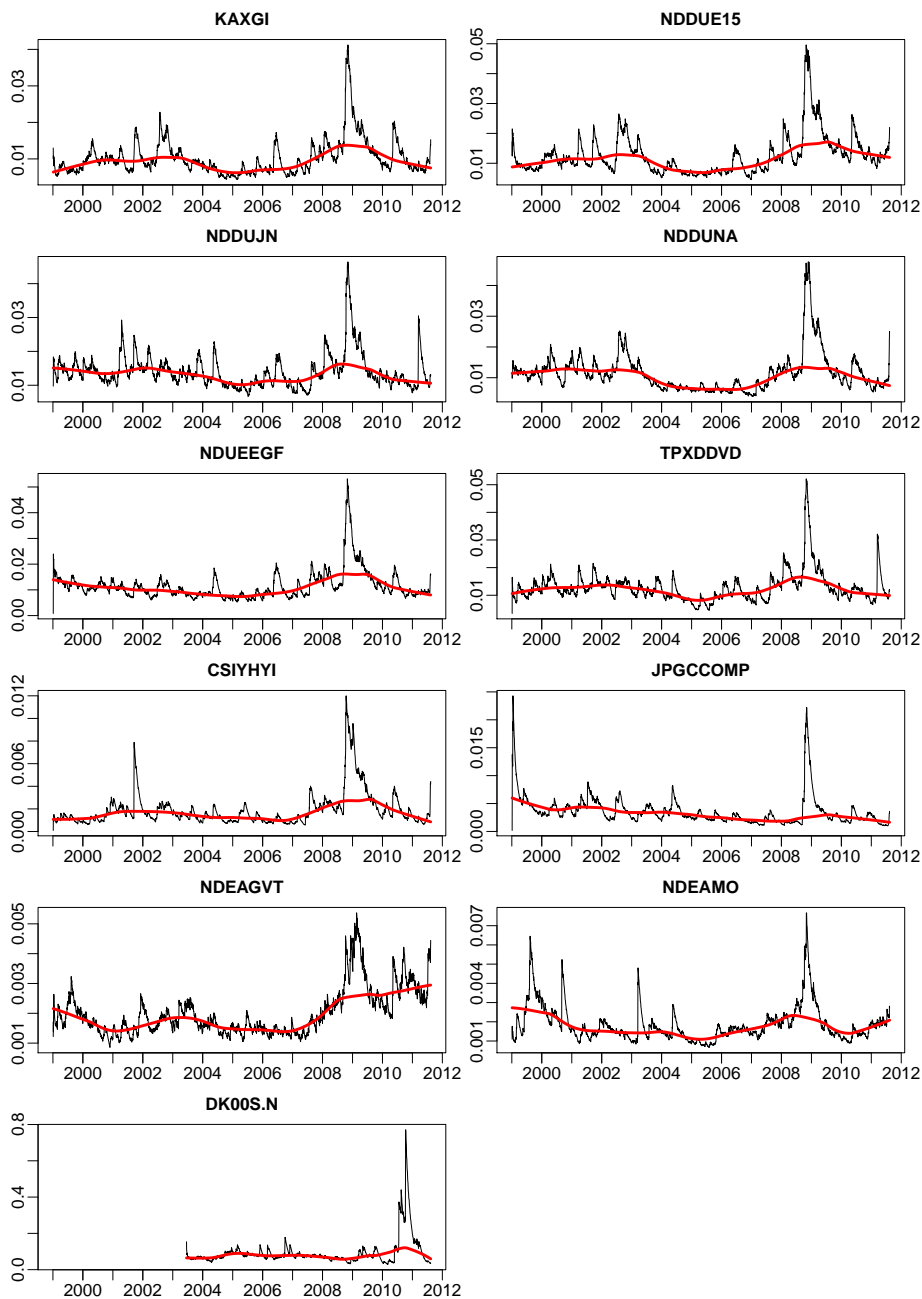


Figure 3.4: The standard deviation of the log returns estimated recursively as a function of time using exponential weights. The red line is a smoothing of the recursively estimated standard deviation showing the local mean level.

of the indices themselves shows the same behaviour as for the log returns. This appears from figure A.1 in appendix A on page 89. Such processes with non-constant variance are said to be conditionally heteroscedastic, when there is no evidence that the overall variance, or the unconditional variance, is changing. This is seen to be the case.

The resemblance between the indices in figure 3.2 is not as evident from figure 3.3 and 3.4, but it is still there. The most striking similarity is obviously the large peak in volatility around 2009 that all the log return series have in common. The volatility of the LIBOR-index does however first peak in late 2010. From a comparison of figure 3.2 and 3.4 it is seen that the largest peaks in volatility occur, when the indices are at a local minimum. In periods where the indices are growing, the volatility is found to be below average.

The log return processes are only weakly stationary of order one, since the mean values are stationary, but the variances (the second moments) are still not stationary.[14]

3.3 The Distribution of Data

Density histograms of the daily log return values of the eleven indices are displayed in figure 3.5 together with the theoretical density of a normal distribution with the same mean and variance as the data (red lines). Except for CSIHYYI and LIBOR, all the log return distributions seem somewhat close to the normal distribution. It should be noted that the raw index values have less resemblance with the normal distribution compared to the log returns. This is evident from the histograms of the index data in figure A.2 in appendix A on page 89. For that reason, the focus will be solely on the distribution of the log returns.

The mean values of the log returns appear to be zero, but a t-test rejects this null hypothesis at a 5%-significance level for the four bond indices that all have positive mean returns. This is evident from table 3.1, where the estimated mean values are listed together with the p -values of the t-test. Only the TPXDDVD and the LIBOR-index have negative mean returns. These conclusions agree with what appeared from the recursive mean values in figure 3.2 on page 18.

The estimated standard deviations in table 3.1 also support the conclusions drawn from figure 3.2 on page 18 in that the four bond indices are less volatile

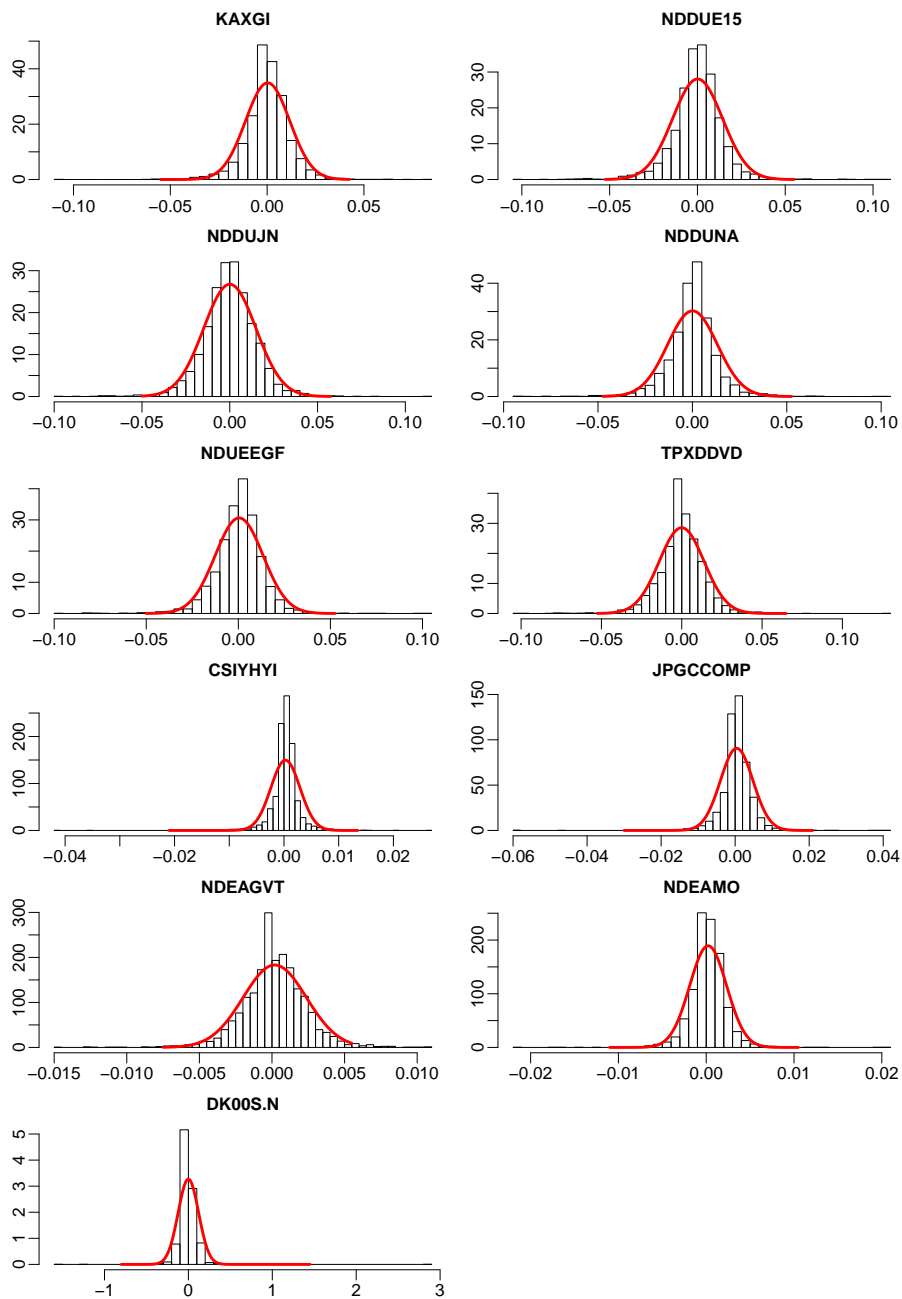


Figure 3.5: Density histograms of the daily log return values of the eleven indices together with the theoretical density of a normal distribution with the same mean and variance as the log return data (red lines).

	$\frac{\hat{\mu}}{10^{-4}}$	p -value	$\frac{\hat{\sigma}}{10^{-2}}$	$\hat{\gamma}_1$	$\hat{\gamma}_2$	JB
KAXGI	2.66	0.18	1.14	-0.39	6.45	5800
NDDUE15	0.96	0.70	1.42	-0.08	7.07	6854
NDDUJN	0.44	0.87	1.49	-0.15	3.96	2157
NDDUNA	0.51	0.82	1.32	-0.23	7.62	7993
NDUEEGF	4.44	0.05	1.30	-0.44	7.35	7510
TPXDDVD	-0.57	0.82	1.40	-0.36	6.81	6436
CSIYHYI	2.68	0.00	0.27	-2.54	40.4	227623
JPGCCOMP	4.20	0.00	0.44	-1.64	33.0	150322
NDEAGVT	1.92	0.00	0.22	-0.28	3.60	1821
NDEAMO	2.20	0.00	0.21	-0.35	13.9	26710
DK00S.N	-2.33	0.93	12.2	5.44	170	2587839

Table 3.1: Estimated mean value, $\hat{\mu}$, of the daily log returns, associated p -value of t-test of the null hypothesis: $\mu = 0$, estimated standard deviation, $\hat{\sigma}$, estimated skewness, $\hat{\gamma}_1$, estimated excess kurtosis, $\hat{\gamma}_2$, and the Jarque-Bera test statistic, JB .

than the stock indices, and the LIBOR-index is much more volatile than both the stock and the bond indices.

The distributions appear a bit skew compared to the normal distribution. This observation is supported by the estimated skewness,

$$\hat{\gamma}_1 = \frac{\frac{1}{N} \sum_{t=1}^N (r_t - \hat{\mu}_r)^3}{\hat{\sigma}_r^3}, \quad (3.3)$$

in table 3.1. In fact, the distributions are all left skew with exception of the LIBOR-index, which is very right skew.

What is more apparent from figure 3.5 is that all the distributions have more mass centered around the mean and in the tails compared to the normal distribution. Such distributions with excess kurtosis,

$$\hat{\gamma}_2 = \frac{\frac{1}{N} \sum_{t=1}^N (r_t - \hat{\mu}_r)^4}{\hat{\sigma}_r^4} - 3, \quad (3.4)$$

are said to be leptokurtic.² It appears from table 3.1 that the estimated excess kurtoses for the daily log returns are all much greater than zero.

²The normal distribution has kurtosis 3.

3.3.1 Testing for normality

The Jarque-Bera test statistic

$$JB = N \left(\frac{\hat{\gamma}_1^2}{6} + \frac{\hat{\gamma}_2^2}{24} \right) \quad (3.5)$$

is used to test whether the estimated skewness and excess kurtoses are significant. The Jarque-Bera test statistic is asymptotically χ^2 -distributed with two degrees of freedom.[6] The 95%-quantile for the χ^2 -distribution with two degrees of freedom is 6. The probability of the log return data being normally distributed is therefore nil according to the Jarque-Bera test, since all the estimated test statistics are much larger than 6. This conclusion is also backed by a Shapiro-Wilk test for normality returning p -values of less than 2.2×10^{-16} for all the log return series, which rejects any resemblance with the normal distribution. This is above all caused by the excess kurtoses.

3.3.2 Outliers

It is a general characteristic of financial returns, that there are too many extreme values compared to the normal distribution, which results in the very large kurtoses observed. It is however important to distinguish between outliers and extreme observations in this connection; extreme observations deviate considerably from the group mean, but may still represent meaningful conditions, and thus should not be disregarded. It is important to include these extreme observations in the model estimation in order for the scenarios to be as realistic as possible. The scenarios should after all preferably reflect the possibility of such extreme events, since they are seen to occur.

The log return graph on page 20 is a good means of detecting outliers, since the presence of outliers would lead to sudden spikes in the returns. There is no reason to investigate the observations, where all the graphs spike around the same point in time like late 2008, since these simultaneous peaks are most likely caused by extreme conditions in the financial markets and therefore should not be disregarded.

The two Japanese indices (NDDUJN and TPXDDVD) are seen to have spikes in the first part of 2011, which are not recovered in the rest of the indices. The large fluctuations took place in the period from the 14th March to the 17th March, at the time where Japan was hit by earthquakes and a tsunami.

The CSIYHYI has a very large negative deviation in the last half of 2001. It is a large decline from Friday the 14th September to Monday the 17th September. This large decline is probably caused by the events on the 11th September in USA that year.

There are no obvious explanation for the spikes JPGCCOMP shows right at the beginning of the period as well as the spikes NDEAMO has in 2000 and 2003, but the amplitude of these spikes is smaller than the spikes that occurred during the latest financial crisis. There is no need to remove these large returns, when it is not known with certainty that they are outliers.

Originally the NDEAGVT-index had some unexplainable peaks in 2002. This turned out to be caused by mistakes by the people manually entering the data. These outliers have been corrected in the presented data.

The LIBOR-index shows large fluctuations in the last part of 2010. The interest rate was very low at that point in time and therefore even minor absolute changes resulted in large deviations in the index, since log return is a relative measure (see eq. (3.2)).

3.3.3 Robust Moments

The presence of this many extreme events is what causes the tests for normality to reject any resemblance with the normal distribution, even though there appeared to be some resemblance based on the histograms on page 23. With more than three thousand observations it doesn't take more than a few extreme observations to reject normality, since the applied tests are very sensitive to outliers.

The sensitivity to outliers is recognised from the sample mean, and since the conventional measures of skewness and kurtosis are essentially based on sample averages, they are also sensitive to outliers. Moreover, the impact of outliers is greatly amplified in the conventional measures of skewness and kurtosis due to the fact that they are raised to the third and fourth powers.

If the most extreme observations were disregarded, the similarity with the normal distribution would be more convincing. Therefore it would make sense to estimate the moments of the distributions using measures that are more robust towards the presence of extreme values.

The median could be used instead of the mean as a measure of location. The interquartile range could be used as a measure of the dispersion instead of the standard deviation. As an alternative measure of skewness, the following estimate, which is also based on quantiles, could be applied:

$$\hat{\gamma}_1^{robust} = \frac{Q_3 + Q_1 - 2Q_2}{Q_3 - Q_1}, \quad (3.6)$$

where $Q_1 = F^{-1}(0.25)$ is the lower quartile, $Q_2 = F^{-1}(0.5)$ is the the median etc., and F is the cumulative distribution function of the process in question. This robust measure of skewness is zero for a symmetric distribution like the normal distribution, one for extreme right skew distributions, and minus one for extreme left skew distributions.

Finally, as a measure of excess kurtosis the following estimate could be applied:

$$\hat{\gamma}_2^{robust} = \frac{F^{-1}(0.975) - F^{-1}(0.025)}{Q_3 - Q_1} - 2.91, \quad (3.7)$$

which is zero for the standard normal distribution.[9]

The robust estimates for the log return series appear from table 3.2. When comparing these estimates to the non-robust, conventional measures from table 3.1 on page 24, it is first noticed that the measures of location and dispersion are of the same magnitude. None of the medians are negative, and especially for the NDUEEGF-index, the median is much more positive than the mean value (4.44×10^{-4}). Furthermore, the dispersion of the LIBOR-index is more than halved, when quantified by the interquartile range instead of the standard deviation, due to the many extreme returns observed for this index.

The estimated skewness is close to zero for all the indices. When looking at the robust measure, there are actually more right skew than left skew distributions. More importantly, the estimated kurtoses are much smaller than the non-robust estimates in table 3.1 on page 24. All the distributions still have excess kurtosis, but it is only a small amount.

The best description of the log returns is probably obtained by some mixture distribution, where the majority of the returns are normally distributed, and then there are a few extreme returns from a completely different distribution.[9] This makes perfectly sense in that most of the observed extreme returns could be explained by a specific event like 9/11, the bankruptcy of Lehmann Brothers, the Tsunami in Japan etc., and most of these events were caused by outer circumstances.

	$\frac{Q_2}{10^{-4}}$	$\frac{Q_3 - Q_1}{10^{-2}}$	$\frac{\hat{\gamma}_1^{robust}}{10^{-2}}$	$\hat{\gamma}_2^{robust}$
KAXGI	1.59	1.11	7.12	1.25
NDDUE15	4.41	1.40	1.06	1.16
NDDUJN	1.25	1.68	-0.11	0.54
NDDUNA	4.06	1.19	-3.81	1.47
NDUEEGF	11.7	1.32	-6.04	0.84
TPXDDVD	0.00	1.43	6.20	0.82
CSIYHYI	3.84	0.17	2.98	2.78
JPGCCOMP	4.53	0.33	5.45	1.95
NDEAGVT	1.17	0.23	9.41	0.77
NDEAMO	1.90	0.20	9.89	1.15
DK00S.N	0.00	6.20	-6.56	2.54

Table 3.2: Robust estimates of location, Q_2 , dispersion, $Q_3 - Q_1$, skewness, $\hat{\gamma}_1^{robust}$, and excess kurtosis, $\hat{\gamma}_2^{robust}$, of the daily log returns.

The conclusion is, that even though there is no doubt that the log returns are not normally distributed as a whole, it is not pointless to assume some resemblance with the normal distribution after all.

3.3.4 The Distribution of Weekly Log Returns

Looking at weekly log returns instead of daily log returns doesn't change the conclusion. From table 3.3 it appears that the median and the dispersion is slightly larger for the weekly log returns. The distributions are a more skew compared to the daily log returns, and all of them, except for JPGCCOMP, to the left.

The largest difference is that only CSIYHYI, JPGCCOMP, and LIBOR have excess kurtoses above one, when using the robust estimate (3.7). The reason being that the weekly sampling smooths the impact of the extreme events.

Density histograms of the weekly log returns can be found in figure A.3 on page 92 in appendix A. The histograms support the conclusion, that also the weekly log returns have resemblances with the normal distribution.

	$\frac{Q_2}{10^{-4}}$	$\frac{Q_3-Q_1}{10^{-2}}$	$\frac{\hat{\gamma}_1^{robust}}{10^{-2}}$	$\hat{\gamma}_2^{robust}$
KAXGI	42.0	2.82	-10.6	0.58
NDDUE15	36.1	3.17	-14.6	0.83
NDDUJN	4.48	3.57	-0.93	0.19
NDDUNA	18.8	2.74	-8.69	0.95
NDUEEGF	44.2	3.36	-7.38	0.59
TPXDDVD	20.0	3.65	-3.20	-0.11
CSIYHYI	21.4	0.67	-11.7	2.57
JPGCCOMP	27.6	0.98	0.62	1.30
NDEAGVT	10.8	0.59	-2.81	0.55
NDEAMO	13.8	0.56	-2.92	0.80
DK00S.N	0.00	6.48	-13.9	6.77

Table 3.3: Robust estimates of location, Q_2 , dispersion, $Q_3 - Q_1$, skewness, $\hat{\gamma}_1^{robust}$, and excess kurtosis, $\hat{\gamma}_2^{robust}$, of the weekly log returns.

3.4 Correlations

There are two types of correlation to consider, namely intra- and inter-series correlation. Autocorrelation is the cross-correlation of a time series with itself. It is a way of finding repeating patterns in data. The autocorrelation of a stochastic process describes the correlation between values of the process at different points in time as a function of the two times. If the process is stationary, then the autocorrelation will only be a function of the time difference, usually referred to as the lag, k .

The autocovariance of a stationary time series X_1, \dots, X_N can be estimated as

$$C_{XX}(k) = \frac{1}{N} \sum_{t=1}^{N-|k|} (X_t - \bar{X})(X_{t+|k|} - \bar{X}), \quad (3.8)$$

where $|k| = 0, 1, \dots, N-1$.

From this, the autocorrelation of the process can be estimated as

$$\hat{\rho}(k) = \frac{C(k)}{C(0)}, \quad (3.9)$$

implying that $\hat{\rho}(0) \equiv 1$.

For a white noise process $\hat{\rho}(k)$ is asymptotically normally distributed with mean

value and variance

$$E[\hat{\rho}(k)] \simeq 0, \quad k \neq 0, \quad (3.10)$$

$$\text{Var}[\hat{\rho}(k)] \simeq \frac{1}{N}, \quad k \neq 0. \quad (3.11)$$

Hence, an approximate 95% confidence interval for the null hypothesis of independency is $\pm 2/\sqrt{N}$. [14]

3.4.1 Autocorrelation in Log Returns

The blue bins in figure 3.6 are the estimated autocorrelation for the daily log returns for the first ten lags. The dashed blue lines make up an approximate 95% confidence interval for the null hypothesis of no correlation. For most of the indices, there are seen to be significant correlation at lag 1, and for some of the series there are also significant correlation at lag 2. The CSIYHYI-index is different from the rest since it has significant positive correlation at all ten lags.

In order to find a time series model that fits the series, it is beneficial to get rid of some of the observed autocorrelation. Since the objective is to generate scenarios five years forward in time, there is no point in looking at daily values. If the data is instead sampled on a weekly basis, then the autocorrelation function looks like the red bins in figure 3.6. Only the CSIYHYI- and the LIBOR-index have significant correlation at lag 1, when looking at weekly log returns. Some of the weekly log return series appear to have significant correlation at lag 2, but it should be remembered that the dashed lines only make up a 95% confidence interval after all. It is only the CSIYHYI- and the LIBOR-index that could not reasonably be assumed to be independent, when looking at weekly log returns.

Though the majority of the weekly log returns series don't exhibit any significant autocorrelation when looking at the entire period, it is relevant to investigate, whether the autocorrelation is constant, or if it is significant in some shorter periods.

The blue lines in figure 3.7 are the estimated autocorrelations at lag 1 for the weekly log return series estimated for a rolling window of 26 weeks. The red lines are the same estimates for the squared log returns. If the weekly log returns are a GARCH(p, q)-process, as often assumed in econometrics because of the observed volatility clustering (see section 4.2 on page 51), then the squared residuals follow an ARMA(m, p)-process, with $m = \max(p, q)$. [6]

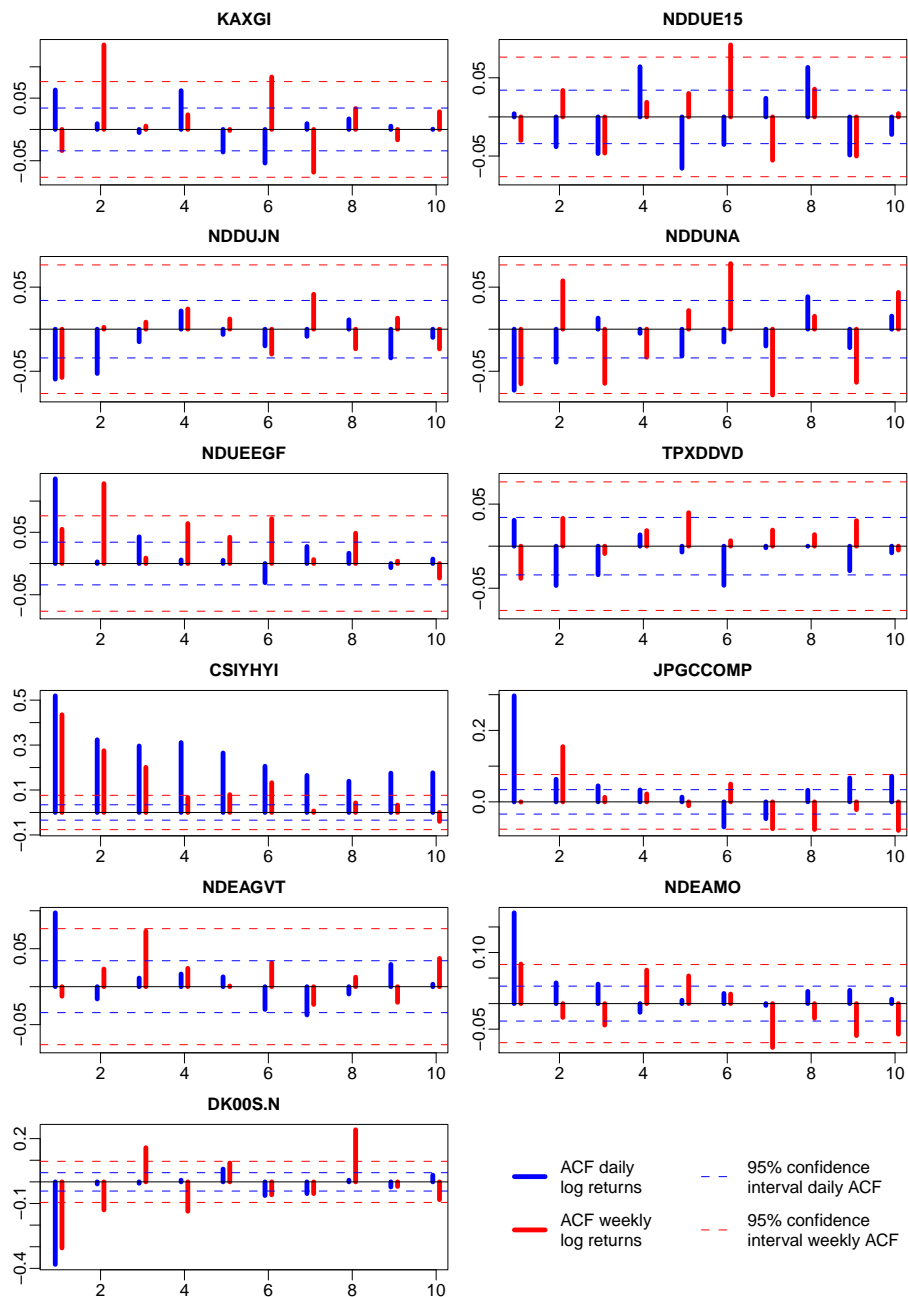


Figure 3.6: The autocorrelation function for the daily log returns (blue) and the weekly log returns (red) estimated for lag 1 through 10. The dashed lines make up an approximate 95% confidence interval for the null hypothesis of no autocorrelation.

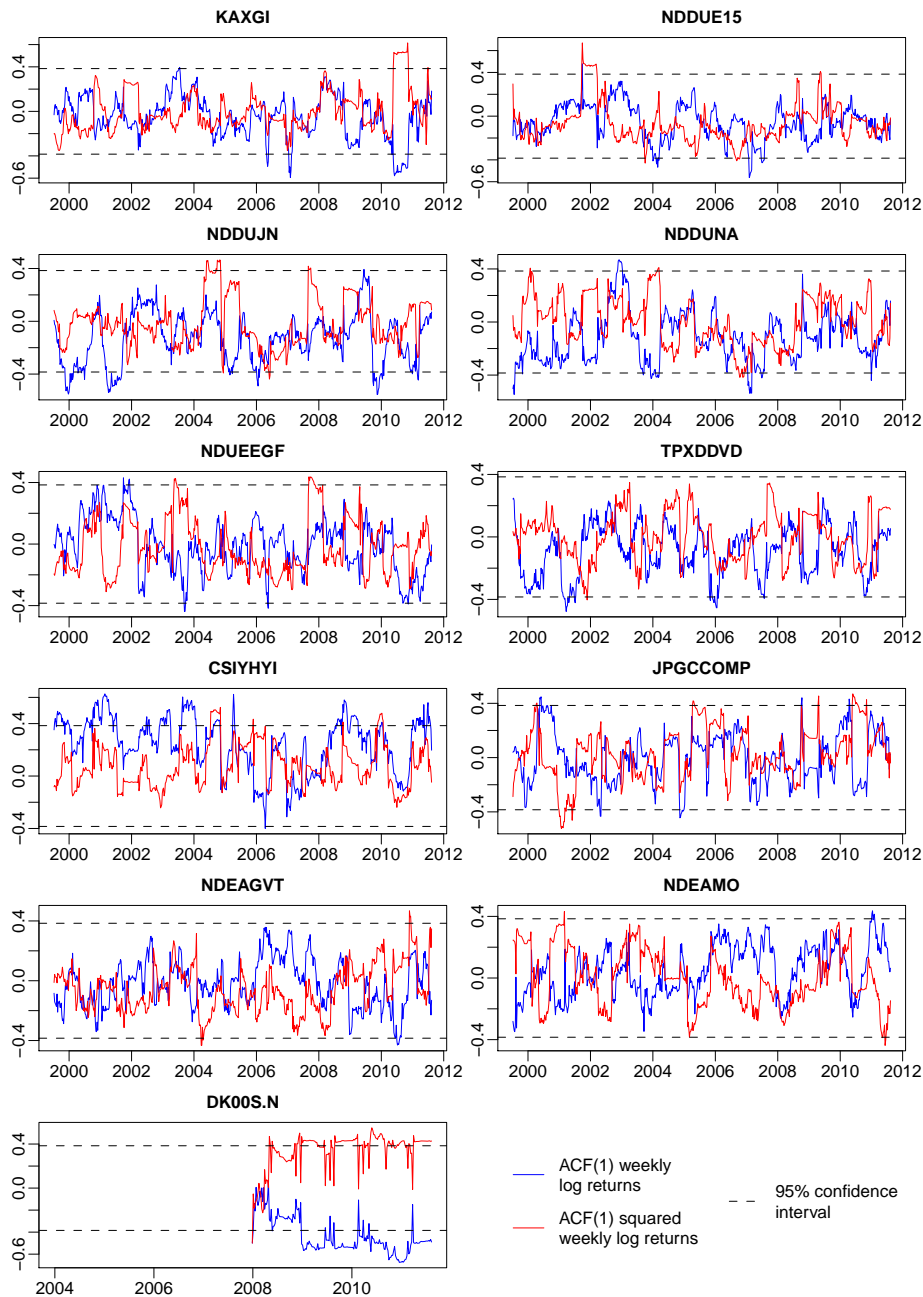


Figure 3.7: The autocorrelation function at lag 1 for the weekly log returns (blue) and the squared weekly log returns (red) estimated using a rolling window of 26 weeks. The dashed lines make up a 95% confidence interval for the null hypothesis of independency.

The autocorrelations at lag 1 are seen to be far from constant. For most of the indices, the ACF is seen to fluctuate in-between the dashed lines that make up an approximate 95% confidence interval most of the time. Some of the indices do, however, exhibit significant autocorrelation in shorter periods. For instance around year 2010 some of the indices appear to have significant autocorrelation. The CSIYHYI- and the LIBOR-index are different in that they have significant first order autocorrelation in a large part of the period, which is consistent with the estimates in figure 3.6. The conclusion remains the same, whether it is the weekly log returns or the squared weekly log returns that are focused on, except for CSIYHYI, where the squared returns are less significant.

It is evident from figure 3.7 that the autocorrelations are fluctuating, but looking at weekly data, it could reasonably be assumed that the autocorrelation is insignificant for the entire period for most of the indices except for CSIYHYI and LIBOR.

3.4.2 Cross-correlation

The cross-correlation of two stochastic processes describes the correlation between values of the processes at different points in time as a function of the two times. If the processes are stationary, then the cross-correlation will only be a function of the time-lag.

The cross-covariance for lag $k = 0, 1, \dots, N-1$ between two time series X_1, \dots, X_N and Y_1, \dots, Y_N can be estimated as

$$C_{XY}(k) = \frac{1}{N} \sum_{t=1}^{N-k} (X_t - \bar{X})(Y_{t+k} - \bar{Y}) \quad (3.12)$$

and similarly for negative lags

$$C_{XY}(-k) = \frac{1}{N} \sum_{t=1}^{N-k} (X_{t+k} - \bar{X})(Y_t - \bar{Y}). \quad (3.13)$$

Based on the estimated cross-covariance function, the cross-correlation function is estimated as

$$\hat{\rho}_{XY}(k) = \frac{C_{XY}(k)}{\sqrt{C_{XX}(0)C_{YY}(0)}}. \quad (3.14)$$

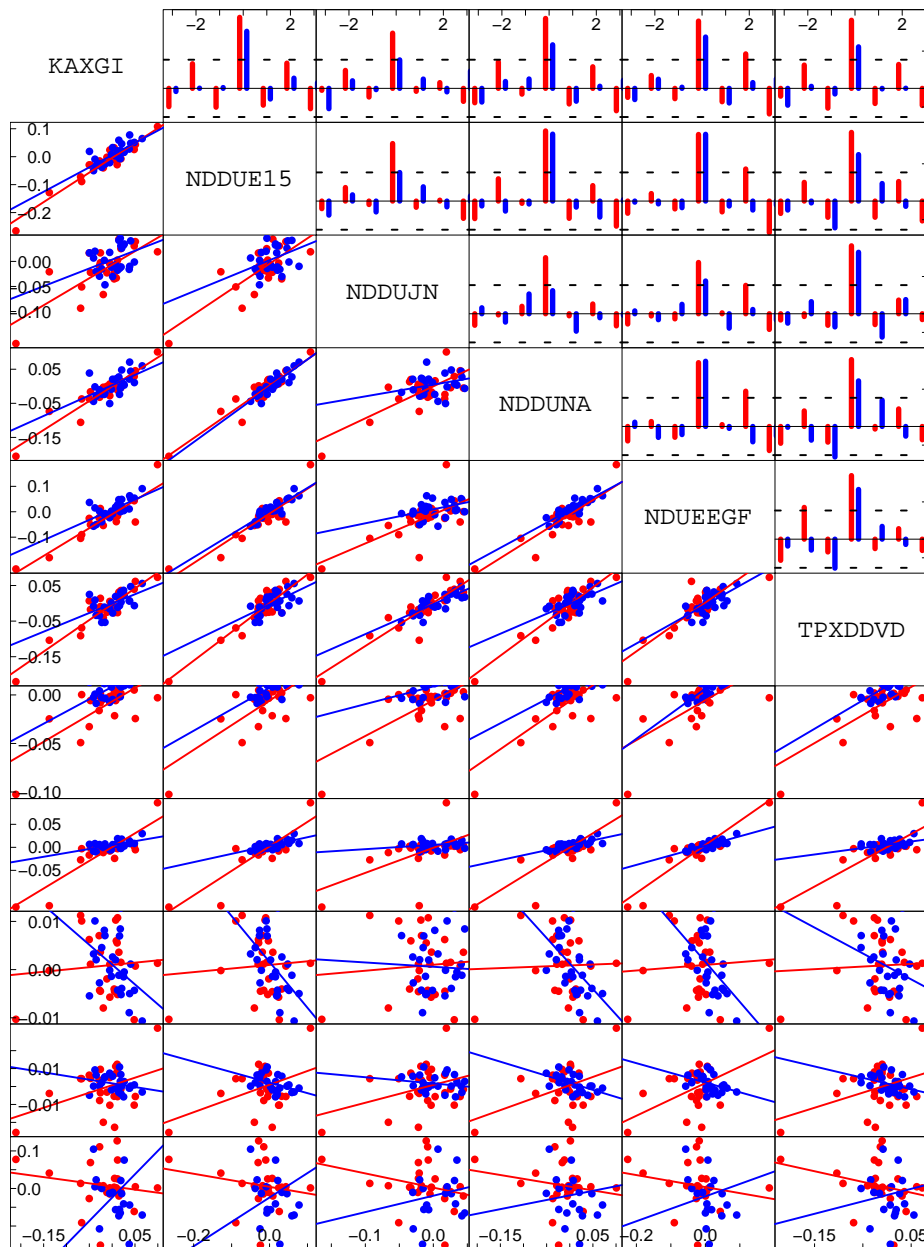
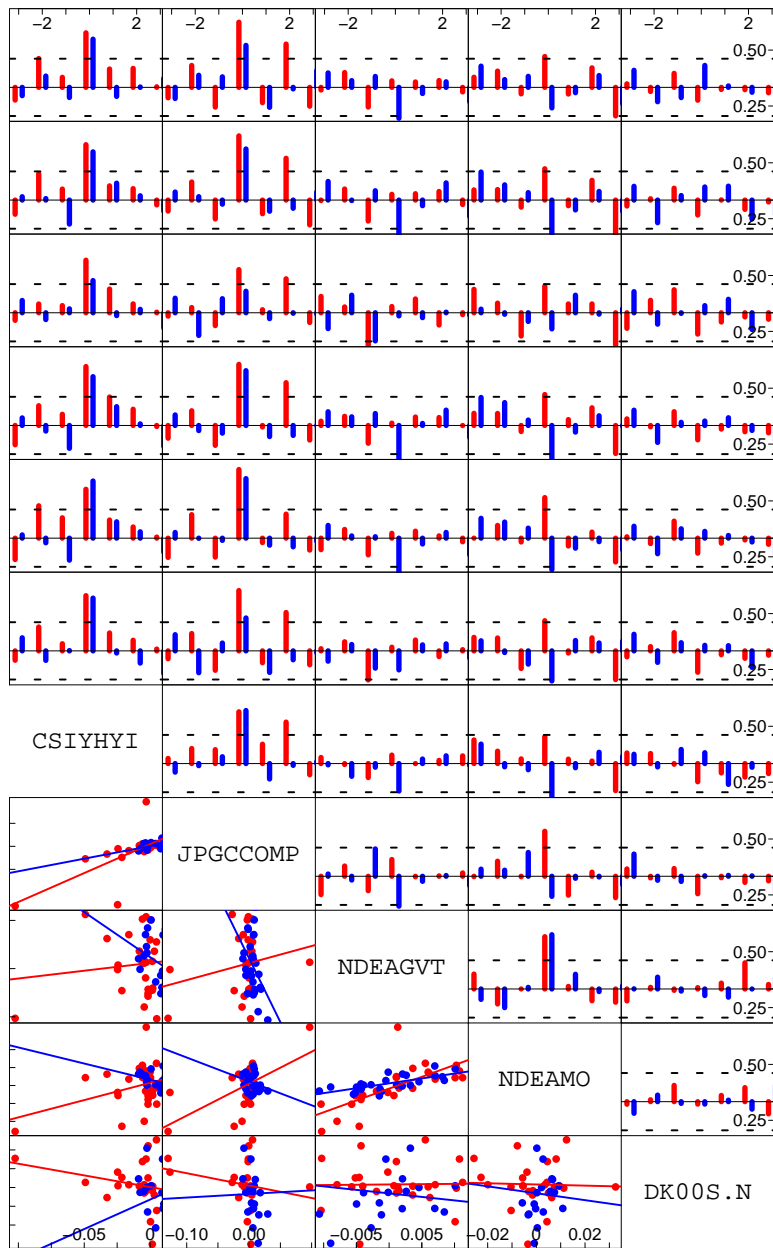


Figure 3.8: The upper panel shows the estimated cross-correlation functions for the weekly log returns for lag -3 through 3. The blue bins are estimated from the period from the beginning of may 2009 and 26 weeks forward, while the red lines are estimated from the period from the beginning of may 2008 and 26 weeks forward. The dashed horizontal lines make up a 95% confidence interval for the cross correlation. The lower panel shows scatterplots of the weekly log returns together with the best linear fits.



The estimated cross-correlation between two mutually uncorrelated time series is asymptotically normally distributed with mean value and variance

$$E[\widehat{\rho}_{XY}(k)] \simeq 0, \quad (3.15)$$

$$\text{Var}[\widehat{\rho}_{XY}(k)] \simeq \frac{1}{N}, \quad (3.16)$$

conditioned that at least one of the processes is white noise. Thus, $\pm 2/\sqrt{N}$ also makes up an approximate 95% confidence interval for the cross-correlation.[14]

3.4.3 Cross-correlation between Weekly Log Returns

From the very first plot of the indices it was evident that the series exhibit a great deal of similarity. For that reason, it is expected that there should be significant cross-correlation between many of the indices. This presumption turns out to be correct which appears from figure 3.8.

The upper panel shows the estimated cross correlation functions for the weekly log returns for lag -3 through 3. The blue bins are estimated from the period from the beginning of may 2009 and 26 weeks forward, while the red lines are estimated from the period from the beginning of may 2008 and 26 weeks forward. The dashed horizontal lines make up a 95% confidence interval for the cross correlation. The lower panel shows scatterplots of the weekly log returns from the same two periods together with the best linear fit for each period.

The two periods have been chosen carefully as exemplifications of a steep decline and incline respectively. There are three main conclusions that should be drawn from the figure. First of all, it appears from the graphs that sampling on a weekly basis almost removes all higher order cross-correlation. Second of all, it should be noticed that there is significant cross-correlation at lag 0, usually referred to as the correlation, between most of the series. Third of all, the level of cross-correlation at lag 0 is in many cases not the same for the two different estimation periods.

The correlation between some of the series is almost the same for the two periods, while for other series the correlation is less in the upturn, and in some cases it is even in-significant during the upturn, while it is significant during the downturn.

The correlation between the two Nordea indices (NDEAGVT and NDEAMO) and some of the stock indices is seen to change from significantly negative during

the downturn to in-significant or even significantly positive during the upturn. This is a remarkable difference.

From the scatter plots in the lower panel it is seen that the blue points tend to be centered in very small areas, while the red points tend to be spread out over larger areas. In for instance the scatter plot of NDDUNA against NDDUJN, the blue points form some kind of cluster, while the red points appears to be a much better fit for a straight line. In some of the scatterplots, the two lines are almost coinciding, while in other plot the two lines have completely different slopes. This agrees with the conclusions drawn from the CCF-plots.

3.5 Principal Component Analysis

Principal component analysis (PCA) is a procedure for reducing the dimensionality of a set of correlated random variables whilst keeping most of the explanatory power. It uses an orthogonal transformation to convert a set of observations of possibly correlated variables into a set of values of uncorrelated variables called principal components. The number of principal components is by definition less than or equal to the number of original variables.

The transformation is defined in such a way that the first principal component has as high a variance as possible, meaning that it accounts for as much of the variability in the dataset as possible. The succeeding components in turn have the highest variance possible under the constraint that each one has to be orthogonal to (uncorrelated with) the preceding components.[30]

Expressing the data in terms of principal components is a way of highlighting the similarities and differences, making it possible to identify the underlying structures in the dataset. It is the correlations that determine how parsimoniously the PCA will be able to summarise the data. Higher correlations mean greater redundancy, and greater redundancy results in more variation extracted in fewer components. A large number of high correlations may however indicate that several variables will weight heavily in each principal component, thus complicating interpretation.[16]

3.5.1 The Mathematical Foundation

Most of the characteristics of principal components are simply statistical reformulations of results from linear algebra. Let

$$\mathbf{X} = \begin{pmatrix} X_1 \\ \vdots \\ X_n \end{pmatrix} \quad (3.17)$$

be a multi-dimensional, standardised, stochastic variable with variance-covariance matrix

$$\text{Var}[\mathbf{X}] = \boldsymbol{\Sigma}_{\mathbf{X}}, \quad (3.18)$$

which has only ones in the diagonal. $\boldsymbol{\Sigma}_{\mathbf{X}}$ is by definition always symmetric and positive semi-definite.[14] The standardisation of the variables X_1, \dots, X_n corresponds to analysing the empirical correlation matrix instead of the empirical variance-covariance matrix.

The eigenvalues of $\boldsymbol{\Sigma}_{\mathbf{X}}$ in descending order are denoted

$$\lambda_1 \geq \dots \geq \lambda_n, \quad (3.19)$$

with corresponding orthonormal eigenvectors $\mathbf{q}_1, \dots, \mathbf{q}_n$.

The matrix

$$\mathbf{Q} = [\mathbf{q}_1 \cdots \mathbf{q}_n], \quad (3.20)$$

which has the eigenvectors as columns, is an orthogonal matrix, implying that

$$\mathbf{Q}^{-1} = \mathbf{Q}^T. \quad (3.21)$$

The vector of principal components, \mathbf{PC} , is then defined as the result of the orthogonal transformation

$$\mathbf{PC} = \begin{pmatrix} PC_1 \\ \vdots \\ PC_n \end{pmatrix} = \mathbf{Q}^T \mathbf{X}. \quad (3.22)$$

Hence, the i 'th principal component of \mathbf{X} is the projection $PC_i = \mathbf{q}_i^T \mathbf{X}$ of \mathbf{X} onto the direction of the i 'th eigenvector. It follows that the principal components are uncorrelated and the variance of the i 'th component is λ_i :

$$\text{var}[\mathbf{PC}] = \text{var}[\mathbf{Q}^T \mathbf{X}] = \mathbf{Q}^T \boldsymbol{\Sigma}_{\mathbf{X}} \mathbf{Q} = \begin{bmatrix} \lambda_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \lambda_n \end{bmatrix}. \quad (3.23)$$

It hereby follows, that the first principal component is the linear combination of the original variables that has the largest variance, since λ_1 is the largest eigenvalue. Also, the total variance of the original variables, X_1, \dots, X_n , is seen to equal the sum of the variances of the principal components, since they both equal one.[4, 5, 10]

It can be useful to think of the orthogonal transformation (3.22) as a rotation from the high-dimensional data space into a lower-dimensional principal component space. The less eigenvectors that are used in the transformation, the less the dimension of the principal component space will be. The data can of course be rotated back from the principal component space through multiplication with \mathbf{Q} from left in equation (3.22). If less than n eigenvectors are used in the transformation, then the reconstructed data will be a smoothed version of the original data due to the loss of some variation. In this way, PCA is also a means of compressing data.

3.5.2 Assumptions

PCA assumes that the observations represent an independent, random sample from a multivariate normal distribution. The multivariate normality assumption is only important, if the purpose of the analysis is to do statistical inference. Otherwise, it is sufficient if the observations represent an independent, random sample from a distribution with constant mean and variance. The sample has to be independent and random in order to obtain an unbiased estimate of the principal components. The correlations also have to be constant, since the correlation matrix would otherwise be changing in time, which would lead to non-constant principal components. The principal components are guaranteed to be independent only if these assumptions are met.

The essence of PCA is to find the new basis, which is the linear combination of the original basis, that best expresses the data set. An important underlying assumption is that linear relationships exist among the variables so that they can be combined in a linear fashion to create principal components. Pronounced non-linearities in the dataset result in large errors, when the relationships among the variables are linearised, and usually make it impossible to extract a large explanatory power in only a few principal components.

PCA is sensitive to the relative scaling of the original variables, which is why the variables should be standardised prior to the analysis. Extreme observations will exert undue pull on the direction of the component axes and therefore

strongly affect the ordination. As mentioned previously, it is important to distinguish between outliers and extreme observations, since the latter represents exceptional, yet meaningful conditions, and therefore should not be disregarded. Anyway, it is important to be aware of the strong influence of these points on the linearisation.

3.5.3 PCA on Weekly Log Returns

The weekly log returns cannot be assumed to be an independent sample from a multivariate normal distribution with constant mean and variance as both CSIY-HYI and the LIBOR-index showed signs of significant autocorrelation (appeared from figure 3.6 on page 31). Furthermore, the correlation between some of the indices changes over time (appeared from figure 3.8 on page 34), resulting in non-constant principal components. Despite this, PCA is still a means of exploring the underlying structures in the data set, without doing any statistical inference.

Table 3.4 summarises the standard deviation, the proportion of variance, and the cumulative proportion of variance explained by the estimated principal components for the eleven log return series. The analysis is based on the weekly data from 16th of June 2003 and onwards, since observations of the LIBOR-index are only available from that point on.

	PC1	PC2	PC3	PC4	PC5	PC6
Standard deviation	2.37	1.37	1.00	0.92	0.68	0.59
Proportion of Variance	0.51	0.17	0.09	0.08	0.04	0.03
Cumulative Proportion	0.51	0.68	0.77	0.85	0.89	0.92
	PC7	PC8	PC9	PC10	PC11	
Standard deviation	0.51	0.47	0.43	0.35	0.29	
Proportion of Variance	0.02	0.02	0.02	0.01	0.01	
Cumulative Proportion	0.94	0.96	0.98	0.99	1.00	

Table 3.4: Standard deviation, proportion of variance, and cumulative proportion of variance explained by the principal components estimated from the weekly log returns from 16th of June 2003 and onwards.

The standard deviation listed for each principal component is the squareroot of the corresponding eigenvalue, meaning that it is the amount of variation explained by that component. Thus, the first principal component is seen to explain a total standard deviation of 2.37 corresponding to 51% of the total

variance of the dataset. The second principal component explains 17%, and it takes eight components to explain more than 95% of the total variation.

The first seven eigenvectors are listed in table 3.5. The eigenvector coefficients represent the weights that is applied to the variables in the transformation. The coefficients are therefore sometimes referred to as the loadings.

	$\hat{\mathbf{p}}_1$	$\hat{\mathbf{p}}_2$	$\hat{\mathbf{p}}_3$	$\hat{\mathbf{p}}_4$	$\hat{\mathbf{p}}_5$	$\hat{\mathbf{p}}_6$	$\hat{\mathbf{p}}_7$
KAXGI	-0.37	-0.05	0.03	-0.20	0.04	-0.35	0.77
NDDUE15	-0.38	-0.05	0.04	-0.17	0.31	-0.21	-0.18
NDDUJN	-0.31	0.05	-0.06	0.68	-0.08	0.01	-0.08
NDDUNA	-0.37	-0.07	0.03	-0.17	0.27	-0.27	-0.45
NDUEEGF	-0.39	0.01	0.02	-0.03	0.29	0.18	-0.12
TPXDDVD	-0.35	-0.02	-0.07	0.51	-0.07	0.14	0.18
CSIYHYI	-0.32	0.10	0.02	-0.22	-0.83	-0.23	-0.18
JPGCCOMP	-0.30	0.31	0.03	-0.32	-0.09	0.72	0.00
NDEAGVT	0.11	0.64	0.07	0.14	0.02	-0.36	-0.19
NDEAMO	0.00	0.68	0.06	-0.02	0.19	-0.02	0.22
DK00S.N	0.01	-0.09	0.99	0.10	-0.03	0.04	0.01

Table 3.5: Estimates of the first seven eigenvectors based on the weekly log returns from 16th of June 2003 and onwards.

The first principal component can be interpreted as some kind of general level. Most of the indices have almost equal weight in this component except for the two Nordea bond indices and the LIBOR-index. This agrees with the conclusion from figure 3.2 on page 18, where these three indices also stood out. For the same reason the two Nordea indices achieve high weight on the second component, while the LIBOR-index has weight 0.99 on the third component. A weight of almost one, when none of the other indices have weights above 0.07, indicates that the LIBOR-index is completely uncorrelated with the rest of the indices. The two Japanese indices achieve high weight on the fourth component, while the CSIYHYI achieves a very high weight on the fifth component. JPGCCOMP and KAXGI are weighted heavily on the sixth and seventh component respectively.

Deciding on the number of components

The crucial point in PCA is the decision on how many principal components to retain in order to get a sufficient description of the variation in the data.

Typically, only those components that have a meaningful interpretation will be kept. Retaining enough components to explain at least 95%, which is the usual threshold used in statistics, often results in keeping way to many components. After all, there is no point in keeping components that explain way less than what is explained by the same component in a truly random distribution.

One criterion for selecting the number of principal components to retain is the broken stick criterion. This is based on the eigenvalues from random data. If the total variation is randomly distributed among components, then the scree plot should exhibit a broken stick distribution. A scree plot is a plot of the estimated eigenvalues as a function of the component number. Observed eigenvalues that exceed the eigenvalue expected under the broken stick distribution are considered meaningful and should be retained for interpretation. The eigenvalues under the broken stick distribution³ are

$$\lambda_i^* = \sum_{k=i}^L \frac{1}{k}. \quad (3.24)$$




Figure 3.9: Scree plot of the eigenvalues from the principal components analysis of the weekly log return data (black) and from random data (red) against the rank order

The eigenvalues of the broken stick distribution are plotted together with the estimated eigenvalues in figure 3.9. According to the broken stick criterion, only the first component should be kept.[16] This would lead to only 51% of the total variance being explained. For most applications, this is simply to little.

Another approach is to decide based on the screeplot alone how many components to retain. The smaller eigenvalues, representing random variation, tend to lie along a straight line. Thus, one could look for a turning point in the scree plot, where the curve flattens out. In the present case, this is after the fourth component. The third and the fourth components have almost identical eigenvalues, but these two points are not on line with the following points. Retaining the first four components, which altogether account for acceptable 85% of the total variance, seems more reasonable. See [8] for further criteria.

³If a stick of unit length is broken into L pieces, then the expected length of the i 'th longest segment is $\frac{1}{L} \sum_{k=i}^L \frac{1}{k}$.

Communalities

Having decided on the number of components to retain, it should be measured, how well each variable is explained by the retained components. This measure is called the communality. The communality of the j 'th variable is

$$c_j = \sum_{i=1}^L q_{i,j}^2 \lambda_i, \quad (3.25)$$

where $q_{i,j}$ is j 'th element of eigenvector \mathbf{q}_i , and the $L \in [1; n]$ is the number of retained principal components. The communality is the proportion of a variable's variance that is accounted for by the principal components. The communalities will of course be one, if all the principal components were retained.[16]

The communalities for the eleven series based on retention of four principal components appear from table 3.6. CSYHYI has the lowest communality with 64% of the variance being explained, while the LIBOR-index is explained perfectly. This is due to the fact that the third component almost exclusively explained the LIBOR-index. If the third component was excluded, then the communality of the LIBOR-index would be close to zero, while the rest of the communalities would be unaffected. There is no point in including the LIBOR-index in the PCA, when it is completely uncorrelated with the rest of the indices.

KAXGI	NDDUE15	NDDUJN	NDDUNA	NDUEEGF	TPXDDVD
0.79	0.86	0.95	0.81	0.84	0.92
CSYHYI	JPGCCOMP	NDEAGVT	NDEAMO	DK00S.N	
0.64	0.78	0.87	0.87	1.00	

Table 3.6: Communalities for the eleven weekly log return series based on the first four principal components.

Finally, it should be emphasised that the conclusions drawn from the PCA should be accepted with caution, since the weekly log returns don't meet the assumptions of PCA, as already discussed. There is however no doubt left, that the LIBOR-index is very different from the rest, while the rest of the indices have a lot in common.

CHAPTER 4

Modelling

This chapter is concerned with finding and calibrating a time series model to the index series. In the choice among approaches, it is important to emphasise that the aim is to use the model for generating scenarios. Therefore, a model that is more readily applied to generating scenarios is preferred to a more exhaustive model with a larger explanatory power, if it is able to describe the variation in the data sufficiently.

It is only the stock and bond market indices that are modelled, since it appeared from chapter 3 that the LIBOR-index is uncorrelated with the rest of the indices. Thus, the LIBOR-index should have its own distinct model, which would not contribute to the present analysis. Furthermore, the lack of observations of the LIBOR-index from the first four and a half years means that this period would have to be disregarded in the model estimation, if this index was included in a multivariate model.

As the last thing in the chapter, the performance of the chosen model is tested on the data in an in-sample test. It will be the weekly log returns that are modelled, since these are mean stationary and almost free of autocorrelation.

4.1 Choosing a Regime Approach

From chapter 3 it appeared that the main challenge is that everything in the dataset seems to be changing over time. Since there is no doubt about the strong correlation between many of the series, the chosen model should benefit from the correlations in order to improve the model performance. Given that the correlations between the indices aren't constant, the only option is some kind of threshold model.

A threshold model is useful when modelling a system, where the dynamical behaviour depends on the actual state.[15] From figure 3.8 on page 34 it appeared that there is as a minimum a difference between the behaviour during upturns and downturns.

Since threshold models can be very complicated to estimate, the approach taken here is to partition the overall period into shorter periods, where correlations can be assumed to be constant. In each discrete regime the dynamics of the system is then approximated by some 'simple' model.

4.1.1 The Number of Regimes

The first thing to decide on is the number of different regimes. The partitioning into regimes should be based on the primary trend, which is the long-term movement of prices, lasting from several months to several years. Otherwise the number of regimes would be too large for the approach to be feasible. Within each regime there will then be intermediate trends caused by short-term deviations from the underlying trend line. These deviations are less important, since they are eliminated, when prices revert back to the trend line.[2] The daily fluctuations can of course deviate from the intermediate trend, but this has no importance, when looking at a five year time horizon. Looking at weekly data also has the advantage that it smooths these minor trends.

Based on the analysis of the cross-correlation function in section 3.4, it is expected that as a minimum there will be a significant difference between upturn and downturn. The typical choice, when dealing with financial data, is four different types of regimes; expansion, peak, recession, and trough.[2]

Peak is the period, where the economy is running at its highest. After a peak

follows a recession, where the economy experiences negative growth. The recession leads to a trough, where the economy is at its lowest. After a trough comes a period of expansion, where the economy experiences growth. The expansion then leads to a peak and so on. This recurring pattern of contraction and recovery is called the business cycle.

The simplest choice is obviously to partition into upturn and downturn. The rather abrupt changes observed between expansion and recession advocate this approach, since there is no need for intermediate regimes. The more different regimes that are applied, the more turningpoints are needed, and the harder it will be to identify the turningpoints. The approach that will be taken here is to partition the period into upturn and downturn, and then afterwards assess, whether this partitioning is sufficient.

Having decided upon the number of different regimes, the decision on the actual number of regimes still remains. The question is, whether each upturn and downturn, respectively, can be assumed to be identical. It seems reasonable to expect that the dynamics underlying the two types of regimes would be the same for each respective phase. On the other hand, it appeared from the data analysis that almost everything in the dataset is changing over time.

Once again the approach taken will be the simplest one, that is joining the upturns and downturns into two different regimes. This reduces the number of models that have to be estimated and avoids the choice of which upturn and downturn to simulate from, when generating scenarios.

4.1.2 OECD Composite Leading Indicators

As already mentioned, the economy recurrently experiences periods of upturn and downturn. The length and depth of those phases can be irregular. In a two regime model, the transition points across cycles are called peaks and troughs; a peak is the transition from the end of an expansion to the start of a contraction. A trough occurs at the bottom of a recession just as the economy enters a recovery. Here peaks and troughs are points in time and not periods of time like before, when they referred to regimes.[2] More precisely, a cycle is the time span separating two turning points of the same nature (two peaks or two troughs). The time span between a peak and a trough is referred to as the phase length.

Given the cyclical nature of the economy, it is not surprising that to some extend

the cycle can be predicted. Leading economic indicators are economic time series that tend to rise or fall in advance of the rest of the economy.[2] A leading indicator is said to have a leading relationship with respect to the turning points of a reference series. A composite leading indicator, CLI, is constructed by combining two or more leading indicators into a single index number.




Figure 4.1: The normalised composite leading indicator together with a normalised index of industrial production for the total OECD area. The vertical lines indicate the turning points of the CLI: Trough 1998M10, Peak 2000M2, Trough 2001M9, Peak 2007M5, Trough 2009M2, Peak 2011M2.

term trend with a median of six months and a standard deviation of three and a half month for the total OECD area.[21]

The construction of the OECD CLI is a complicated matter. First reference series and component series are chosen. Reference series are economic time series that exhibit similar cyclical fluctuations to those of the business cycle, but which precede those of the business cycle. The typical choice is an index of industrial production. The component series have to show a leading relationship with the reference series and the relationship should have an economic explanation. Also, series with a broad coverage of economic activity are preferred to narrowly-defined series.

The potential leading indicators are divided into four groups: *Early stage indicators* such as new orders and construction approvals. *Rapidly responsive indicators* such as average hours worked, profits, and stocks. *Expectation-sensitive indicators* measuring, or sensitive to, expectations, such as stock prices, raw material prices, and data from business surveys concerning production, the general economic situation, consumer confidence etc. And finally *prime mover*

¹Organisation for Economic Co-operation and Development.

The OECD¹ system of composite leading indicators was developed in the 1970's to predict turningpoints in a reference series chosen as a proxy for economic activity. Fluctuations in economic activity are measured as the variation in the output gap, typically represented by the movements in the GDP around its long-term trend. The normalised CLI in figure 4.1 shows a leading relationship with the normalised index of industrial production (the reference series) for the total OECD area. Since 1961 the turningspoints of the OECD CLI has consistently been preceding those of the economic activity relative to long-

indicators relating to monetary policy and foreign economic developments such as money supply and terms of trade.

Some potential indicators might belong to more than one of the groups. The idea behind dividing potential leading indicators into the four above mentioned groups is to balance the composition of a given CLI and the cyclical events that can impact it by including indicators from all four categories.

When the leading components have been selected they are filtered for outliers, seasonal patterns, long term trends etc. in order to clarify the cyclical patterns. This also includes a smoothing that removes high frequency noise. Then the turningpoints are detected using the Bry-Boschan algorithm. This routine specifies a minimum duration of five months for a phase and fifteen months for a cycle. The last step in the filtering process is a normalisation to secure a common scale.

The cyclical performance of each candidate component series is then evaluated using different statistical methods. A component is only considered for inclusion if the detected turning points have suitable leading properties. The best performing component series are selected based on the length and the consistency of their lead and the cyclical conformity between the indicators and the reference series.

The selected leading indicators are then composed to one composite leading indicator. Each component receives the same weight. It should be noted that the normalisation of the components described above, introduces an implicit weighting of the component series, with the series being weighted by the inverse of their standard deviation. OECD calculates CLIs for different countries and zones. In the zone aggregates the CLIs themselves are weighted reflecting the country weights.[20]

4.1.3 Partitioning According to OECD CLI

Instead of partitioning the data period into upturn and downturn based on a subjective assessment, the partitioning will be done according to the CLI for the total OECD area.² No matter how the partitioning is done, it will give rise

²The OECD Total Area covers the following 29 countries: Australia, Austria, Belgium, Canada, Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Japan, Korea, Luxembourg, Mexico, Netherlands, New Zealand, Norway, Poland, Por-

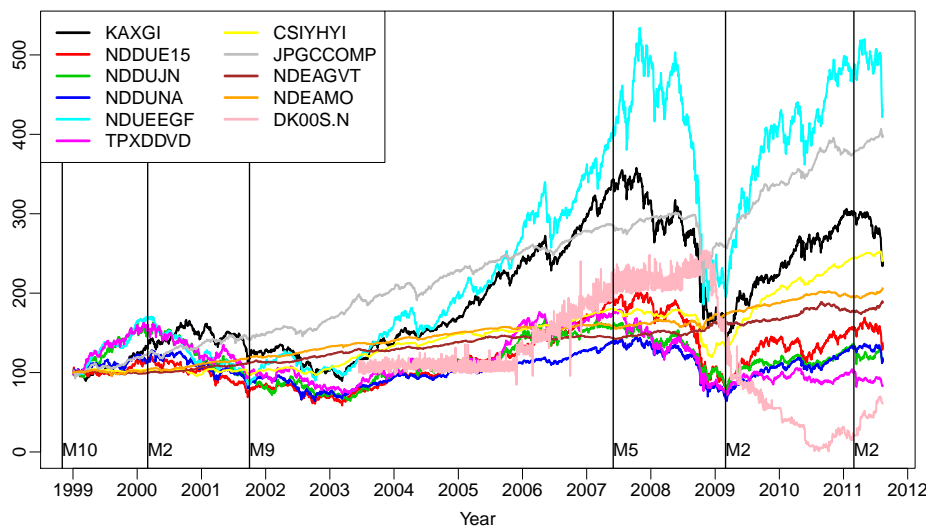


Figure 4.2: The eleven indices scaled to start at index 100 on the 1st of January 1999. The vertical lines indicate the partition into upturns and downturns based on the turning points of the CLI for the total OECD area.

to an error as a consequence of the indices representing different parts of the world, and the turning points being different from one country to another. For instance, the turning point marking the end of the latest financial crisis was at the end of October 2009 in Denmark, while it was at the end of June 2009 in the United States. The total area CLI is chosen since the indices included in this project represent a large part of the world.

The CLI can be used to partition the indices into regimes even though the CLI tends to precede turning points in economic activity by six months. The reason being that the financial markets react on expectations based on such indicators and therefore also precede the turning points in the economic activity. This is supported by the fact that stock indices are grouped both as rapidly responsive indicators and expectation-sensitive indicators in the construction of the CLI.

Figure 4.2 shows the eleven indices, scaled to start at index 100 on the 1st of January 1999, partitioned into regimes according to the turning points of the CLI for the total OECD area (see figure 4.1).

The OECD CLIs are revised regularly, partly because countries revise many
tugal, Slovak Republic, Spain, Sweden, Switzerland, Turkey, United Kingdom, and United States.

informations after they have been published, partly because of changes in the computation method, and partly because the filtering of the series means that new information can change previous numbers. The partitioning was done according to the turning points of the normalised CLI for the total OECD area based on data extracted from the OECD database on the 5th of January 2012.

The partitioning seems overall decent. Especially the turning point at the end of May 2007 would be moved slightly if it was decided based on a subjective assessment instead. It is however evident that the indices don't share the same turning points. It appears that there are intermediate trends like a steep decline for a couple of months during an upturn and vice versa, but the restrictions on phase and cycle length are perfectly fine, since otherwise the number of regimes would be too large.

Once again it appears that the choice of disregarding the LIBOR-index is correct, since this index seen to have moved opposite the rest of the indices for the last couple of years.

4.2 Choosing a Model

Having decided on the partitioning, the next choice is the model class to be used within each regime. A linear time series model like ARIMA is not applicable, since it is unable to account for the non-constant variance observed (see figure 3.4 on page 21).

Only one of the ten indices in question (CSIYHYI) exhibited significant autocorrelation, when looking at weekly log returns. Assuming that there is no autocorrelation, a possibility could be a multivariate random walk.

The downside of using a multivariate model is that the scenario generation is complicated as all the indices have to be simulated simultaneously. Emphasising that the aim is to generate scenarios, there is a preference towards simulating principal components instead, since these are independent and consequently can be simulated independently, still assuming that there is no autocorrelation. Modelling principal components also has the advantage that it offers an easy way of reducing the dimensionality of the simulation space.

The principal components are conditionally heteroscedastic just like the log returns, since the principal components are simply a rotation of the log returns.

Since the log returns displayed signs of volatility clustering (figure 3.3 on page 20 and 3.4 on page 21), it is expected that a conditional heteroscedastic model is necessary for modelling the principal components. Conditional heteroscedasticity means that the conditional variance depends on past observations.

The conditional heteroscedasticity doesn't conflict with the assumption of PCA, that the log returns originate from a multivariate normal distribution with constant mean and variance, since the unconditional variance is assumed to be constant. The crucial point is whether the correlations can be assumed to be constant within the regimes.

4.2.1 ARCH and GARCH Models

The Generalised AutoRegressive Conditional Heteroscedastic model (GARCH-model) of order (p, q) is given by

$$y_t = \mu_t + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_t^2) \quad (4.1)$$

$$\sigma_t^2 = \omega + \sum_{i=1}^q \alpha_i \varepsilon_{t-i}^2 + \sum_{j=1}^p \beta_j \sigma_{t-j}^2, \quad (4.2)$$

where $\omega > 0$, $\alpha \geq 0$, $\beta \geq 0$, and $\mu_t = \mathbb{E}[y_t | Y_{t-1}]$ is the conditional mean of y_t . It appeared from figure 3.3 on page 20 that the log returns are mean stationary, thus $\mu_t = \mu$. The model is weakly stationary if $\alpha + \beta < 1$. In the special case $p = 0$, the model is referred to as an ARCH(q)-model.

The GARCH-model assumes that the conditional variance σ_t^2 is a linear function of previous squared errors and previous variances. Large fluctuations in y_t induces a larger variance, and thus a larger probability of large fluctuations in the future.

The unconditional variance of the process is

$$\text{Var}[y_t] = \frac{\omega}{1 - \alpha - \beta}. \quad (4.3)$$

The model is based on the assumption that the conditional distribution of y_t is normal. This is seen to be consistent with the no-arbitrage principle, as the best prediction of the next return is simply its unconditional mean:

$$y_{t+1|t} = \mathbb{E}[y_{t+1} | Y_t] = \mathbb{E}[y_{t+1}] = \mu_t. \quad (4.4)$$

It is possible to make other assumptions about ε_t than the normal distribution, but this is rarely done in practice. The reason being that the maximum likelihood estimator based on the normal distribution has the huge advantage that it is robust to misspecification of the conditional distribution, if the conditional mean, μ_t , and the conditional variance, σ_t^2 , are correctly specified.[13]

4.2.2 The Likelihood of the Model

The likelihood function of the model is the probability density function of the observed outcome $Y_N = (y_N, \dots, y_1)$ given the parameters θ :

$$L(\theta) = f(Y_N|\theta) = f(y_1) \prod_{i=2}^N f(y_i|Y_{i-1}, \theta). \quad (4.5)$$

The conditional distribution is assumed to be

$$f(y_t|Y_{t-1}) = \frac{1}{\sqrt{2\pi\sigma_t^2}} \exp\left(-\frac{1}{2} \frac{\varepsilon_t^2}{\sigma_t^2}\right), \quad (4.6)$$

where $\varepsilon_t = y_t - \mu_t$.

Since the likelihood function is a product of N terms involving the exponential function of a negative value, it is often convenient to consider the logarithm of the likelihood function instead of the likelihood function itself, since this function is easier to evaluate on a computer:

$$\ln L(\theta) = \sum_{t=1}^N \ln L_t(\theta), \quad (4.7)$$

where

$$\ln L_t(\theta) = -\frac{1}{2} \ln(2\pi) - \frac{1}{2} \ln \sigma_t^2 - \frac{1}{2} \frac{\varepsilon_t^2}{\sigma_t^2} \quad (4.8)$$

is the conditional distribution of y_t given Y_{t-1} .

For a GARCH(p, q)-model, there is no knowledge of the distribution of the first $p + q$ observations. This problem is tackled by looking at the conditional distribution of y_N, \dots, y_{p+q+1} given Y_{p+q} corresponding to $t = p + q + 1, \dots, N$ in (4.7). This has no influence on an optimisation of the likelihood function, since these terms would be constant.

4.3 Estimating Principal Components

Table 4.1 summarises the standard deviation, the proportion of variance, and the cumulative proportion of variance explained by the estimated principal components for the ten centered and scaled weekly log return series during the three upturns from figure 4.2 on page 50. The centering and scaling factors for both upturn and downturn appear from table A.1 on page 93 in appendix A.

	PC1	PC2	PC3	PC4	PC5
Standard deviation	2.11	1.39	1.07	0.82	0.75
Proportion of Variance	0.44	0.19	0.12	0.07	0.06
Cumulative Proportion	0.44	0.64	0.75	0.82	0.87
	PC6	PC7	PC8	PC9	PC10
Standard deviation	0.64	0.56	0.49	0.44	0.31
Proportion of Variance	0.04	0.03	0.02	0.02	0.01
Cumulative Proportion	0.91	0.95	0.97	0.99	1.00

Table 4.1: Standard deviation, proportion of variance, and cumulative proportion of variance explained by the principal components estimated from the joined expansive periods according to figure 4.2 on page 50.

The first seven associated eigenvectors are listed in table 4.1.

	$\widehat{\mathbf{p}}_1$	$\widehat{\mathbf{p}}_2$	$\widehat{\mathbf{p}}_3$	$\widehat{\mathbf{p}}_4$	$\widehat{\mathbf{p}}_5$	$\widehat{\mathbf{p}}_6$	$\widehat{\mathbf{p}}_7$
KAXGI	-0.35	0.05	0.18	-0.15	0.63	-0.59	-0.09
NDDUE15	-0.41	0.04	0.19	-0.37	-0.08	0.09	-0.07
NDDUJN	-0.32	-0.10	-0.63	0.10	-0.07	-0.04	0.03
NDDUNA	-0.38	0.03	0.21	-0.47	-0.21	0.28	0.48
NDUEEGF	-0.40	0.01	0.06	0.06	-0.28	0.16	-0.79
TPXDDVD	-0.35	-0.00	-0.58	0.04	0.05	-0.04	0.20
CSIYHYI	-0.31	-0.20	0.24	0.54	0.45	0.53	0.14
JPGCCOMP	-0.24	-0.38	0.30	0.41	-0.50	-0.48	0.23
NDEAGVT	0.16	-0.62	-0.09	-0.17	0.07	0.12	-0.05
NDEAMO	0.06	-0.65	-0.03	-0.32	0.12	-0.01	-0.14

Table 4.2: Estimates of the first seven eigenvectors based on the joined expansive periods according to figure 4.2 on page 50.

The squareroots of the eigenvalues and the eigenvectors for the joined recessive periods appears from table 4.3 and 4.4.

It takes four principal components to explain 82% of the variation during the

	PC1	PC2	PC3	PC4	PC5
Standard deviation	2.33	1.36	0.95	0.72	0.65
Proportion of Variance	0.55	0.19	0.09	0.05	0.04
Cumulative Proportion	0.55	0.73	0.82	0.87	0.92
	PC6	PC7	PC8	PC9	PC10
Standard deviation	0.50	0.47	0.42	0.35	0.28
Proportion of Variance	0.02	0.02	0.02	0.01	0.01
Cumulative Proportion	0.94	0.96	0.98	0.99	1.00

Table 4.3: Standard deviation, proportion of variance, and cumulative proportion of variance explained by the principal components estimated from the joined recessive periods according to figure 4.2 on page 50.

joined upturn, while three principal components are enough to explain the same proportion of variation during the joined downturn. This is not surprising, since the cross-correlations in many cases appeared stronger during the downturn than the upturn (see figure 3.8 on page 34). This difference in the eigenvalues, and consequently the explanatory power, of the first principal components is further evidence of the difference between the two types of regimes.

	$\hat{\mathbf{p}}_1$	$\hat{\mathbf{p}}_2$	$\hat{\mathbf{p}}_3$	$\hat{\mathbf{p}}_4$	$\hat{\mathbf{p}}_5$	$\hat{\mathbf{p}}_6$	$\hat{\mathbf{p}}_7$
KAXGI	-0.38	0.05	-0.17	0.08	-0.22	0.37	-0.27
NDDUE15	-0.38	0.02	-0.19	0.37	-0.11	0.11	0.09
NDDUJN	-0.31	0.04	0.70	-0.11	-0.08	-0.01	0.01
NDDUNA	-0.36	0.09	-0.20	0.34	-0.43	-0.55	0.08
NDUEEGF	-0.38	-0.00	-0.07	0.10	0.32	0.59	0.13
TPXDDVD	-0.35	0.07	0.52	-0.03	0.12	-0.18	-0.02
CSIYHYI	-0.31	-0.02	-0.26	-0.84	-0.28	-0.02	-0.02
JPGCCOMP	-0.32	-0.27	-0.24	-0.09	0.66	-0.38	0.24
NDEAGVT	0.06	-0.67	0.12	0.03	-0.34	0.16	0.60
NDEAMO	-0.06	-0.68	0.03	0.09	0.01	-0.07	-0.69

Table 4.4: Estimates of the first seven eigenvectors based on the joined recessive periods according to figure 4.2 on page 50.

Comparing the eigenvectors of the two regimes it is seen that there are certain differences, but the overall interpretation of the first four components is the same: the first principal component is the general level. Most of the indices have almost equal weight in this component except for the two Nordea bond indices. The two Nordea indices achieve high weights on the second component. The two Japanese indices achieve high weights on the third component, while CSIYHYI has by far the highest weight on the fourth component.

4.3.1 Deciding on the Number of Components

As mentioned previously, the crucial point in PCA is the decision on how many principal components to retain. It is of course possible to retain all ten components and simply use the principal component analysis to obtain independency between the components that is simulated. But why not use the analysis to reduce the number of quantities that should be simulated, when this can be done without losing anything but some apparent random disturbance.




Figure 4.3: Scree plot of the eigenvalues against the rank order for random data and for the principal components of the weekly log returns partitioned into regimes.

According to the broken stick criterion, only the first two principal components should be kept. This appears from figure 4.3, where the eigenvalues of the broken stick distribution are plotted together with the estimated eigenvalues for the two regimes. Keeping two principal components would lead to an average explanatory power of 64% during the upturns, which seems insufficient for the purpose of generating scenarios.

Based on the screeplot alone it appears that three components should be retained, since the curves flatten out after the third component. Retaining the first three components would result in average explanatory powers of 75% and 82% respectively, which is more likely to be sufficient for the purpose of generating scenarios. The difference in explanatory power between the two regimes reflects that correlations are weaker during upturns resulting in more unexplainable random behaviour.

4.3.2 Communalities

Upon deciding on the number of components to retain, the communalities should be investigated to see how well each variable is explained by the retained principal components.

The communalities for the ten indices based on retention of three principal components during the upturns appear from table 4.5. The corresponding estimates

for the downturns appear from table 4.6. It should be noted that the communalities are averages, since the conditional variance of the principal components is expected to be non-constant.

KAXGI	NDDUE15	NDDUJN	NDDUNA	NDUEEGF
0.60	0.78	0.93	0.68	0.72
TPXDDVD	CSIYHYI	JPGCCOMP	NDEAGVT	NDEAMO
0.92	0.56	0.63	0.85	0.82

Table 4.5: Communalities for the ten weekly log return series during the joined upturn based on the three four principal components.

The index with the lowest communality is the CSIYHYI-index with only 56% of the variance being explained during upturns and 60% during downturns. The obvious explanation is that the CSIYHYI-index has a high weight on the fourth component, which is the first component not retained. Also the KAXGI and the JPGCCOMP-index have very low communalities during upturns. The low communalities during upturns are acceptable since they have an explanation in the amount of random behaviour.

KAXGI	NDDUE15	NDDUJN	NDDUNA	NDUEEGF
0.83	0.84	0.96	0.76	0.81
TPXDDVD	CSIYHYI	JPGCCOMP	NDEAGVT	NDEAMO
0.94	0.60	0.75	0.86	0.88

Table 4.6: Communalities for the ten weekly log return series during the joined downturn based on the first three principal components.

Only the CSIYHYI-index is not sufficiently explained neither during upturns nor downturns by the retained components. This could be the result of disregarding the significant autocorrelation observed in this index series. With most of the communalities for the rest of the indices being much above 70%, the number of components retained seems sufficient.

4.3.3 Assumptions

PCA assumes that there is no autocorrelation within the log return series since the observations represent an independent, random sample from a multivariate normal distribution. The normality assumptions has already been discussed thoroughly in section 3.3 and seems acceptable, but the CSIYHYI has significant

autocorrelation. Accepting the independency assumption means accepting a one in ten error, which has to be emphasised. The risk of the principal components not being perfectly independent is accepted, since there is no obvious explanation for the observed weekly autocorrelation.

The linearity assumption appears unproblematic, as it would be hard to extract an explanatory power as high as 80% in only a few principal components if there were pronounced non-linearities.

The ten log return series were standardised prior to the PCA in order to handle the sensitivity of the analysis to the relative scaling of the original variables. Finally, the influence of the many extreme observations on the direction of the component axes has to be taken into account, but as discussed in chapter 3, these extreme observations should be included since realistic scenarios have to contain a possibility of such extreme events.

4.4 Estimating Model Parameters

The parameters in the GARCH-model (4.2) are estimated through a maximisation of the conditional log likelihood function using the function “garch” from the “tseries” package in R. The code used can be found in appendix B.2 on page 109.

It has to be mentioned that there is a flaw in the estimation of the parameters that arises from the merging of the different upturns and downturns respectively. The first $p + q$ observations after each merging point should be disregarded in the evaluation of the conditional likelihood function. This is not done due to the technical challenges that is involved in writing a new parameter estimation function. The effect of not disregarding a few points should be negligible given the large number of observations within each regime.

4.4.1 Parameter Estimates

The parameter estimates of the optimal ARCH-models appear from table 4.7. The principal components are found to be best described by an ARCH(1)-model. This means that the conditional variances are independent of previous variances.

	$\hat{\omega}$	SD ($\hat{\omega}$)	$\hat{\alpha}_1$	SD ($\hat{\alpha}_1$)	p_{JB}	p_{LB}
PC_1^{up}	3.519	0.23	0.221	0.06	0.00	0.47
PC_2^{up}	1.493	0.12	0.225	0.05	0.00	0.95
PC_3^{up}	1.041	0.08	*0.092	0.07	0.08	0.74
PC_1^{down}	1.978	0.30	0.732	0.13	0.00	0.69
PC_2^{down}	1.365	0.14	0.261	0.10	0.01	0.69
PC_3^{down}	0.894	0.09	*0.015	0.09	0.00	0.98

Table 4.7: Estimates of the parameters in (4.2) along with the standard deviation of the estimates and p -values for a Jarque-Bera test for normality of the residuals and a Ljung-Box test for independency among the squared residuals. * denotes estimates that are statistically insignificant at 95%-level.

$\mu_t = 0$ in all cases since the principal components are computed for the standardised weekly log returns (the standardisation factors appear from table A.1 on page 93 in appendix A).

Most of the stated parameters are statistically significant on a 95%-confidence level. This does however not apply to the α_1 parameter for PC_3 . The third principal components is simply a random walk with variance ω .

p_{JB} is the p -value of a Jarque-Bera test of whether the residuals can be assumed to be normally distributed, as they should be. The test is seen to reject the possibility of the residuals being normally distributed in all but one case, namely PC_3^{up} . This is not surprising, since it is the same conclusion that was drawn in section 3.3, when the same test was applied to the log return series, despite the resemblance noticed by the robust moments.

p_{LB} is the p -value of a Ljung-Box test for independency among the squared residuals. It is seen to be very likely that the squared residuals are independent, as they should be. This suggests that the models are valid. If the neglected autocorrelation in the CSIYHYI-index compromised the independency assumption, this test should have shown.

The 95%-confidence intervals for the parameter estimates for PC_2 and PC_3 during upturn and downturn are seen to overlap, meaning that the behaviour of these principal components is not significantly different from upturn to downturn.³ This is not the case for the parameter estimates for PC_1 . The parameters

³If the model assumptions are fulfilled, the parameter estimates are approximately t -distributed, meaning that an approximate 95%-confidence interval for a parameter estimate is $\pm 2 \cdot \text{SD}$, when the number of degrees of freedom is large.

for this component, which is interpreted as the general level of log returns, are significantly different from upturn to downturn. This is the final evidence of the difference between the two regimes.

The merging of the individual upturns and downturns into two regimes leads to a smoothing. The expansive regime will not be as expansive as the most expansive of the three upturns, and conversely the recessive regime will not be as recessive as the worst decline observed during the two downturns. In this way, the differences between upturn and downturn are smoothed, which could be part of the explanation for the similarities observed between some of the principal components.

4.4.2 Considering Stationarity

It has been shown that if a GARCH-model is estimated on data that contain switches in the data-generating parameters of the conditional variance equation (4.2) that are not accounted for, then the sum of the estimated parameters converges to one.[7] Therefore it is worth noting that none of the estimated models are close to being non-stationary, since $\hat{\alpha}_1$ is clearly less than one. This suggests that the partitioning based on the OECD leading indicators is reasonable even though it is far from perfect.

Also, if the flaw in the estimation of the parameters that arose from not disregarding the first observation after each merging point was important, it should have shown through parameters converging to one.

In the process of modelling the data several partitionings have been tried. Among those, partitioning according to the turning points of the OECD reference series. The parameter estimates were found to be very sensitive to the partitioning. Minor changes in the turning points led to completely different model parameters. The impact on the generated scenarios would probably be less apparent, since the model class was always the same. The important point to make is that the estimated models are no way near being non-stationary, which suggest that the chosen partitioning is adequate. Then the sensitivity to the partitioning is less important.

4.5 Assessing Model Performance

Before the estimated models are applied to generate scenarios, they are tested in a so-called in-the-sample test. This is done through simulations of the data-period based on the estimated model. The simulated time series should resemble the true data for the model to be applicable for generating scenarios forward in time. There are not enough observations available to also do an out-of-sample test, when the purpose is to generate scenarios five years forward in time.

The 5%, 25%, 50%, 75%, and 95%-quantiles of the 10,000 simulated time series are shown in figure 4.4 together with the observed time series (red). The lowest dashed line is the 5%-quantile, the second lowest is the 25%-quantile, the median is the solid black line etc. The phase length of each up- and downturn was kept fixed to the observed value in the simulations.

The CSYHYI-index is clearly the index that is worstly described by the model. From 2004 until 2010 the median of the simulations deviates considerably from the observed time series. After 2010 the median and the observed time series are reunited.

The model seems unable to capture the magnitude of all the observed fluctuations, but the observed values are seen to reunite with the median shortly after each deviation. With only a few exceptions, the fluctuations are within the interquartile range of the simulated data. Only three indices (KAXGI, TPXD-DVD, and CSYHYI) are seen to have fluctuations that lie outside the 5% or the 95%-quantile at some point.

The simplicity of the model comes with the cost of a smoothing as already discussed. Even with a more complicated model, it is unlikely that the median would reflect the observed magnitude of fluctuations. The important point is that there should be some scenarios that reflect the possibility of such steep inclines and declines as observed in the data. All in all, the results are found to be decent, so the model seems applicable, since the observed values lie within the interquartile range most of the time.

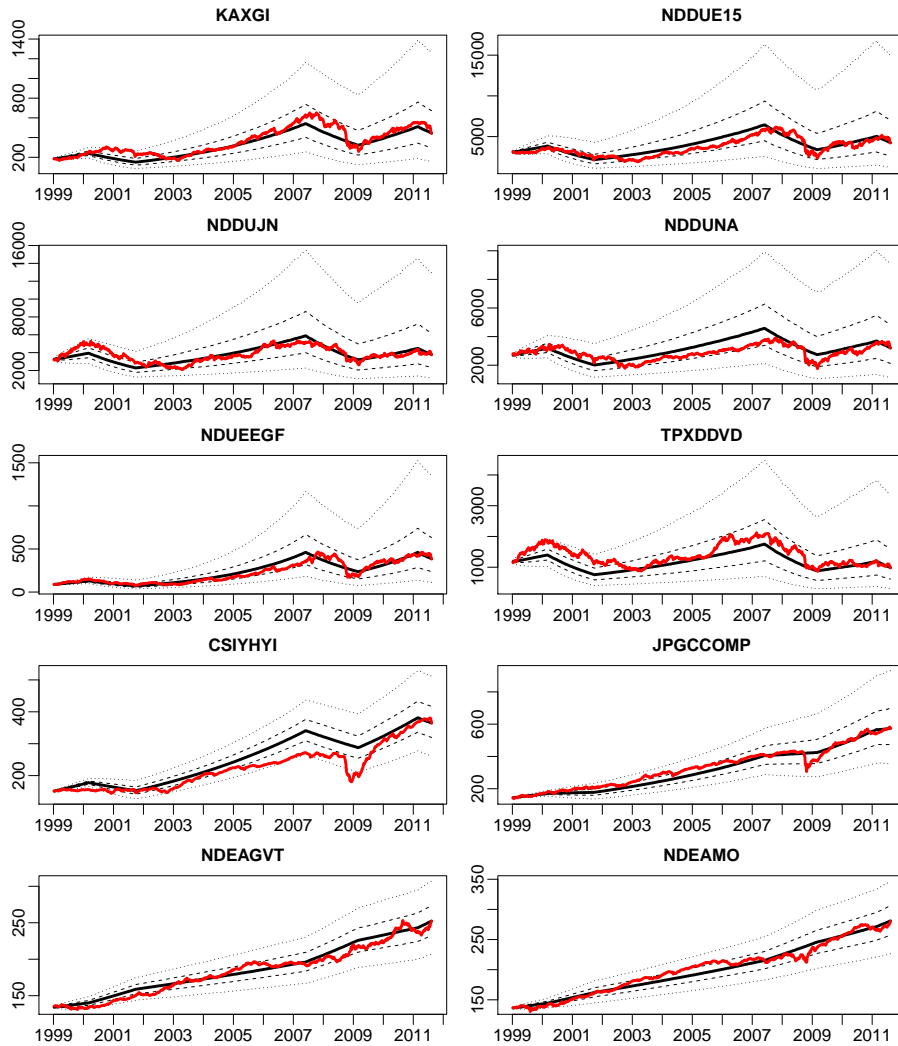


Figure 4.4: The 5%, 25%, 50%, 75%, and 95%-quantiles of the 10,000 simulated time series together with the true datapoints (red). The simulations are done keeping the phase length of each up- and downturn fixed to the observed value.

CHAPTER 5

Scenario Generation

Having chosen a model with emphasis on the purpose of generating scenarios makes the scenario generation a lot easier. In the chosen regime model, which fitted decently with the data, each principal component can be simulated independently. The challenge that remains is the choice of which regime to simulate from.

Each scenario is a weekly realisation of the value of all ten indices five years forward in time from the 12th of August 2011. The generated scenarios are visualised using quantiles, since these are robust to the presence of a few very extreme scenarios. The interquartile range is very important as a measure of volatility since the collection of scenarios should capture the range of variations in the indices that is likely to occur.

There are also generated scenarios using a bootstrapping procedure developed based on the observations from the data analysis chapter. The bootstrapped scenarios provide a standard of comparison for the scenarios generated with the model from chapter 4.

5.1 The Choise of Regime

For each time step in the scenario generation there has to be a choice of which regime to simulate from. This choice should be made in a way that is consistent with the no-arbitrage principle. Given that the economy was in a recession on the 12th of August 2011, where the scenarios begin, almost all the scenarios should begin with values simulated from the downturn model.

The length of the current phase is a positive random variable T chosen to have exponential distribution with rate λ , where λ is a positive parameter. T then has probability density

$$f(t) = \lambda \exp(-\lambda t), \quad t \geq 0. \quad (5.1)$$

The exponential distribution is chosen, because T then by definition has the memoryless property:

$$P(T > t + s | T > t) = P(T > s), \quad (s \geq 0, t \geq 0). \quad (5.2)$$

Given that a crisis has lasted t , the chance that it will last a further time s is the same as the chance that it would last s in the first place.[24] In other words, the probability of shifting regime is the same each week. This is perfectly consistent with the no-arbitrage principle.

The parameter λ is chosen as the reciprocal of the empirical average length of the phase in question:

$$E[T] = SD[T] = \frac{1}{\lambda} = \hat{\mu}_T. \quad (5.3)$$

The theoretical median of T is

$$\text{Median}[T] = \frac{1}{\lambda} \ln 2. \quad (5.4)$$

The exponential distribution is a continous distribution. In the scenario generation, T has to be rounded in order to be a discrete number of weeks. In this way, the probability of $T = 0$ is finite, meaning that there is a chance that the crisis ended on the 12th of August 2011 and therefore some scenarios should begin with values simulated from the upturn model. This probability, which is the probability that the current regime ended last week, is of course very little.

5.1.1 Estimating the Parameter λ

The parameter λ has to be estimated for both upturn and downturn. The estimates of λ will have a large influence on the generated scenarios. Hence, estimating the λ 's solely based on the three upturns and two downturns contained in the data period is too erratic.

The λ parameters are instead estimated from the historical phase lengths of

Peak/trough	Time	Length upturn	Length downturn
P	1957M3		
T	1958M6		15 months
P	1960M3	21 months	
T	1963M1		34 months
P	1964M2	13 months	
T	1965M8		18 months
P	1969M6	46 months	
T	1971M8		26 months
P	1973M11	27 months	
T	1975M6		19 months
P	1979M12	54 months	
T	1982M12		36 months
P	1984M9	21 months	
T	1986M12		27 months
P	1990M6	42 months	
T	1993M8		38 months
P	1995M1	17 months	
T	1996M4		15 months
P	1997M11	19 months	
T	1999M1		14 months
P	2000M8	19 months	
T	2001M12		16 months
P	2008M2	86 months	
T	2009M4		14 months
P	2011M3	23 months	
	Mean	32.3 months	22.7 months
	Median	22.0 months	18.5 months
	Std. deviation	21.2 months	9.1 months

Table 5.1: Historical turning points and phase lengths of the reference series for the total OECD area.[22]

the reference series for the total OECD area since 1957. The reason for using the turning points of the reference series instead of the CLI-series is simply the availability of the data. A comparison of figure 4.1 on page 48 and figure 4.2 on page 50 gives no reason to assume that this should introduce any bias.

From table 5.1 it appears that the mean phase length has been 32.3 months for upturns and 22.7 months for downturns equalling 140 weeks and 98 weeks respectively.¹ Thus, the average upturns has lasted ten months longer than the average downturns. This agrees with the majority of the past twelve years being characterised as expansion. The probability of having to simulate from an upturn should consequently be larger than for a downturn.

The large difference in the average phase lengths is not reflected in the median. Given the empirical averages the medians should be 22.4 and 15.7 months respectively according to (5.4). The empirical medians are seen not to be too far from the expected values, which supports the distribution assumption. The empirical standard deviations are however much smaller than the averages, which does not agree with an exponential distribution (5.3). Overestimating the standard deviation of the phase lengths by applying an exponential distribution assumptions is not a problem, since it only leads to scenarios that capture a slightly broader range of variation.

From the estimated average phase lengths in weeks the following estimates of λ are obtained: $\lambda_{up} = 1/140$ and $\lambda_{down} = 1/98$.

5.2 Generating Scenarios

The algorithm used for generating scenarios is summarised in pseudo code in figure 5.1. The number of scenarios is $N_{scenario} = 10,000$ and the total number of weeks in each scenario is $N_{week} = 5 \times 52 = 260$. I_{up} is a binary variable indicating whether the economy is currently in an upturn or not. I_{up} is initialised as zero in all simulations as the economy was in a recession on the 12th of August 2011, where the scenarios begin. Also the number of weeks simulated in the current scenario, N , is initially zero.

As long as the number of weeks simulated in the current scenario is less than 260, a random number n specifying the length in weeks of the current regime

¹ $\frac{32.3 \times 52}{12} = 140$ and $\frac{22.7 \times 52}{12} = 98$

```

for  $i$  in  $1 : N_{scenario}$ 

  Initialise  $N = 0$  and  $I_{up} = 0$ 

  while  $N < N_{week}$ 

    Generate parameter  $n \in \text{Exponential}(I_{up}\lambda_{up} + (1 - I_{up})\lambda_{down})$  and round off to nearest integer

    Simulate principal components  $\min(n, N_{week} - N)$  steps using up- or downturn model depending on  $I_{up}$ 

    Transform simulated components to weekly log returns

    Set  $N = N + n$  and  $I_{up} = \neg I_{up}$ 

  endwhile

  Combine simulated weekly log returns to one time series

endfor

Transform weekly log return series to weekly index values

```

Figure 5.1: Scenario generation algorithm in pseudo code.

is generated from an exponential distribution with parameter λ depending on the current regime. n is then rounded to the nearest integer. Then the ARCH-models describing the first three principal components of the current regime are simulated n steps or the remaining number of steps in the current scenario.

After the simulated components have been transformed to weekly log returns, the number of weeks simulated, N , is updated and the indicator variable I_{up} is negated, such that the next regime is the opposite type. If the number of weeks simulated is less than 260, the procedure is repeated starting with generating a new n . When $N = 260$ the simulated log returns are pasted to one time series. After the simulation of $N_{scenario}$ scenarios, all the log return series are transformed to weekly index values.

5.3 Assessing Scenarios

Ten of the 10,000 scenarios generated using the pseudo code in figure 5.1 on the preceding page are visualised in figure 5.2 together with the 5%, 25%, 50%, 75%, and 95%-quantiles (black lines). The median is the solid black line. The red line is the scenario containing the maximum drawdown for each index. The median, interquartile range, and return is summarized for each 52 weeks in table A.3 on page 94 in appendix A.

The ten scenarios visualised are stratified according to the final value ranging from the scenario ending at the lowest value to the 999th highest value. The scenarios ending at the highest values are not visualised because they end at extremely high values and consequently mess up the scale of the second axis. If KAXGI ends at a value of 20,000 and the other indices likewise, then it does not really matter which index that was invested in, since the profit will be huge anyway. The influence of these few very optimistic scenarios on the quantiles is not a problem, since the quantiles are robust.

The distribution of the final values of the scenarios indexed by the value on the 12th of August 2011 are displayed in figure 5.3 excluding the 1% highest values. The 5%, 25%, 50%, 75%, and 95%-quantiles are also shown (blue lines).

In the process of modelling the data and generating scenarios several approximations that influence the accuracy of the scenarios have been made. This includes the error that arises from the smoothing due to the weekly sampling of a continuous process. From the communalities in table 4.5 and 4.6 on page 57 it appeared that the worst described index is the CSYHYI with communalities of about 60%. Consequently, the scenarios for this particular index cannot be assumed to be very accurate. As far as the rest of the indices, there have not been any indications that accuracy is compromised.

Some of the scenarios are seen to be very volatile, while others are more steady. Some scenarios contain large setbacks and others contain steep inclines, just as the real data. The scenarios for the stock indices are seen to be much more volatile than for the bond indices. This observation is also consistent with the observed behaviour of the real data. By choosing ARCH-models the scenarios are also guaranteed to exhibit some level of volatility clustering.

None of the indices provide a possibility of a riskless profit. Only the medians of the bond scenarios end up at values higher than the starting value, with

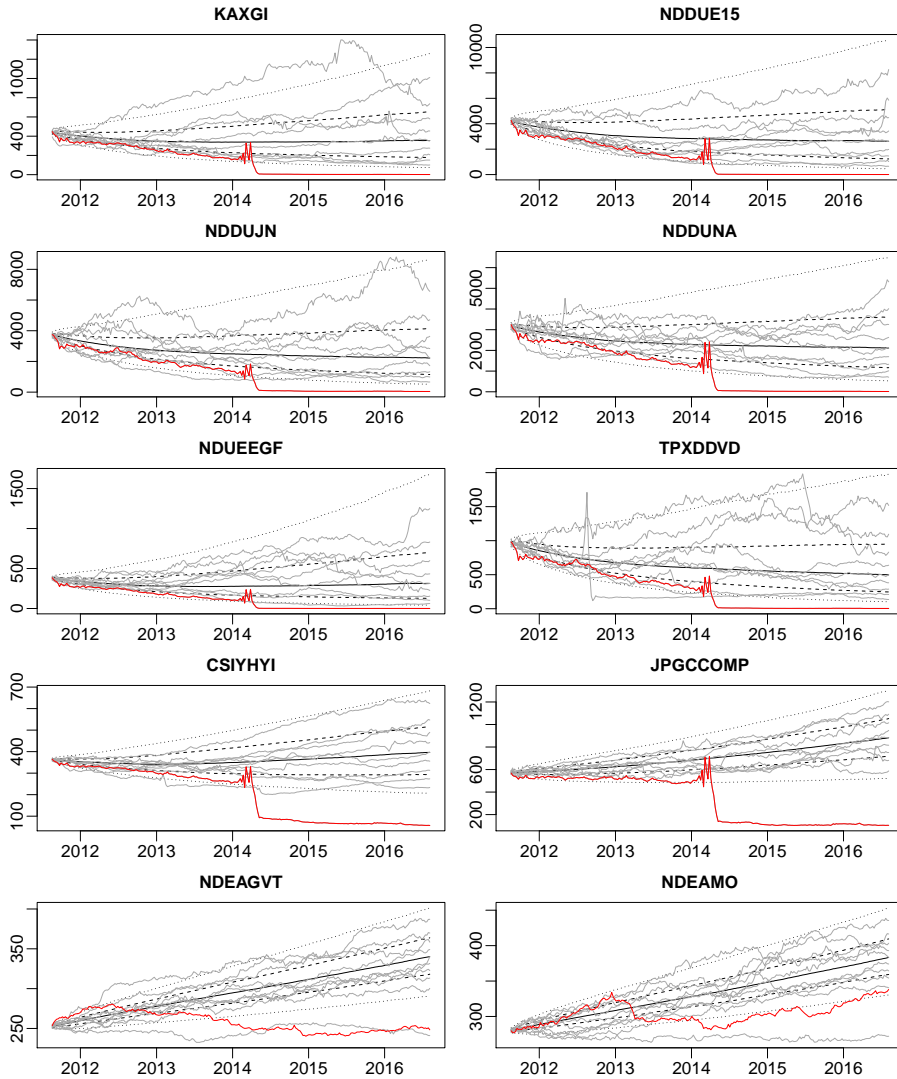


Figure 5.2: Ten of the 10,000 simulated scenarios (grey) together with the 5%, 25%, 50%, 75%, and 95%-quantiles (black), and the maximum drawdown (red).

JPGCCOMP offering the highest median return (54%). This is despite the fact that all indices except for TPXDDVD have provided positive total returns over the twelve year data period. The average yearly returns through the data period can be found in table A.2 on page 93 in appendix A. This is the result of most of the scenarios beginning in a downturn with an average length of 15 months

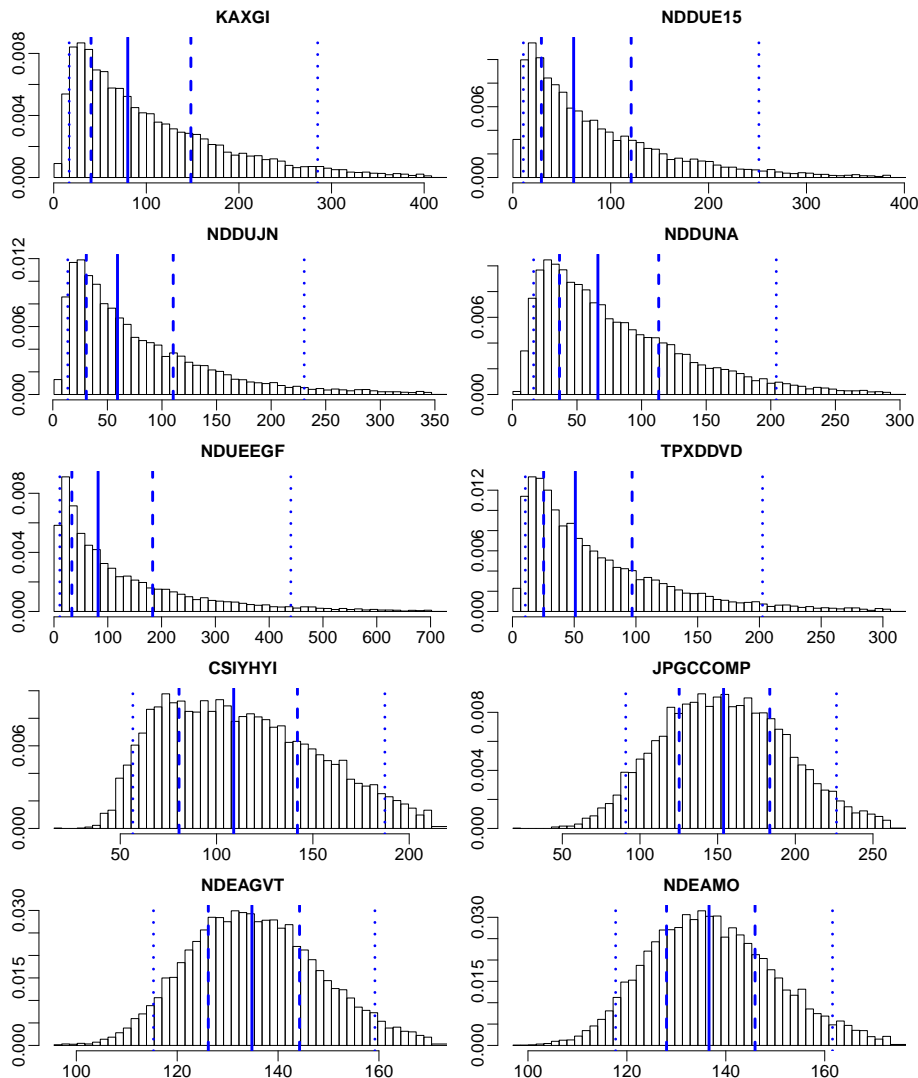


Figure 5.3: Density histograms of the final values of the 10,000 simulated scenarios excluding the 1% highest final values. The values are indexed by the value of the indices on the 12th of August 2011. The blue lines are the 5%, 25%, 50%, 75%, and 95%-quantiles.

out the total scenario length of 60 months.

The scenarios for the two Danish bond indices (NDEAGVT and NDEAMO)

are closest to offering a riskless profit with only a few scenarios containing a very limited loss. For instance, the probability of an investment in NDEAGVT resulting in a gain is more than 95% and worst case is loosing less than 15% of the initial investment. This is however balanced by the observation that the more secure bond indices provide no way near the same profit opportunities as the more volatile stock indices. The median returns of the Danish bond indices are very high compared to the low risk, but this reflects the fact that both indices have almost doubled in value over the past twelve years.

Since the data was modelled as a multivariate distribution using principal components, consistency cannot be troublesome either. The only way consistency can be compromised with this approach is due to lack of accuracy. And as already argued, there are no signs of lack of accuracy. For instance, the scenario containing the maximum drawdown is seen to decline rapidly in the first part of 2014 for all the indices except the stable Danish bond indices. The drop happens at the same point in time for all the indices due to the strong correlations. The magnitude of the drop is however different from one index to another. To summarise, both accuracy, correctness, and consistency seems approved.

```
Initialise  $MDD = 0$  and  $peak = 0$ 

for  $i$  in  $1 : N$ 

  if  $x_i > peak$ 
     $peak = x_i$ 

  else
     $DD = (peak - x_i) / peak$ 

    if  $DD > MDD$ 
       $MDD = DD$ 

    endif

  endif

endif

endfor
```

Figure 5.4: Pseudo code for computing the maximum drawdown of a variable X .

5.3.1 Maximum Drawdown

The maximum drawdown is the largest relative decline from a historical peak in the index value. The pseudo code for determining the maximum drawdown of a variable X is shown in figure 5.4.[31]

The scenario containing the maximum drawdown was visualised for each index in figure 5.2 on page 69. Maximum drawdown is a risk measure and is sometimes used instead of the standard deviation. The maximum drawdown is particularly interesting in connection with stop-loss mechanisms. The size of the maximum drawdown is important in order to secure that the stop-loss mechanisms are never activated, since an activation often causes a huge loss.

The 5%, 25%, 50%, 75%, and 95%-quantiles of the distribution of the maximum drawdowns for the 10,000 generated scenarios are shown in figure 5.5 together with the worst case maximum drawdown (red). The median is the solid black line. The worst case maximum drawdown is seen to increase quite rapidly in the beginning because the scenarios begin in a downturn.

Generally speaking, what characterises a scenario is how many shocks there will come, and when they will come. For instance, if a shock is defined as a drawdown of minimum 40%, then probability that there will come at least one shock within the next five years is more than 50% for KAXGI, while the chance that there will come such a shock before 2013 is more than 25%.

5.4 Bootstrapping

The simplest bootstrap is to sample each observation in the scenarios at random from the weekly log returns with replacement. The strength of this method, besides being very simple, is that it captures the historical correlation between the indices perfectly, implying that consistency is guaranteed. The great weakness is that the method is unable to capture time dependent effects such as cross-correlations at other lags than 0, autocorrelation, and volatility clustering.

It has already been assumed that there is no significant cross-correlation at other lags than 0, nor any significant autocorrelation, but the volatility clustering observed has to be reproduced for the scenarios to be correct. A possible

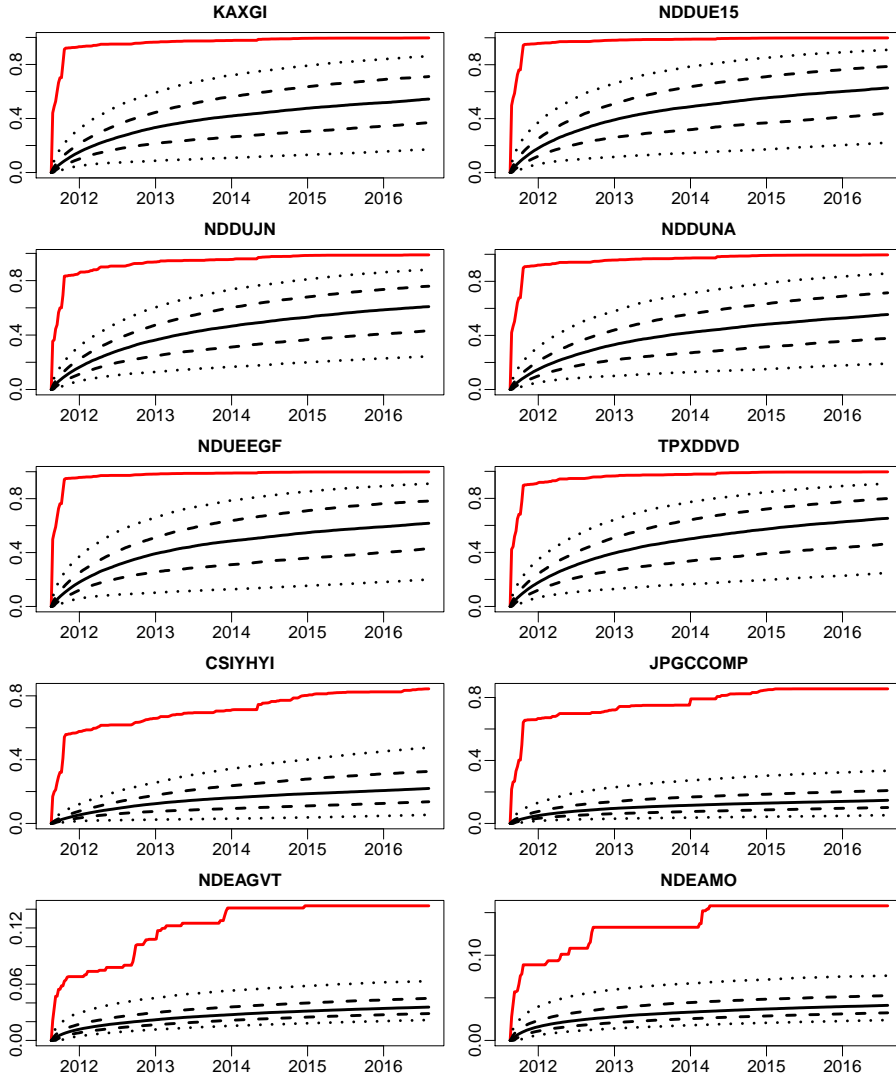


Figure 5.5: The 5%, 25%, 50%, 75%, and 95%-quantiles of the distribution of the maximum drawdowns for the 10,000 simulated scenarios as percentages together with the worst case maximum drawdown (red).

solution could be, instead of sampling completely at random, to specify probability weights so that it is most likely that each sample is within a few weeks of the previous sample.

The question is, how wide the week span can be, if the bootstrapped scenarios are supposed to reproduce volatility clustering correctly. Laurini suggests that the maximum number of days extreme information is related is 50 bank days corresponding to ten weeks.[11] Thus, probability weights that makes it most likely that the next sample will be within ten weeks of the previous sample should be adequate.

A possibility is to apply normal probability weights with the last row index as mean and a standard deviation of five. In this way, the probability that the next sample will be drawn within ten weeks of the last sample is almost 95%. The first sample in each scenario will then most likely be drawn from the last ten observations in the dataset. This bootstrapping procedure is summarised in pseudo code in figure 5.6. It is actually a random walk in the (row) index with reflection at the endpoints of the interval $[0; 658]$.

```

for  $i$  in  $1 : N_{scenario}$ 

  Sample  $index_1$  from interval  $[0; 658]$  assuming normal
  probability weights with  $\mu = 658$  and  $\sigma = 5$ 

  for  $j$  in  $2 : N_{week}$ 

    Sample  $index_j$  from interval  $[0; 658]$  assuming normal
    probability weights with  $\mu = t_{j-1}$  and  $\sigma = 5$ 

  endfor

  Single out weekly log returns according to indices in
   $index$  vector

endfor

Transform weekly log return series to index observations

```

Figure 5.6: Bootstrapping algorithm in pseudo code.

The initialisation of the bootstrapping algorithm is an issue on its own. The approach used in the pseudo code, where the first sample is most likely to be drawn within ten weeks of the last observation in the dataset, is the most meaningful. The problem with this approach is that since the current recession has only lasted a few months, the probability of escaping it within a few weeks will be too large. As discussed in section 5.1, the probability of escaping the current regime should be independent of how long the crisis has lasted.

5.5 Bootstrapping Performance

The result of bootstrapping 10,000 scenarios using this procedure appears from figure 5.7. Ten of the 10,000 bootstrapped scenarios are visualised together with the 5%, 25%, 50%, 75%, and 95%-quantiles (black lines). The median is the solid black line. The red line is the scenario containing the maximum drawdown for each index. The ten scenarios visualised are stratified according to the final value ranging from the scenario ending at the lowest value to the 999th highest.

From a comparison of these scenarios with those in figure 5.2 on page 69 it appears that the bootstrapped scenarios are more smooth, which could indicate that the standard deviation in the bootstrapping algorithm is chosen too large to get proper representation of the observed volatility clustering. This could also be the reason why only a few of the visualised bootstrapped scenarios have steep inclines and declines like the real data. Still, the quantile bands are seen to be wider for the bootstrapped scenarios meaning that these are more dispersed. Once again, the scenarios for the stock indices are more volatile than for the bond indices.

Density histograms of the final values of the scenarios indexed by the value on the 12th of August 2011 are displayed in figure 5.8 excluding the 1% highest final values. The 5%, 25%, 50%, 75%, and 95%-quantiles are also shown.

The distributions of the final values of the bootstrapped scenarios are slightly different from those in figure 5.3 on page 70. Among others due to the larger standard deviation of the bootstrapped scenarios. Apart from the larger standard deviations, the shape of most of the distributions appear to be the same except for CSIHYYI and JPCCOMP. This conclusion is supported by the plot in figure A.4 on page 96 in appendix A of the first 99 quantiles of the final values of the scenarios generated using the model as a function of the first 99 quantiles of the final values of the bootstrapped scenarios.

Most of the distributions are shifted to the right for the bootstrapped scenarios, meaning that the median returns of the bootstrapped scenarios are larger. The reason is that the initialisation of the bootstrapping procedure at the end of the data period implies that the probability of sampling from the relatively steep inclines in the periods from 2003 to 2008 and 2009 to 2010 will be large compared to the probability of sampling from the short intermediate downturn.

The median, interquartile range, and return is summarized for each 52 weeks in table A.4 on page 95 in appendix A. The yearly returns of the bootstrapped

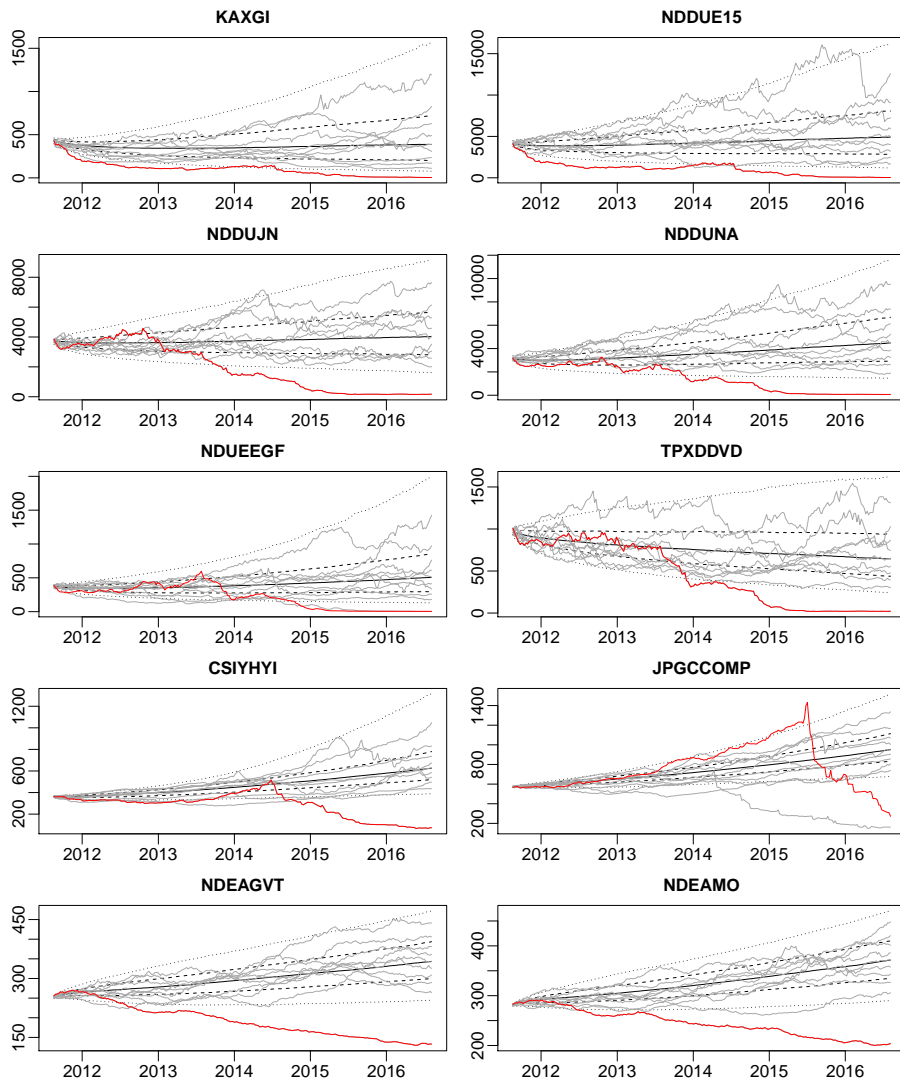


Figure 5.7: Ten of the 10,000 bootstrapped scenarios (grey) together with the 5%, 25%, 50%, 75%, and 95%-quantiles (black), and the maximum drawdown (red).

scenarios are seen to be much closer to the historical averages in table A.2 on page 93, than for the scenarios generated using the model.

None of the indices provide a possibility of a riskless profit according to the bootstrapped scenarios. Only the medians of the KAXGI and the TPXDDVD

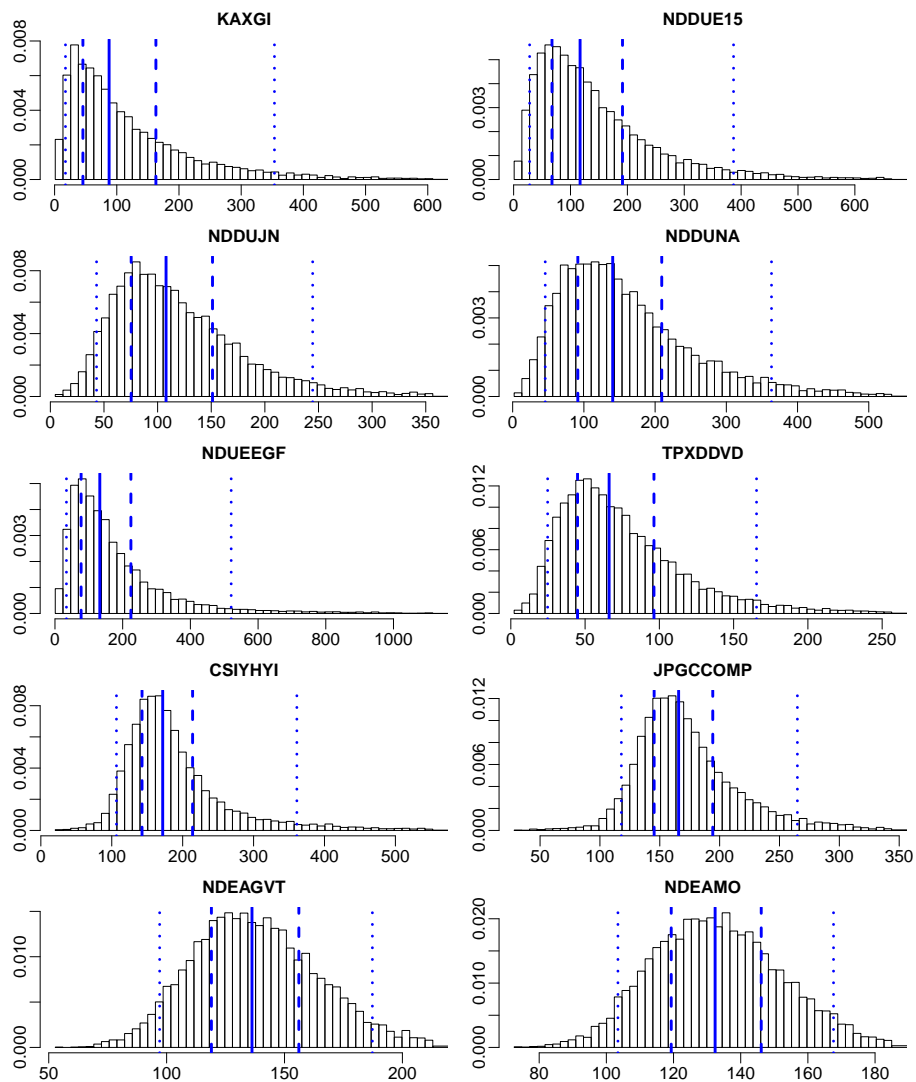


Figure 5.8: Density histograms of the final values of the 10,000 bootstrapped scenarios excluding the 1% highest final values. The values are indexed by the value of the indices on the 12th of August 2011. The blue lines are the 5%, 25%, 50%, 75%, and 95%-quantiles.

scenarios end up at values lower than the starting value. This time CSIIHYI offers the highest median return (72%).

The error that arises from the smoothing due to the weekly sampling of a continuous process is also present, when bootstrapping. But the bootstrapping procedure does not include further approximations that could compromise accuracy. As already argued, bootstrapping captures the historical correlation perfectly, which guarantees consistency, unless accuracy is compromised.

5.5.1 Maximum Drawdown

The 5%, 25%, 50%, 75%, and 95%-quantiles of the distribution of the maximum drawdowns for the 10,000 bootstrapped scenarios are shown in figure 5.9 together with the worst case maximum drawdown (red). The median is the solid black line.

The worst case maximum drawdown is seen not to increase as rapidly as in figure 5.5 on page 73. A part of the reason is that even though the bootstrapped scenarios most likely also begin in a downturn, the probability of escaping the downturn within a few weeks is much higher since the current downturn has only lasted a few months. However, it could also be evidence of the lack of volatility clustering implying that the bootstrapped scenarios do not have the same rapid drawdowns. Even though the maximum drawdowns are at least as severe, they might not happen as rapidly.

A large difference to the distribution of maximum drawdown for the scenarios generated using the model is the two Nordea indices. These are much more volatile when bootstrapping, which leads to more severe drawdowns than in figure 5.5 on page 73.

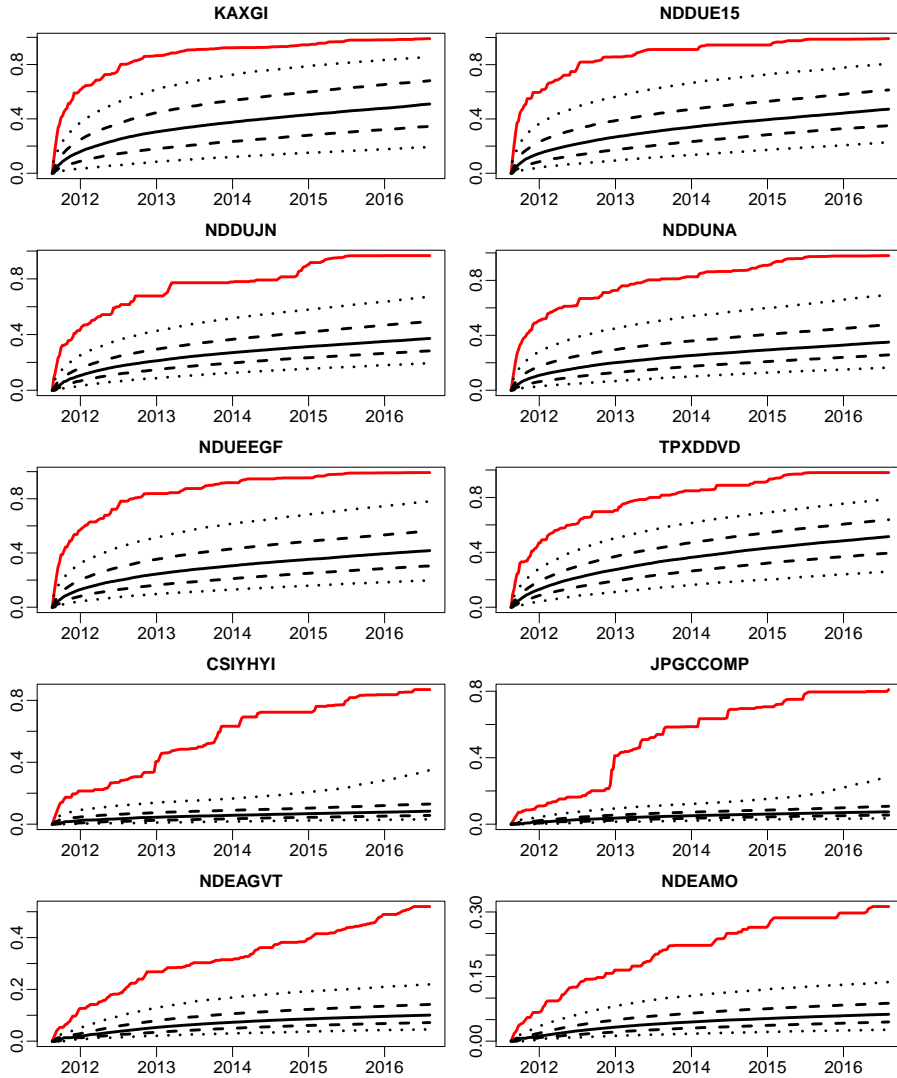


Figure 5.9: The 5%, 25%, 50%, 75%, and 95%-quantiles of the distribution of the maximum drawdowns for the 10,000 bootstrapped scenarios as percentages together with the worst case maximum drawdown (red).

CHAPTER 6

Summary and Conclusions

The distribution of the log returns was found to be a mixture of a normal distribution and a few extreme observations originating from a completely different distribution. Due to the presence of these extreme values, traditional measures rejected any resemblance with the normal distribution. Yet, more robust measures were able to find certain similarities.

The index data had a growing mean trend as well as a changing variance, autocorrelation, and cross-correlation. A log return transformation was applied to handle the growing mean and weekly sampling eliminated most of the autocorrelation.

The LIBOR-index was left out from the rest of the analysis, since it was evident from the data analysis that it was completely independent of the rest of the indices. LIBOR is a rate and should therefore have its own type of model. Furthermore, observations of the LIBOR-index were only available from the last part of the data period.

The main conclusion from the data analysis was that a model with constant coefficients would not be able to account for the changes in the underlying

dynamics of the dataset. A regime model, where upturns and downturns were modelled independently, was applied to handle the changing cross-correlation, and ARCH-models were employed within each regime to model the changing variance. Instead of modelling the weekly log returns the first three principal components of each regime were modelled.

The data period was partitioned into regimes based on the composite leading indicator for the total OECD area. Regardless of the method, the partitioning gives rise to an error as a consequence of the indices representing different parts of the world with different turning points. If the number of regimes or the partitioning had been insufficient, the estimated ARCH-models would have converged towards not being stationary.

Based on the estimated regime model's ability to reproduce the index data, the model was accepted for the scenario generation purpose. 10,000 scenarios were generated using the exponential distribution to simulate the length of each regime in a way that was perfectly consistent with the no-arbitrage principle and made initialisation easy.

The median returns of the simulated scenarios for the stock indices were much below the average returns over the data period as a result of the scenarios beginning in a recessive period. This was not the case for the more stable bond indices. The scenarios were found to satisfy the no-arbitrage principle and to reproduce the possibility of future crises, like those experienced in the estimation period, correctly.

A bootstrapping procedure was implemented to obtain a standard of comparison for the generated scenarios. The typical shortfall of bootstrapping is the inability to capture time dependent effects like volatility clustering and consequently the inability to reproduce the crises observed in the data. The bootstrapping procedure applied used normal probability weights centered around the last sample to capture the volatility clustering.

The initialisation of the bootstrapping procedure turned out to be quite problematic, since it had a large influence on the scenarios. The initialisation at the end of the data period was problematic, since the probability of escaping the current crisis within a few weeks was too large.

The bootstrapped scenarios were found to have median returns much closer to the average returns over the data period. The bootstrapped scenarios appeared to be smoother, but the overall dispersion of the bootstrapped scenarios was

larger, indicating an insufficient reproduction of volatility clustering. Still, the shape of most of the distributions of the final values was found to be the same as for the simulated scenarios. The distributions were simply more dispersed and shifted to the right towards the more positive historical returns.

6.1 Discussion

The scenarios simulated using the regime model were most successful at reproducing the steep inclines and declines observed in the real data. The ability to reproduce realistic crises appeared to be connected with the ability to reproduce volatility clustering correctly. The bootstrapping procedure could most likely be optimised by tuning the standard deviation of the normal distribution used to draw samples.

Furthermore, the bootstrapped scenarios were less capable of taking the initial situation into account, resulting in median returns that were (too) close to the historical averages.

In his thesis, Emil Ahlmann Østergaard implemented a bootstrapping procedure with uniform probability weights centered around the last sample index, also with a standard deviation of five weeks. The resulting scenarios were not very different from those obtained with normal probability weights.[25]

Østergaard also modelled the data using a regime based approach, but instead of merging the individual upturns and downturns, he modelled each regime independently. This led to the interesting observation that within each regime the majority of the principal components behaved like random walks and not ARCH-processes. This implies that the volatility clusters were coinciding with the regimes.

The scenarios generated by Østergaard using his regime model were not very different from those generated in the present thesis, at least not when comparing the distribution of the final values and maximum drawdowns. The main difference observed was a tendency towards slightly more steep inclines and declines in Østergaard's scenarios. The reason being that the merging of regimes done in this thesis causes a smoothing since the overall downturn cannot be as steep as the steepest downturn observed and vice versa.

A more thorough quantitative testing would be needed to decide whether the increased complexity of Østergaard's model resulted in scenarios that actually performs better than those generated with merged regimes.

6.2 Possibilities for Future Work

The finding by Østergaard that volatility clusters tend to overlap with the regime partitioning could be used to develop a bootstrapping procedure based on regimes. Sampling from one regime at a time should result in correct representation of the observed volatility clustering, at least if the overlap is large enough. A regime based procedure would also allow for the use of the exponential distribution to simulate the length of each regime, which would eliminate the initialisation difficulties.

Possible extensions of the work also include a quantification of the decision on the number of regimes and the partitioning of data for instance through application of likelihood ratio tests. This could entail an extension of the model to a Markov switching regime model. Inclusion of further asset classes in the model is also an obvious possibility for future work.

The inclusion of extreme values in the model estimation mainly results in a larger standard deviation. A study of the distribution of the most extreme returns could lead to an improved procedure for generating scenarios: if data was described in terms of a mixed distribution, a possibility could be to sample most observations from a model estimated from the normally distributed returns and then a few observations from another model estimated from an extreme value distribution.

Developing and performing a thorough testing procedure including out-of-the-sample and maybe live-sample testing of the model would also be a relevant problem for future work in the light of the intention to apply the generated scenarios to an asset allocation decision.

Bibliography

- [1] Preben Bertelsen and Anders Skytte Aalund. Revidering af nordeas indeksunivers. *Nordea Quantitative Research*, November 2006.
- [2] Zvi Bodie, Alex Kane, and Alan J. Marcus. *Investments and Portfolio Management*. McGraw-Hill/Irwin, 9 edition, 2011. ISBN 978-007-128914-6. Global Edition.
- [3] Michael J. Crawley. *Statistics - An introduction using R*. Wiley, 2005. ISBN 978-0-470-02298-6.
- [4] Jens Eising. *Lineær ALGEBRA*. Institut for Matematik, Danmarks tekniske universitet, 1 edition, 1999. ISBN 87-88764-39-7.
- [5] Bjarne Kjær Ersbøll and Knut Conradsen. *An Introduction to Statistics, Vol. 2*. Danmarks Tekniske Universitet, 6 edition, 2007.
- [6] Jürgen Franke, Wolfgang Karl Härdle, and Christian Matthias Hafner. *Statistics of Financial Markets*. Springer, 3 edition, 2011. ISBN 978-3-642-16520-7.
- [7] Eric Hillebrand. Neglecting parameter changes in garch models. *Journal of Econometrics*, 129:121–138, February 2005. URL <http://harrisd.net/papers/ARCHSV/FIGARCH/Variance%20Breaks/Hillebrand2005JoE.pdf>.
- [8] Donald A. Jackson. Stopping rules in principal component analysis: A comparison of heuristical and statistical approaches. *Ecology*, 74(8):2204–2214, 1993. URL <http://labs.eeb.utoronto.ca/jackson/ecology%2074.pdf>.
- [9] Tae-Hwan Kim and Halbert White. On more robust estimation of skewness and kurtosis. *Finance Research Letters*, 1:56–73, 2004.

- [10] Erwin Kreyszig. *Advanced Engineering Mathematics*. Wiley, 9 edition, 2006. ISBN 0-471-72897-7.
- [11] Fabrizio Laurini. Clusters of extreme observations and extremal index estimate in garch processes. *Studies in Nonlinear Dynamics & Econometrics*, 8(2):1–21, 2004.
- [12] Financial Times Lexicon. Embi. 2011. URL <http://lexicon.ft.com/Term?term=EMBI>. As of 24 October 2011.
- [13] Jesper Lund. Arch/garch modeller, April 2006. URL <http://www.jesperlund.com/mfok06/garch.pdf>.
- [14] Henrik Madsen. *Time Series Analysis*. Chapman & Hall, 2008. ISBN 978-1-4200-5967-0.
- [15] Henrik Madsen, Jan Nygaard Nielsen, Erik Lindström, Mikkel Baadsgaard, and Jan Holst. *Statistics in Finance*. Lund University, Lund Institute of Technology - Centre for Mathematical Sciences, 2004.
- [16] Kevin McGarigal, Sam Cushman, and Susan Stafford. *Multivariate Statistics for Wildlife and Ecology Research*. Springer, 1 edition, 2000. ISBN 978-0-387-98891-7.
- [17] J. P. Morgan. Jpm global high yield bond fund. 2011. URL <http://documents.financialexpress.net/literature/1804971.pdf>.
- [18] MSCI. Index definitions. January 2012. URL <http://www.msci.com/products/indices/tools/index.html>.
- [19] MSCI. Index country membership. January 2012. URL http://www.msci.com/products/indices/tools/index_country_membership/.
- [20] OECD. Oecd system of composite leading indicators. November 2008. URL <http://www.oecd.org/dataoecd/26/39/41629509.pdf>.
- [21] OECD. Methodological notes. August 2011. URL <http://www.oecd.org/dataoecd/22/47/44728410.pdf>.
- [22] OECD. Oecd total area - reference turning points. December 2011. URL http://www.oecd.org/document/6/0,3746,en_2649_34349_35726918_1_1_1_1,00.html.
- [23] NASDAQ OMX. Rules for the construction and maintenance of the nasdaq omx and oslo børs all-share, benchmark, tradable and sector indexes. Version 1.9.2, March 2010. URL https://indexes.nasdaqomx.com/docs/Methodology_NASDAQ_OMX_and_Oslo_Bors_Equity_Indexes.pdf.
- [24] Jim Pitman. *Probability*. Springer, 2006. ISBN 0-387-97974-3.

-
- [25] Emil Ahlmann Østergaard. Scenario generation for financial market indices. Bachelor's thesis, Technical University of Denmark, January 2012. IMM-BSc-2012-3.
- [26] Wikinvest. Libor. October 2011. URL <http://www.wikinvest.com/rate/LIBOR>.
- [27] Wikipedia. Asset allocation. 2011. URL http://en.wikipedia.org/wiki/Asset_allocation. Last modified on 22 August 2011 at 06:38.
- [28] Wikipedia. London interbank offered rate. 2011. URL http://en.wikipedia.org/wiki/London_Interbank_Offered_Rate. Last modified on 16 September 2011 at 16:40.
- [29] Wikipedia. Money market. 2011. URL http://en.wikipedia.org/wiki/Money_market. Last modified on 4 October 2011 at 11:22.
- [30] Wikipedia. Principal component analysis. 2011. URL http://en.wikipedia.org/wiki/Principal_component_analysis. Last modified on 4 December 2011 at 14:19.
- [31] Wikipedia. Drawdown (economics). 2011. URL [http://en.wikipedia.org/wiki/Drawdown_\(economics\)](http://en.wikipedia.org/wiki/Drawdown_(economics)). Last modified on 21 August 2011 at 13:24.

APPENDIX A

Additional Graphs and Tables

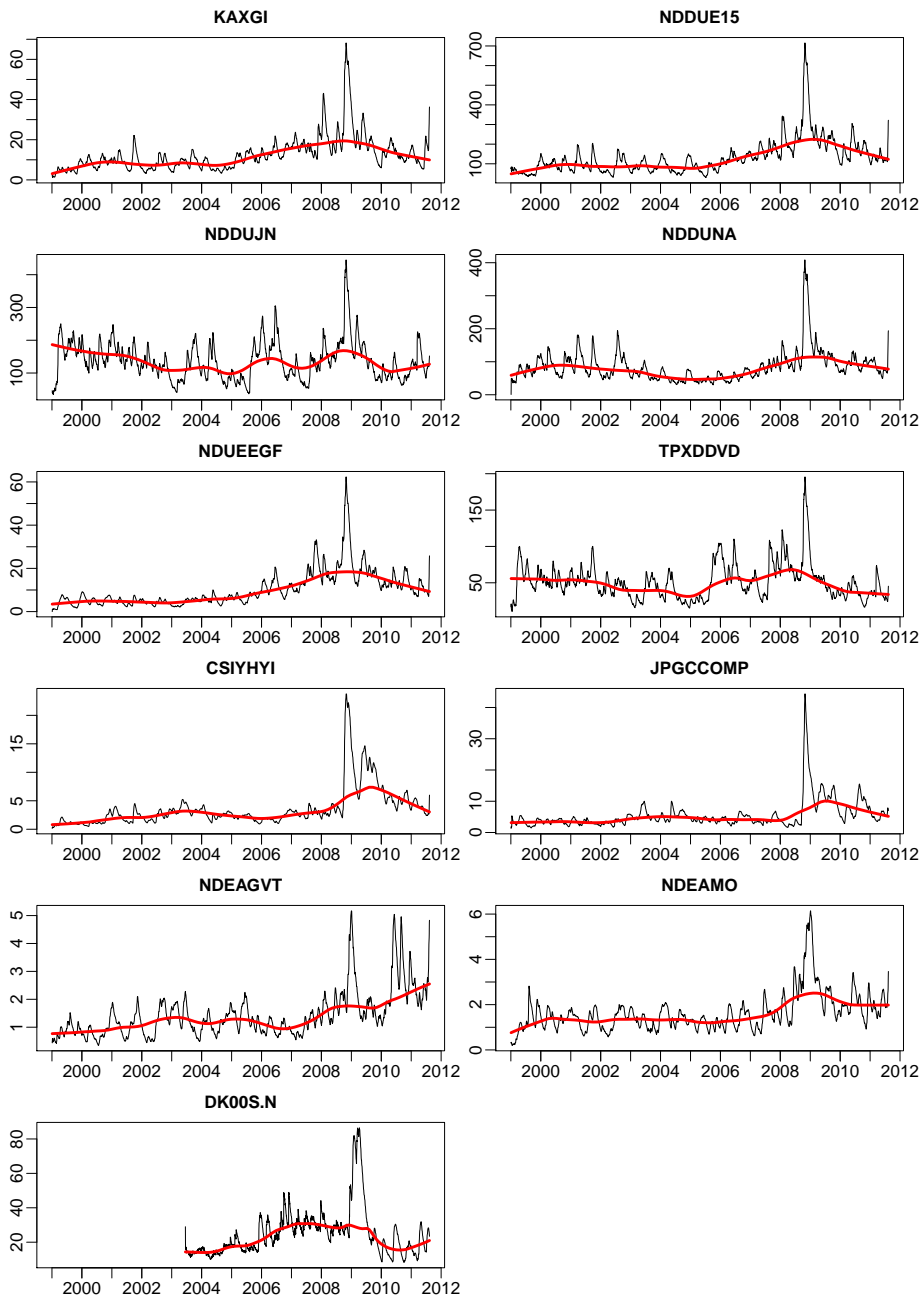


Figure A.1: The standard deviation of the eleven indices estimated recursively as a function of time using exponential weights (black lines). The weight assigned to the t 'th observation at time T is $\lambda_{t,T} = c \cdot 0.95^{T-t}$, where $t \leq T$ and $c = \frac{1-0.95}{1-0.95^T}$ is a normalising factor. The red line is a smoothing of the recursively estimated standard deviation showing the local mean level.

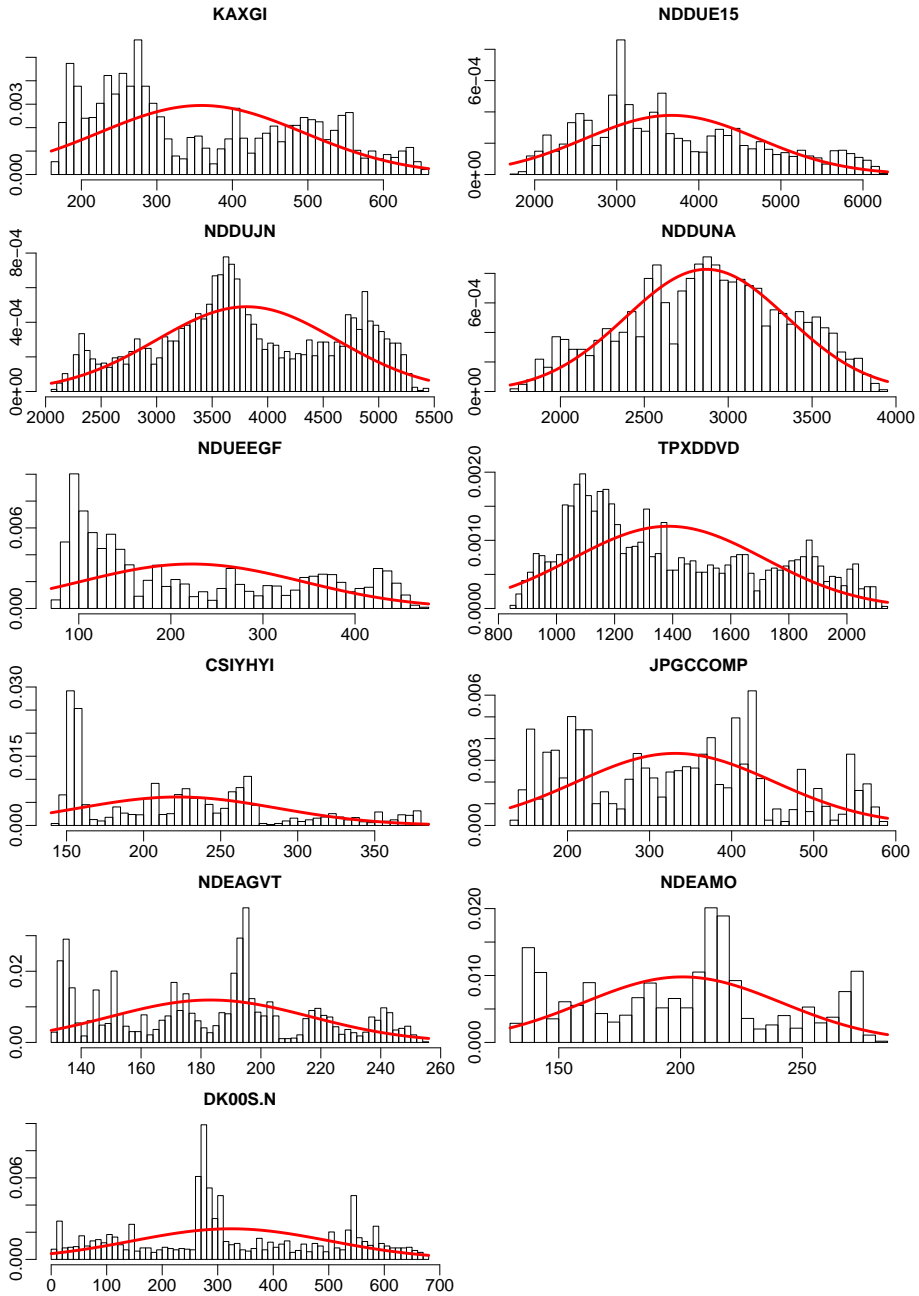


Figure A.2: Density histograms of the daily values of the eleven indices together with the theoretical density of a normal distribution with the same mean and variance as the index data (red lines).

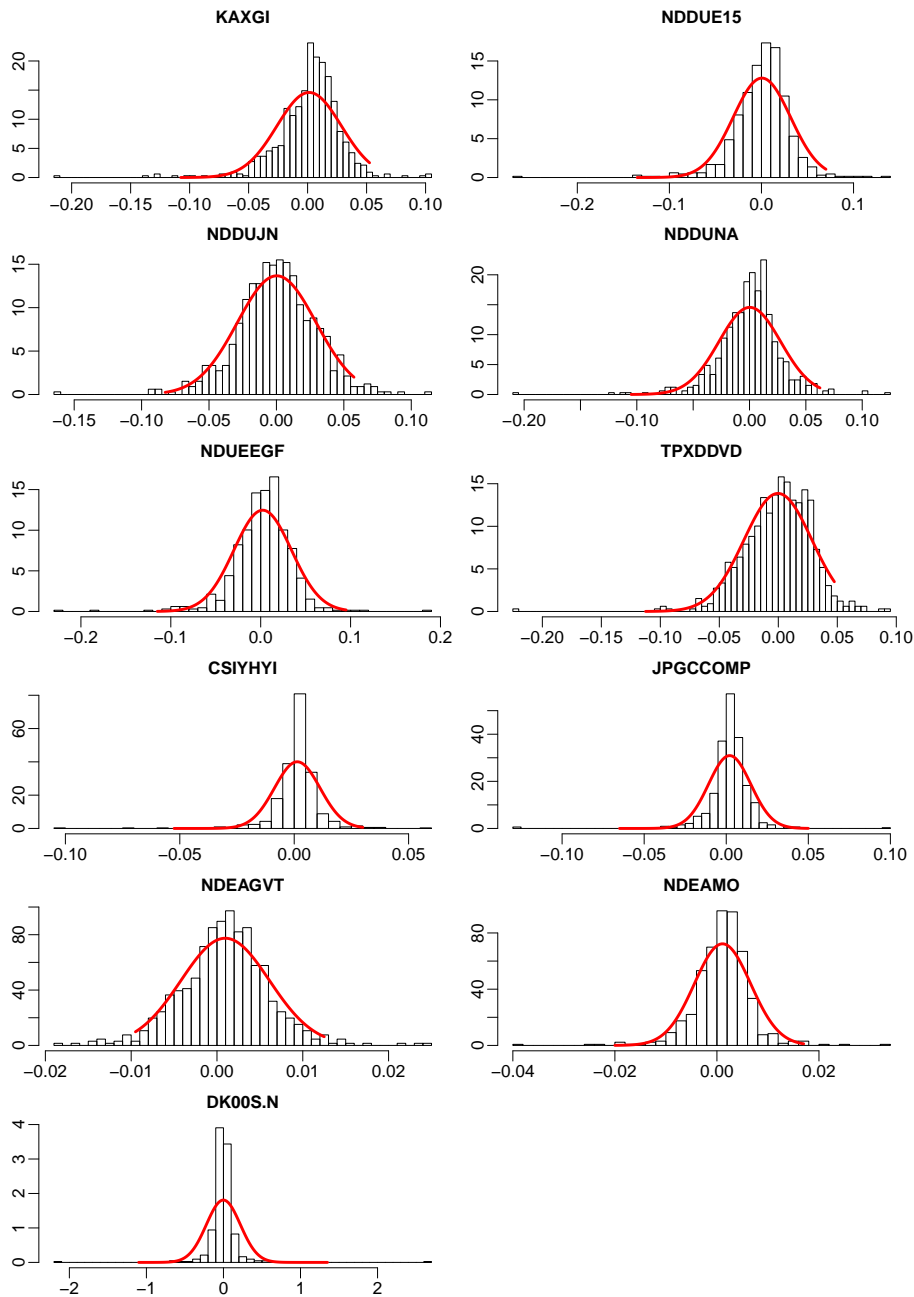


Figure A.3: Density histograms of the weekly log return values of the eleven indices together with the theoretical density of a normal distribution with the same mean and variance as the weekly log return data (red lines).

	Upturn		Downturn	
	Mean	Std. deviation	Mean	Std. deviation
KAXGI	0.0044	0.022	-0.0057	0.036
NDDUE15	0.0037	0.024	-0.0071	0.043
NDDUJN	0.0032	0.027	-0.0067	0.032
NDDUNA	0.0028	0.022	-0.0057	0.036
NDUEEGF	0.0063	0.025	-0.0073	0.042
TPXDDVD	0.0028	0.025	-0.0076	0.035
CSIYHYI	0.0027	0.007	-0.0019	0.014
JPGCCOMP	0.0028	0.010	-0.0005	0.018
NDEAGVT	0.0007	0.005	-0.0015	0.006
NDEAMO	0.0010	0.005	-0.0014	0.007

Table A.1: Estimated mean and standard deviation of the ten log return series during upturn and downturn respectively.

	Total return	Yearly return
KAXGI	140%	6.9%
NDDUE15	37%	2.5%
NDDUJN	16%	1.2%
NDDUNA	18%	1.3%
NDUEEGF	331%	11.5%
TPXDDVD	-17%	-1.5%
CSIYHYI	141%	6.9%
JPGCCOMP	298%	10.9%
NDEAGVT	88%	5.0%
NDEAMO	106%	5.7%

Table A.2: The total return over the period from 1st of January 1999 to 12th of August 2011 and the corresponding average yearly return assuming continuous compounding.

	KAXGI			NDDUE15		
Weeks	Q_2	$Q_3 - Q_1$	R_{1Y}	Q_2	$Q_3 - Q_1$	R_{1Y}
52	363	144	-18%	3269	1477	-22%
104	339	242	-7%	2903	2281	-11%
156	340	328	+0%	2766	2908	-5%
208	344	404	+1%	2678	3459	-3%
260	353	477	+3%	2619	3865	-2%
	NDDUJN			NDDUNA		
Weeks	Q_2	$Q_3 - Q_1$	R_{1Y}	Q_2	$Q_3 - Q_1$	R_{1Y}
52	2914	1110	-22%	2594	935	-19%
104	2555	1747	-12%	2344	1465	-10%
156	2406	2228	-6%	2236	1872	-5%
208	2295	2662	-7%	2163	2183	-3%
260	2220	2983	-3%	2113	2450	-2%
	NDUEEGF			TPXDDVD		
Weeks	Q_2	$Q_3 - Q_1$	R_{1Y}	Q_2	$Q_3 - Q_1$	R_{1Y}
52	297	147	-22%	730	301	-25%
104	275	258	-7%	624	455	-15%
156	282	365	+3%	574	559	-8%
208	296	471	+5%	534	641	-7%
260	314	574	+6%	497	699	-7%
	CSIYHYI			JPGCCOMP		
Weeks	Q_2	$Q_3 - Q_1$	R_{1Y}	Q_2	$Q_3 - Q_1$	R_{1Y}
52	344	52	-5%	607	93	+6%
104	344	102	+0%	659	150	+9%
156	359	147	+4%	724	206	+10%
208	376	189	+5%	799	268	+10%
260	397	224	+6%	882	334	+10%
	NDEAGVT			NDEAMO		
Weeks	Q_2	$Q_3 - Q_1$	R_{1Y}	Q_2	$Q_3 - Q_1$	R_{1Y}
52	271	14	+8%	301	18	+7%
104	288	23	+6%	320	27	+6%
156	305	31	+6%	340	34	+6%
208	322	38	+6%	361	42	+6%
260	340	46	+6%	383	50	+6%

Table A.3: Median, interquartile range, and return the last 52 weeks for each 52 weeks since 12th of August 2011, where the generated scenarios begin.

	KAXGI			NDDUE15		
Weeks	Q_2	$Q_3 - Q_1$	R_{1Y}	Q_2	$Q_3 - Q_1$	R_{1Y}
52	348	143	-21%	3883	1603	-8%
104	344	236	-1%	4095	2444	+5%
156	357	330	+4%	4328	3374	+6%
208	376	423	+5%	4666	4323	+8%
260	387	519	+3%	4925	5227	+6%
	NDDUJN			NDDUNA		
Weeks	Q_2	$Q_3 - Q_1$	R_{1Y}	Q_2	$Q_3 - Q_1$	R_{1Y}
52	3600	1020	-4%	3084	1008	-3%
104	3676	1531	+2%	3380	1639	+10%
156	3801	1992	+3%	3719	2272	+10%
208	3925	2414	+3%	4135	3000	+11%
260	4041	2841	+3%	4484	3761	+8%
	NDUEEGF			TPXDDVD		
Weeks	Q_2	$Q_3 - Q_1$	R_{1Y}	Q_2	$Q_3 - Q_1$	R_{1Y}
52	347	127	-9%	837	254	-14%
104	368	208	+6%	772	344	-7%
156	405	305	+10%	726	415	-6%
208	451	420	+11%	688	466	-5%
260	507	563	+10%	645	502	-6%
	CSIYHYI			JPGCCOMP		
Weeks	Q_2	$Q_3 - Q_1$	R_{1Y}	Q_2	$Q_3 - Q_1$	R_{1Y}
52	384	42	+5%	629	47	+10%
104	427	75	+11%	690	86	+10%
156	483	118	+13%	764	137	+11%
208	551	176	+14%	851	202	+11%
260	626	260	+14%	950	281	+12%
	NDEAGVT			NDEAMO		
Weeks	Q_2	$Q_3 - Q_1$	R_{1Y}	Q_2	$Q_3 - Q_1$	R_{1Y}
52	273	33	+8%	300	25	+7%
104	288	49	+5%	314	37	+5%
156	305	64	+6%	331	49	+5%
208	323	78	+6%	351	62	+6%
260	343	94	+6%	371	75	+6%

Table A.4: Median, interquartile range, and return the last 52 weeks for each 52 weeks since 12th of August 2011, where the bootstrapped scenarios begin.

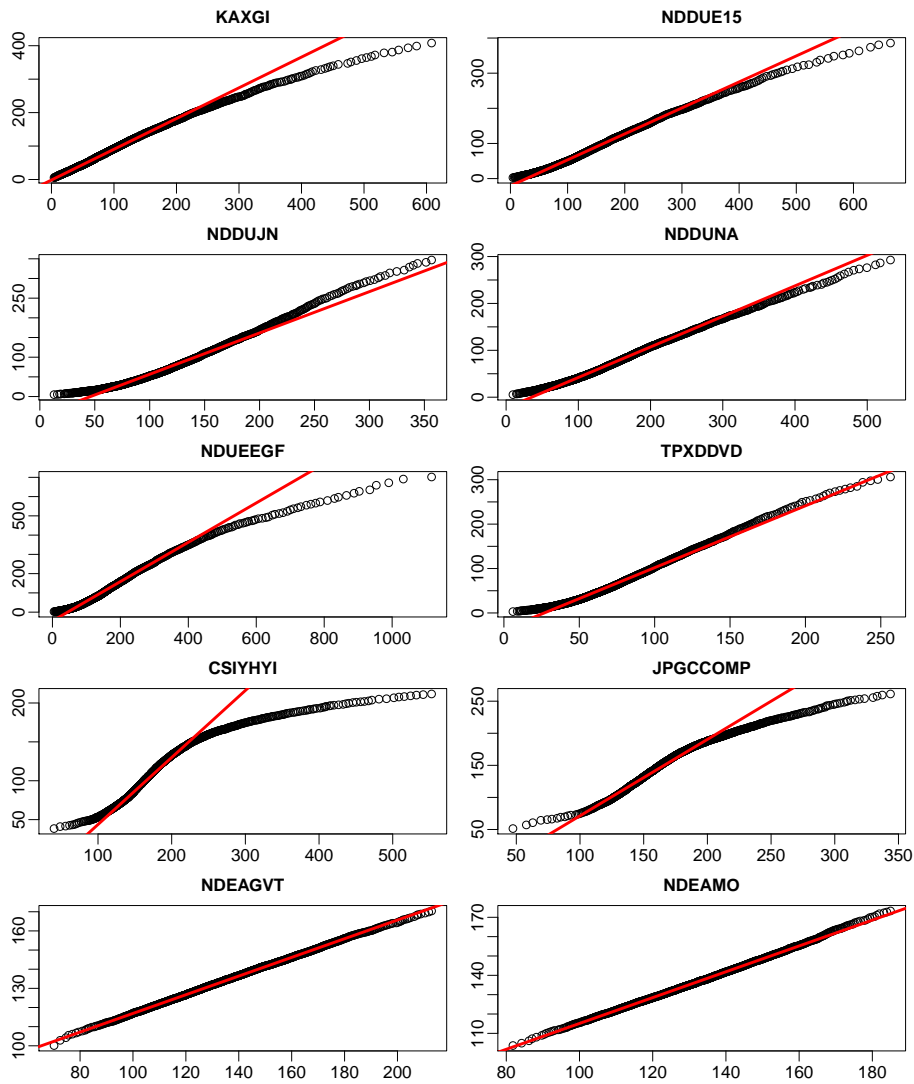


Figure A.4: The first 99 quantiles of the distribution of the final values of the scenarios generated using the model as a function of the first 99 quantiles of the final values of the bootstrapped scenarios. The red line passes through the first and third quartiles.

APPENDIX B

R-code

B.1 R-code for Chapter 3

```
#####
## Filename: DataAnalysis.R                      Date: 20-01-2012 ##
## Author: Peter Nystrup s082634                ##
## Description: R-code for Chapter 3            ##
#####

#Loading packages
library(tseries)
library(SDMTools)
library(TSA)

#Reading in the data
data<-read.csv2('Data.csv',header=T);
data$Date<-as.Date(data$Date,"%d-%m-%Y")
data$DKOOS.N[1:1161]<-NA
data$DKOOS.N<-as.numeric(data$DKOOS.N)
n=length(data[,1])

##Data Reconstruction##
#Saving original NDUEEGF-data
original_data=data$NDUEEGF[1:520]

#Reconstruction using linear interpolation
data$NDUEEGF[c(F,diff(data$NDUEEGF[1:520])==0,rep(F,n-520))]=NA
sd.dev=sd(diff(na.remove(data$NDUEEGF[1:520])),na.rm=T)/sqrt(22)
set.seed(12356789)
data$NDUEEGF[1:520]=approx(1:n,data$NDUEEGF,1:520)$y+rnorm(520,
sd.dev)

#Plot of the original and the reconstructed data
pdf(file='DataReconstruction.pdf',height=7,width=14)
par(mar=c(2.8,1.5,0.3,0.5),cex=3.3,mgp=c(1.8,0.5,0),cex.main=1)
plot(data$Date[1:520],data$NDUEEGF[1:520],xlab='Year',
type='l',lwd=3,col='red')
lines(data$Date[1:520],original_data,lwd=3)
axis.Date(1,at=seq(as.Date("1999/1/1"),as.Date("2001/1/1"),
"months"),tcl=NA,labels=F)
legend('topleft',legend=c('Original data','Reconstructed data'),
col=1:2,lty=1,lwd=3,x.intersp=0.2,bty='n',seg.len=0.8)
dev.off()
```

```

##Considering Stationarity##
#Plot of index data
pdf(file='Data.pdf',height=6*7,width=2*14.8)
par(mfrow=c(6,2),mar=c(1.5,1.5,2,1),cex=3.3,mgp=c(3,0.5,0),
cex.main=1)
for(i in 2:12){
rmean=numeric(n)
for (j in 1:n){
rmean[j]=mean(data[1:j,i],na.rm=T)
}
plot(data$Date,data[,i],type='l',lwd=3,main=names(data)[i])
lines(data$Date,rmean,lwd=3,col='red')
axis.Date(1,at=seq(as.Date("1999/1/1"),as.Date("2011/1/1"),
"years"),labels=F)
}
dev.off()

#Recursive estimate of SD
pdf(file='SD estimatet recursively 0.95.pdf',height=6*7,
width=2*14.8)
par(mfrow=c(6,2),mar=c(1.5,1.5,2,1),cex=3.3,mgp=c(3,0.5,0),
cex.main=1)
for(i in 1:11){
rsd=numeric(n)
for (j in 1:n){
w=0.05/(1-0.95^j)*0.95^((j-1):0)
rsd[j]=wt.sd(data[1:j,i+1],w)
}
plot(data$Date,rsd,type='l',lwd=3,main=names(data)[i+1])
axis.Date(1,at=seq(as.Date("1999/1/1"),as.Date("2011/1/1"),
"years"),labels=F)
panel=panel.smooth(data$Date,rsd,span=1/5,lwd=9,pch='')
}
dev.off()

#Defining log return and weekly log return
logr_data=apply(apply(data[-1],2,log),2,diff)
wlogr=apply(apply(data[seq(1,dim(data)[1],5),-1],2,log),2,diff)
wdata=data[seq(1,dim(data)[1],5),]

#Log return plot
pdf(file='Log-return.pdf',height=6*7,width=2*14.8)
par(mfrow=c(6,2),mar=c(1.5,1.5,2,1),cex=3.3,mgp=c(3,0.5,0),

```

```

cex.main=1)
for(i in 1:11){
  rmean=numeric(n-1)
  for (j in 1:(n-1)){
    rmean[j]=mean(logr_data[1:j,i],na.rm=T)
  }
  plot(data$Date[-1],logr_data[,i],type='l',lwd=3,
  main=names(data)[i+1])
  lines(data$Date[-1],rmean,lwd=9,col='red')
  axis.Date(1,at=seq(as.Date("1999/1/1"),as.Date("2011/1/1"),
  "years"),labels=F)
}
dev.off()

#Recursive estimate of SD of log returns
pdf(file='SD of log returns estimatet recursively 0.95.pdf',
height=6*7,width=2*14.8)
par(mfrow=c(6,2),mar=c(1.5,1.5,2,1),cex=3.3,mgp=c(3,0.5,0),
cex.main=1)
for(i in 1:11){
  rsd=numeric(n-1)
  for (j in 1:(n-1)){
    w=0.05/(1-0.95^(j-1))*0.95^((j-1):0)
    rsd[j]=wt.sd(logr_data[1:j,i],w)
  }
  plot(data$Date[-1],rsd,type='l',lwd=3,main=names(data)[i+1])
  axis.Date(1,at=seq(as.Date("1999/1/1"),as.Date("2011/1/1"),
  "years"),labels=F)
  panel=panel.smooth(data$Date[-1],rsd,span=1/5,lwd=9,pch='')
}
dev.off()

##The Distribution of Data##
#Estimating mean and standard deviation of log returns
apply(logr_data,2,mean,na.rm=T)
apply(logr_data,2,sd,na.rm=T)

#Testing mean values
apply(logr_data,2,t.test)

#Estimating skewness and kurtosis
skew=apply(logr_data,2,skewness,na.rm=T)
round(skew,2)

```

```
kurt=apply(logr_data,2,kurtosis,na.rm=T)
round(kurt,2)

#Calculating Jarque-Bera test statistic
nLIBOR=length(na.omit(logr_data[,11]))
JB=c(rep((n-1),10),nLIBOR)*(skew^2/6+kurt^2/24)
round(JB)
qchisq(0.95,2)

#Shapiro-Wilk test for normality
apply(logr_data,2,shapiro.test)

#Histograms of indices
pdf(file='Hist data.pdf',height=6*7,width=2*14.8)
par(mfrow=c(6,2),mar=c(1.5,1.5,2,1),cex=3.3,mgp=c(3,0.5,0),
cex.main=1)
for(i in 2:12){
hist(data[,i],main=names(data)[i],freq=F,breaks=50)
xrange=seq(min(hist(data[,i],breaks=50,plot=F)$breaks),
max(hist(data[,i],breaks=50,plot=F)$breaks),length.out=100)
lines(xrange,dnorm(xrange,mean(data[,i],na.rm=T),sd(data[,i],
na.rm=T)),col='red',lwd=9)
}
dev.off()

#Histograms of log returns
pdf(file='Hist log returns.pdf',height=6*7,width=2*14.8)
par(mfrow=c(6,2),mar=c(1.5,1.5,2,1),cex=3.3,mgp=c(3,0.5,0),
cex.main=1)
for(i in 1:11){
hist(logr_data[,i],main=names(data)[i+1],freq=F,breaks=50)
xrange=seq(min(hist(logr_data[,i],breaks=50,plot=F)$breaks)/2,
max(hist(logr_data[,i],breaks=50,plot=F)$breaks)/2,
length.out=100)
lines(xrange,dnorm(xrange,mean(logr_data[,i],na.rm=T),
sd(logr_data[,i],na.rm=T)),col='red',lwd=9)
}
dev.off()

#Histograms of weekly log returns
pdf(file='Hist wlog returns.pdf',height=6*7,width=2*14.8)
par(mfrow=c(6,2),mar=c(1.5,1.5,2,1),cex=3.3,mgp=c(3,0.5,0),
cex.main=1)
for(i in 1:11){
```

```

hist(wlogr[,i],main=names(data)[i+1],freq=F,breaks=50)
xrange=seq(min(hist(wlogr[,i],breaks=50,plot=F)$breaks)/2,
max(hist(wlogr[,i],breaks=50,plot=F)$breaks)/2,
length.out=100)
lines(xrange,dnorm(xrange,mean(wlogr[,i],na.rm=T),
sd(wlogr[,i],na.rm=T)),col='red',lwd=9)
}
dev.off()

#Estimating robust moments for log returns
quantiles=apply(logr_data,2,quantile,probs=
c(0.025,0.25,0.5,0.75,0.975),na.rm=T)
median=quantiles[3,]
dispersion=quantiles[4,]-quantiles[2,]
rob_skew=apply(quantiles,2,function(x)
(x[4]+x[2]-2*x[3])/(x[4]-x[2]))
rob_kurt=apply(quantiles,2,function(x)
(x[5]-x[1])/(x[4]-x[2])-2.91)

#Estimating robust moments for weekly log returns
quantiles=apply(wlogr,2,quantile,probs=
c(0.025,0.25,0.5,0.75,0.975),na.rm=T)
median=quantiles[3,]
dispersion=quantiles[4,]-quantiles[2,]
rob_skew=apply(quantiles,2,function(x)
(x[4]+x[2]-2*x[3])/(x[4]-x[2]))
rob_kurt=apply(quantiles,2,function(x)
(x[5]-x[1])/(x[4]-x[2])-2.91)

##Correlations##
#ACF daily and weekly log returns
pdf(file='ACF.pdf', height=6*7, width=2*14.8)
par(mfrow=c(6,2),mar=c(1.5,1.5,2,1),cex=3.3,mgp=c(3,0.5,0),
cex.main=1)
for(i in 1:11){
daily=acf(logr_data[,i],plot=F,lag.max=10,na.action=na.omit)
weekly=acf(wlogr[,i],plot=F,lag.max=10,na.action=na.omit)
plot(daily$lag-0.08,daily$acf,type='h',lwd=15,col='blue',
xlab='',ylab='',main=names(data)[i+1],ylim=range(weekly$acf,
daily$acf,-0.08,0.08,na.rm=T))
lines(weekly$lag+0.08,weekly$acf,lwd=15,col='red',type='h')
abline(h=qnorm(0.975)/sqrt(daily$n.used),lty='99',col=4,lwd=3)
abline(h=-qnorm(0.975)/sqrt(daily$n.used),lty='99',col=4,lwd=3)
abline(h=qnorm(0.975)/sqrt(weekly$n.used),lty='99',col=2,lwd=3)
}

```



```

abline(h=-qnorm(0.975)/sqrt(weekly$n.used), lty='99', col=2, lwd=3)
abline(h=0)
}
plot(0,0,type='n',xaxt='n',yaxt='n',bty='n',xlab='',ylab='')
legend('left',legend=c('ACF daily\nlog returns','ACF weekly
log returns'),col=c(4,2),lty=1,lwd=15,bty='n',y.intersp=2)
legend('right',legend=c('95% confidence\ninterval daily ACF',
'95% confidence\ninterval weekly ACF'),col=c('blue','red'),
lty='99',bty='n',y.intersp=2,lwd=3)
dev.off()

#ACF1 weekly log returns
pdf(file='ACF1 weekly log returns.pdf',height=6*7,width=2*14.8)
par(mfrow=c(6,2),mar=c(1.5,1.5,2,1),cex=3.3,mgp=c(3,0.5,0),
cex.main=1)
for(i in 1:11){
n=length(na.omit(wlogr[,i]))
acf1=numeric(n-26)
acf2=numeric(n-26)
for(j in 26:(n-1)){
acf1[j-25]=acf(wlogr[(j-25):j,i],na.action=na.pass,plot=F)$acf[1]
acf2[j-25]=acf(wlogr[(j-25):j,i]^2,na.action=na.pass,plot=F)$
acf[1]
}
plot(wdata[-(1:ifelse(i==11,260,27)),1],acf1,type='l',lwd=3,
col='blue',main=names(data)[i+1],xlab='',ylab='',
ylim=range(acf1,acf2,qnorm(0.975)/sqrt(26),na.rm=T))
lines(wdata[-(1:ifelse(i==11,260,27)),1],acf2,lwd=3,col='red')
abline(h=qnorm(0.975)/sqrt(26),lty='99',lwd=3)
abline(h=-qnorm(0.975)/sqrt(26),lty='99',lwd=3)
}
plot(0,0,type='n',xaxt='n',yaxt='n',bty='n',xlab='',ylab='')
legend('left',legend=c('ACF(1) weekly\nlog returns',
'ACF(1) squared\nweekly log returns'),col=c(4,2),lty=1,lwd=3,
bty='n',y.intersp=2)
legend('right',legend=c('95% confidence\ninterval'),
lty='99',bty='n',y.intersp=2,lwd=3)
dev.off()

#CCF plot
library('PerformanceAnalytics')
pnl.ccf=function(x,y){
ccf_down=ccf(wlogr[488:(488+25),x],wlogr[488:(488+25),y],plot=F,
lag.max=3)

```

```

ccf_up=ccf(wlogr[540:(540+25),x],wlogr[540:(540+25),y],plot=F,
lag.max=3)
plot(ccf_down$lag-0.15,ccf_down$acf,type='h',lwd=50,col='red',
ylim=c(-0.4,1),xaxt='n',yaxt='n',xlab='',ylab='')
lines(ccf_up$lag+0.15,ccf_up$acf,type='h',lwd=50,col='blue')
abline(h=qnorm(0.975)/sqrt(26),lwd=18,lty='49')
abline(h=-qnorm(0.975)/sqrt(26),lwd=18,lty='49')
abline(h=0)}

pnl.plot=function(x,y){
plot(wlogr[488:(488+25),x],wlogr[488:(488+25),y],xaxt='n',
yaxt='n',xlab='',ylab='',col='red',pch=16)
points(wlogr[540:(540+25),x],wlogr[540:(540+25),y],col=4,pch=16)
abline(lm(wlogr[488:(488+25),y]~wlogr[488:(488+25),x]),col=2,
lwd=18)
abline(lm(wlogr[540:(540+25),y]~wlogr[540:(540+25),x]),col=4,
lwd=18)}

pdf(file='CCF1.pdf',height=120,width=6*14.8)
par(mfrow=c(11,6),mar=rep(0,4),cex=9.9,las=1,
oma=numeric(4),tck=0.05)
textplot(names(data)[2],cex=1.5,mar=numeric(4))
pnl.ccf(2,1)
axis(3,seq(-3,3,1),labels=F)
axis(3,c(-2,2),mgp=c(0,-1.2,0),tick=F)
pnl.ccf(3,1)
axis(3,seq(-3,3,1),labels=F)
axis(3,c(-2,2),mgp=c(0,-1.2,0),tick=F)
pnl.ccf(4,1)
axis(3,seq(-3,3,1),labels=F)
axis(3,c(-2,2),mgp=c(0,-1.2,0),tick=F)
pnl.ccf(5,1)
axis(3,seq(-3,3,1),labels=F)
axis(3,c(-2,2),mgp=c(0,-1.2,0),tick=F)
pnl.ccf(6,1)
axis(3,seq(-3,3,1),labels=F)
axis(3,c(-2,2),mgp=c(0,-1.2,0),tick=F)
axis(4,c(-0.25,0,0.5),labels=F)
pnl.plot(1,2)
axis(2,seq(-0.2,0.1,0.1),mgp=c(0,-2,0))
textplot(names(data)[3],cex=1.5,mar=numeric(4))
pnl.ccf(3,2)
pnl.ccf(4,2)
pnl.ccf(5,2)

```

```
pnl.ccf(6,2)
axis(4,c(-0.25,0,0.5),labels=F)
pnl.plot(1,3)
axis(2,seq(-0.1,0,0.05),mgp=c(0,-2,0))
pnl.plot(2,3)
textplot(names(data)[4],cex=1.5,mar=numeric(4))
pnl.ccf(4,3)
pnl.ccf(5,3)
pnl.ccf(6,3)
axis(4,c(-0.25,0,0.5),labels=F)
pnl.plot(1,4)
axis(2,seq(-0.15,0.05,0.1),mgp=c(0,-2.5,0))
pnl.plot(2,4)
pnl.plot(3,4)
textplot(names(data)[5],cex=1.5,mar=numeric(4))
pnl.ccf(5,4)
pnl.ccf(6,4)
axis(4,c(-0.25,0,0.5),labels=F)
pnl.plot(1,5)
axis(2,seq(-0.1,0.1,0.1),mgp=c(0,-2,0))
pnl.plot(2,5)
pnl.plot(3,5)
pnl.plot(4,5)
textplot(names(data)[6],cex=1.5,mar=numeric(4))
pnl.ccf(6,5)
axis(4,c(-0.25,0,0.5),labels=F)
pnl.plot(1,6)
axis(2,seq(-0.15,0.05,0.1),mgp=c(0,-2.5,0))
pnl.plot(2,6)
pnl.plot(3,6)
pnl.plot(4,6)
pnl.plot(5,6)
textplot(names(data)[7],cex=1.5,mar=numeric(4))
pnl.plot(1,7)
axis(2,seq(-0.1,0,0.05),mgp=c(0,-2.5,0))
pnl.plot(2,7)
pnl.plot(3,7)
pnl.plot(4,7)
pnl.plot(5,7)
pnl.plot(6,7)
pnl.plot(1,8)
axis(2,seq(-0.05,0.05,0.05),mgp=c(0,-2.5,0))
pnl.plot(2,8)
pnl.plot(3,8)
```

```
pnl.plot(4,8)
pnl.plot(5,8)
pnl.plot(6,8)
pnl.plot(1,9)
axis(2,seq(-0.01,0.02,0.01),mgp=c(0,-2.5,0))
pnl.plot(2,9)
pnl.plot(3,9)
pnl.plot(4,9)
pnl.plot(5,9)
pnl.plot(6,9)
pnl.plot(1,10)
axis(2,seq(-0.03,0.02,0.01),labels=F)
axis(2,seq(-0.03,0.02,0.02),mgp=c(0,-2.5,0),tick=F)
pnl.plot(2,10)
pnl.plot(3,10)
pnl.plot(4,10)
pnl.plot(5,10)
pnl.plot(6,10)
pnl.plot(1,11)
axis(1,seq(-0.2,0.1,0.05),labels=F)
axis(1,c(-0.15,-0.05,0.05),mgp=c(0,-1.2,0),tick=F)
axis(2,seq(-0.1,0.1,0.05),labels=F)
axis(2,c(0,0.1),mgp=c(0,-1.5,0),tick=F)
pnl.plot(2,11)
axis(1,mgp=c(0,-1.2,0))
pnl.plot(3,11)
axis(1,labels=F)
axis(1,c(-0.1,0),mgp=c(0,-1.2,0),tick=F)
pnl.plot(4,11)
axis(1,labels=F)
axis(1,seq(-0.15,0.1,0.1),mgp=c(0,-1.2,0),tick=F)
pnl.plot(5,11)
axis(1,seq(-0.2,0.1,0.1),mgp=c(0,-1.2,0))
pnl.plot(6,11)
axis(1,seq(-0.15,0.1,0.1),mgp=c(0,-1.2,0))
dev.off()

pdf(file='CCF2.pdf',height=120,width=6*14.8)
par(mfrow=c(11,6),mar=rep(0,4),cex=9.9,las=1,
oma=numeric(4),tck=0.05)
pnl.ccf(7,1)
axis(3,seq(-3,3,1),labels=F)
axis(3,c(-2,2),mgp=c(0,-1.2,0),tick=F)
pnl.ccf(8,1)
```

```
axis(3,seq(-3,3,1),labels=F)
axis(3,c(-2,2),mgp=c(0,-1.2,0),tick=F)
pnl.ccf(9,1)
axis(3,seq(-3,3,1),labels=F)
axis(3,c(-2,2),mgp=c(0,-1.2,0),tick=F)
pnl.ccf(10,1)
axis(3,seq(-3,3,1),labels=F)
axis(3,c(-2,2),mgp=c(0,-1.2,0),tick=F)
pnl.ccf(11,1)
axis(3,seq(-3,3,1),labels=F)
axis(3,c(-2,2),mgp=c(0,-1.2,0),tick=F)
axis(4,c(-0.25,0,0.5),labels=c('0.25','','0.50'),mgp=c(0,-2,0))
textplot(' ',mar=numeric(4))
pnl.ccf(7,2)
pnl.ccf(8,2)
pnl.ccf(9,2)
pnl.ccf(10,2)
pnl.ccf(11,2)
axis(4,c(-0.25,0,0.5),labels=c('0.25','','0.50'),mgp=c(0,-2,0))
textplot(' ',mar=numeric(4))
pnl.ccf(7,3)
pnl.ccf(8,3)
pnl.ccf(9,3)
pnl.ccf(10,3)
pnl.ccf(11,3)
axis(4,c(-0.25,0,0.5),labels=c('0.25','','0.50'),mgp=c(0,-2,0))
textplot(' ',mar=numeric(4))
pnl.ccf(7,4)
pnl.ccf(8,4)
pnl.ccf(9,4)
pnl.ccf(10,4)
pnl.ccf(11,4)
axis(4,c(-0.25,0,0.5),labels=c('0.25','','0.50'),mgp=c(0,-2,0))
textplot(' ',mar=numeric(4))
pnl.ccf(7,5)
pnl.ccf(8,5)
pnl.ccf(9,5)
pnl.ccf(10,5)
pnl.ccf(11,5)
axis(4,c(-0.25,0,0.5),labels=c('0.25','','0.50'),mgp=c(0,-2,0))
textplot(' ',mar=numeric(4))
pnl.ccf(7,6)
pnl.ccf(8,6)
pnl.ccf(9,6)
```

```
pnl.ccf(10,6)
pnl.ccf(11,6)
axis(4,c(-0.25,0,0.5),labels=c('0.25','','0.50'),mgp=c(0,-2,0))
textplot(' ',mar=numeric(4))
textplot(names(data)[8],cex=1.5,mar=numeric(4))
pnl.ccf(8,7)
pnl.ccf(9,7)
pnl.ccf(10,7)
pnl.ccf(11,7)
axis(4,c(-0.25,0,0.5),labels=c('0.25','','0.50'),mgp=c(0,-2,0))
textplot(' ',mar=numeric(4))
pnl.plot(7,8)
axis(2,seq(-0.05,0.05,0.05),labels=F)
textplot(names(data)[9],cex=1.5,mar=numeric(4))
pnl.ccf(9,8)
pnl.ccf(10,8)
pnl.ccf(11,8)
axis(4,c(-0.25,0,0.5),labels=c('0.25','','0.50'),mgp=c(0,-2,0))
textplot(' ',mar=numeric(4))
pnl.plot(7,9)
axis(2,seq(-0.01,0.02,0.01),labels=F)
pnl.plot(8,9)
textplot(names(data)[10],cex=1.5,mar=numeric(4))
pnl.ccf(10,9)
pnl.ccf(11,9)
axis(4,c(-0.25,0,0.5),labels=c('0.25','','0.50'),mgp=c(0,-2,0))
textplot(' ',mar=numeric(4))
pnl.plot(7,10)
axis(2,seq(-0.03,0.02,0.01),labels=F)
pnl.plot(8,10)
pnl.plot(9,10)
textplot(names(data)[11],cex=1.5,mar=numeric(4))
pnl.ccf(11,10)
axis(4,c(-0.25,0,0.5),labels=c('0.25','','0.50'),mgp=c(0,-2,0))
textplot(' ',mar=numeric(4))
pnl.plot(7,11)
axis(1,seq(-0.1,0,0.025),labels=F)
axis(1,c(-0.05,0),labels=c('-0.05','0'),mgp=c(0,-1.2,0),tick=F)
axis(2,seq(-0.1,0.1,0.05),labels=F)
pnl.plot(8,11)
axis(1,labels=F)
axis(1,seq(-0.1,0.05,0.05),mgp=c(0,-1.2,0),tick=F)
pnl.plot(9,11)
axis(1,labels=F)
```

```

axis(1,c(-0.005,0.005),labels=c('-0.005','0.005'),mgp=c(0,-1.2,0))
pnl.plot(10,11)
axis(1,labels=F)
axis(1,seq(-0.02,0.02,0.02),labels=c('-0.02','0','0.02'),
mgp=c(0,-1.2,0),tick=F)
textplot(names(data)[12],cex=1.5,mar=numeric(4))
dev.off()
detach("package:PerformanceAnalytics")

##Principal Component Analysis##
pca=prcomp(na.omit(wlogr),scale=T,center=T)
round(summary(pca)$importance,2)

#Eigenvectors
round(pca$rotation[,1:7],2)

#Scree plot with broken stick criterion
brocken_stick=sapply(1:11,function(x)sum(1/(x:11)))
pdf(file='Scree plot all data.pdf',height=7,width=14)
par(mar=c(2.8,2.8,0.1,0.5),cex=3.3,mgp=c(1.8,0.5,0),cex.main=1)
plot(1:11,pca$sdev^2,type='o',lwd=3,xlab='Component number',
ylab='Eigenvalue',pch=16)
lines(1:11,brocken_stick,type='o',pch=17,lwd=3,col='red')
legend('topright',legend=c('Weekly log return data','Random data'),
col=1:2,lty=1,lwd=3,pch=c(16,17),x.intersp=0.5,bty='n')
dev.off()

#Communalities
apply(pca$rotation[,1:4],1,function(x)sum(x^2*pca$sdev[1:4]^2))
-log(prod(pca$sdev[3:11]^2)/(sum(pca$sdev[3:11]^2)/9)^9)

```

B.2 R-code for Chapter 4

```

#####
## Filename: Modelling.R                               Date: 20-01-2012 ##
## Author: Peter Nystrup s082634                      ##
## Description: R-code for Chapter 4                   ##
#####

##Choosing a Regime Approach##

```

```

#Reading in CLI data
CLI<-read.csv2('OECD.csv',header=T)
CLI$Time<-seq(as.Date('1998-9-1'),as.Date('2011-11-1'),'months')

#Plotting CLI and IIP
pdf(file='CLI plot.pdf',height=7,width=14)
par(mar=c(2.8,1.5,0.3,0.5),cex=3.3,mgp=c(1.8,0.5,0),cex.main=1)
matplot(CLI[,1],CLI[,-1],xaxt='n',type='l',lwd=9,lty=1,
col=c('blue','red'),xlab='Year',ylab='')
axis.Date(1,at=seq(as.Date("1999/1/1"),as.Date("2012/1/1"),
"years"))
abline(v=as.Date("1998/11/1"),lwd=3)
abline(v=as.Date("2000/3/1"),lwd=3)
abline(v=as.Date("2001/10/1"),lwd=3)
abline(v=as.Date("2007/6/1"),lwd=3)
abline(v=as.Date("2009/3/1"),lwd=3)
abline(v=as.Date("2011/3/1"),lwd=3)
legend('bottomleft',c('Normalised CLI','Normalised IIP'),lty=1,
bg='white',lwd=9,col=c('blue','red'),x.intersp=0.2,
seg.len=0.8,y.intersp=0.8)
dev.off()

#Index plot
indeks_data=100*apply(data[,-1],2,function(x)x/na.omit(x)[1])
pdf(file='Indeksplot samlet.pdf',height=8,width=14)
par(mar=c(2.8,1.5,0.3,0.5),cex=1.65,mgp=c(1.8,0.5,0),cex.main=1)
matplot(data$Date,indeks_data,xaxt='n',type='l',lwd=3,lty=1,
col=c(1:8,'brown','orange','lightpink'),xlab='Year',ylab='')
axis.Date(1,at=seq(as.Date("1999/1/1"),as.Date("2012/1/1"),
"years"))
abline(v=as.Date("1998/11/1"),lwd=2)
abline(v=as.Date("2000/3/1"),lwd=2)
abline(v=as.Date("2001/10/1"),lwd=2)
abline(v=as.Date("2007/6/1"),lwd=2)
abline(v=as.Date("2009/3/1"),lwd=2)
abline(v=as.Date("2011/3/1"),lwd=2)
text(as.Date("1998/11/1"),0,pos=4,offset=0,'M10')
text(as.Date("2000/3/1"),0,pos=4,offset=0,'M2')
text(as.Date("2001/10/1"),0,pos=4,offset=0,'M9')
text(as.Date("2007/6/1"),0,pos=4,offset=0,'M5')
text(as.Date("2009/3/1"),0,pos=4,offset=0,'M2')
text(as.Date("2011/3/1"),0,pos=4,offset=0,'M2')
legend('topleft',names(data)[-1],lty=1,bg='white',lwd=4,
col=c(1:8,'brown','orange','lightpink'),y.intersp=1,ncol=2)

```



```
dev.off()

##Estimating Principal Components##
#Modelling upturn
data_up=wlogr[c(1:(which(wdata[,1]=="2000-03-03")-1),
which(wdata[,1]=="2001-09-28"):
(which(wdata[,1]=="2007-06-01")-1),
which(wdata[,1]=="2009-02-27"):
(which(wdata[,1]=="2011-02-25")-1)), -11]

#Eigenvalues
pca_up=prcomp(data_up,scale=T,center=T)
round(summary(pca_up)$importance,2)

#Eigenvectors
round(pca_up$rotation[,1:7],2)

#Modelling downturn
data_down=wlogr[c(which(wdata[,1]=="2000-03-03"):
(which(wdata[,1]=="2001-09-28")-1),which(wdata[,1]=="
"2007-06-01"):(which(wdata[,1]=="2009-02-27")-1),
which(wdata[,1]=="2011-02-25"):dim(wlogr)[1]), -11]

#Eigenvalues
pca_down=prcomp(data_down,scale=T,center=T)
round(summary(pca_down)$importance,2)

#Eigenvectors
round(pca_down$rotation[,1:7],2)

#Communalities
round(apply(pca_up$rotation[,1:3],1,function(x)sum(x^2*
pca_up$sdev[1:3]^2)),2)

round(apply(pca_down$rotation[,1:3],1,function(x)sum(x^2*
pca_down$sdev[1:3]^2)),2)

#Scree plot with broken stick criterion
brocken_stick=sapply(1:10,function(x)sum(1/(x:10)))
pdf(file='Scree plot model.pdf',height=7,width=14)
par(mar=c(2.8,2.8,0.1,0.5),cex=3.3,mgp=c(1.8,0.5,0),cex.main=1)
plot(1:10,pca_down$sdev^2,type='o',xlab='Component number',
ylab='Eigenvalue',pch=16,lwd=3)
lines(1:10,pca_up$sdev^2,type='o',lwd=3,col='blue',pch=15)
```

```

lines(1:10,brocken_stick,type='o',pch=17,lwd=3,col='red')
legend('topright',legend=c('Upturn','Downturn','Random data'),
col=c('blue','black','red'),lty=1,lwd=3,pch=15:17,
x.intersp=0.5,bty='n')
dev.off()

##Estimating Model Parameters##
models_up=list()
models_down=list()
for (i in 1:3){
models_up[[i]]=garch(pca_up$x[,i],order=c(0,1))$coef
models_down[[i]]=garch(pca_down$x[,i],order=c(0,1))$coef
}

##Assessing Model Performance##
#Function that rotates data back to index space from PC space
restore=function(rotated_data,pca,n){
r=na.omit(rotated_data%*%t(pca$rotation[,1:n]))
r=t(apply(r,1,function(x)x*pca$scale))
r=t(apply(r,1,function(x)x+pca$center))
}

#Function that simulates GARCH model and restores data
simulate=function(coefs,nsteps,pca){
simulated_pc=matrix(numeric(nsteps*3),ncol=3)
for (j in 1:3){
simulated_pc[,j]=garch.sim(coefs[[j]],n=nsteps,ntrans=100)
}
restore(simulated_pc,pca,3)
}

#Simulating data
Nweeks=dim(wlogr)[1]
Nsim=10000
simulated_logr=array(numeric(Nweeks*10*Nsim),c(Nweeks,10,Nsim))

for (i in 1:Nsim){
simulated_logr[1:61,,i]=simulate(models_up,61,pca_up)
simulated_logr[62:143,,i]=simulate(models_down,82,pca_down)
simulated_logr[144:439,,i]=simulate(models_up,296,pca_up)
simulated_logr[440:530,,i]=simulate(models_down,91,pca_down)
simulated_logr[531:634,,i]=simulate(models_up,104,pca_up)
simulated_logr[635:658,,i]=simulate(models_down,24,pca_down)
}

```

```

simulated_data=exp(apply(simulated_logr,c(2,3),cumsum))

pdf(file='Testing model.pdf',height=5*6.78,width=2*14.8)
par(mfrow=c(5,2),mar=c(1.5,1.5,2,1),cex=3.3,mgp=c(3,0.5,0),
cex.main=1)
for(i in 1:10){
quantiles=t(apply(simulated_data[,i]*wdata[1,i+1],1,
quantile,probs=c(0.05,0.25,0.5,0.75,0.95)))
matplot(wdata[-1,1],quantiles,type='l',lty=c(3,2,1,2,3),xaxt='n',
lwd=c(3,3,9,3,3),col='black',xlab='',ylab='',
main=names(data)[i+1])
matlines(wdata[-1,1],wdata[-1,i+1],lwd=9,col='red')
axis.Date(1,at=seq(as.Date("1999/1/1"),as.Date("2012/1/1"),
"years"))
}
dev.off()

```

B.3 R-code for Chapter 5

```

#####
## Filename: ScenarioGeneration.R           Date: 20-01-2012 ##
## Author: Peter Nystrup s082634         ##
## Description: R-code for Chapter 5      ##
#####

##The Choise of Regime##
#Estimating historical phase lengths
up=c(21,13,46,27,54,21,42,17,19,19,86,23)
down=c(15,34,18,26,19,36,27,38,15,14,16,14)
mean(up)
median(up)
sd(up)
mean(down)
median(down)
sd(down)

##Generating Scenarios##
#Initialisation
Nweeks=5*52
Nscenario=10000

```

```

simulated_logr=array(numeric(Nweeks*10*Nscenario),
c(Nweeks,10,Nscenario))
lambda_down=1/98
lambda_up=1/140

for (i in 1:Nscenario){
N=0
up_or_dn=0
while (N<Nweeks){
n=round(rexp(1,rate=ifelse(up_or_dn==0,lambda_down,lambda_up)))
n=min(n,Nweeks-N)
if(n!=0){
if (up_or_dn==0){
simulated_logr[(N+1):(N+n),,i]=simulate(models_down,n,pca_down)
}else{
simulated_logr[(N+1):(N+n),,i]=simulate(models_up,n,pca_up)
}}
N=N+n
up_or_dn=ifelse(up_or_dn==0,1,0)
}}

simulated_data=exp(apply(simulated_logr,c(2,3),cumsum))

##Assessing Scenarios##
#Summarizing simulated data
Q2=apply(apply(simulated_data[seq(52,260,52),,],c(1,2),quantile,
probs=0.5),1,function(x)x*wdata[dim(wdata)[1],2:11])
apply(Q2

apply(apply(simulated_data[seq(52,260,52),,],c(1,2),
function(x)(quantile(x,probs=0.75)-quantile(x,probs=0.25))),1,
function(x)x*wdata[dim(wdata)[1],2:11])

#Simulation dates
dates=seq(as.numeric(wdata[length(wdata[,1]),1])+7,
as.numeric(wdata[length(wdata[,1]),1])+52*7*5,7)
dates=sapply(dates,as.Date,origin='1970-1-1',simplify=F)

#Finding maximum drawdown for each index
MDD_location=numeric(10)
for (i in 1:10){
MDD=0
for (j in 1:Nscenario){
peak=0

```

```

for (k in 1:Nweeks){
  if (simulated_data[k,i,j] > peak){
    peak=simulated_data[k,i,j]
  } else {
    DD=(peak-simulated_data[k,i,j])/peak
    if (DD > MDD){
      MDD=DD
      MDD_location[i]=j
    }}}}

#Visualising scenarios
pdf(file='Scenarios.pdf',height=5*7,width=2*14.8)
par(mfrow=c(5,2),mar=c(1.5,1.5,2,1),cex=3.3,mgp=c(3,0.5,0),
    cex.main=1)
for(i in 1:10){
  quantiles=t(apply(simulated_data[,i,]*wdata[dim(wdata)[1],i+1],
    1,quantile,probs=c(0.05,0.25,0.5,0.75,0.95)))
  stratification=simulated_data[,i,sapply(seq(1,Nscenario,
    Nscenario/10),function(x)which(rank(
    simulated_data[Nweeks,i,]==x)))*wdata[dim(wdata)[1],i+1]
  MDD=simulated_data[,i,MDD_location[i]]*
  wdata[length(wdata[,1]),i+1]
  matplot(dates,quantiles,lty=c(3,2,1,2,3),lwd=3,xaxt='n',
    col='black',type='l',xlab='',ylab='',main=names(data)[i+1],
    ylim=range(quantiles,stratification,MDD))
  matlines(dates,stratification,lty=1,col='darkgrey',lwd=3)
  lines(dates,MDD,col='red',lwd=3)
  axis.Date(1,at=seq(as.Date("2012/1/1"),as.Date("2016/1/1"),
    "years"))
}
dev.off()

#Distribution of maximum drawdowns
MDD=array(numeric(Nweeks*10*Nscenario),c(Nweeks,10,Nscenario))
for (i in 1:10){
  for (j in 1:Nscenario){
    peak=0
    for (k in 1:Nweeks){
      if (simulated_data[k,i,j] > peak){
        peak=simulated_data[k,i,j]
      } else {
        DD=(peak-simulated_data[k,i,j])/peak
        if (DD > MDD[k,i,j]){
          MDD[k:Nweeks,i,j]=DD
        }
      }
    }
  }
}

```

```
}}}}}
```

```
pdf(file='Maximum drawdown.pdf',height=5*7,width=2*14.8)
par(mfrow=c(5,2),mar=c(1.5,1.5,2,1),cex=3.3,mgp=c(3,0.5,0),
cex.main=1)
for(i in 1:10){
matplot(dates,apply(MDD[,i,],1,max),type='l',lwd=9,lty=1,
col='red',xaxt='n',xlab='',ylab='',main=names(data)[i+1])
matlines(dates,t(apply(MDD[,i,],1,quantile,
probs=c(0.05,0.25,0.5,0.75,0.95))),lty=c(3,2,1,2,3),
lwd=9,col='black')
axis.Date(1,at=seq(as.Date("2012/1/1"),as.Date("2016/1/1"),
"years"))
}
dev.off()
```

```
pdf(file='Distributions of final values.pdf',height=5*7,
width=2*14.8)
par(mfrow=c(5,2),mar=c(1.5,1.5,2,1),cex=3.3,mgp=c(3,0.5,0),
cex.main=1)
for(i in 1:10){
xlim=quantile(simulated_data[Nweeks,i,]*100,probs=c(0,0.99))
breaks=c(seq(xlim[1],xlim[2],length.out=50),
max(simulated_data[Nweeks,i,]*100))
hist(simulated_data[Nweeks,i,]*100,freq=F,
main=names(data)[i+1],breaks=breaks,xlim=xlim)
abline(v=quantile(simulated_data[Nweeks,i,]*100,probs=c(0.05,
0.25,0.5,0.75,0.95)),col='blue',lwd=9,lty=c(3,2,1,2,3))
}
dev.off()
```

```
##Bootstrapping##
bootstrapped_logr=array(numeric(Nweeks*10*Nscenario),
c(Nweeks,10,Nscenario))
for (i in 1:Nscenario){
tid=numeric(Nweeks)
tid[1]=sample(1:658,1,prob=dnorm(1:658,658,5))
for (j in 2:Nweeks){
tid[j]=sample(1:658,1,prob=dnorm(1:658,tid[j-1],5))
}
bootstrapped_logr[,i]=wlogr[tid,-11]
}
bootstrapped_data=exp(apply(bootstrapped_logr,c(2,3),cumsum))
```

```
##Bootstrapping Performance##
#Summarizing bootstrapped data
Q2=apply(apply(bootstrapped_data[seq(52,Nweeks,52),,],c(1,2),
quantile,probs=0.5),1,function(x)x*wdata[dim(wdata)[1],2:11])

apply(apply(bootstrapped_data[seq(52,Nweeks,52),,],c(1,2),
function(x)(quantile(x,probs=0.75)-quantile(x,probs=0.25))),1,
function(x)x*wdata[dim(wdata)[1],2:11])

#Finding maximum drawdown
MDD_location=numeric(10)
for (i in 1:10){
MDD=0
for (j in 1:Nscenario){
peak=0
for (k in 1:Nweeks){
if (bootstrapped_data[k,i,j] > peak){
peak=bootstrapped_data[k,i,j]
} else {
DD=(peak-bootstrapped_data[k,i,j])/peak
if (DD > MDD){
MDD=DD
MDD_location[i]=j
}}}}
}

#Visualising bootstrapped scenarios
pdf(file='Bootstrapped scenarios.pdf',height=5*7,width=2*14.8)
par(mfrow=c(5,2),mar=c(1.5,1.5,2,1),cex=3.3,mgp=c(3,0.5,0),
cex.main=1)
for(i in 1:10){
quantiles=t(apply(bootstrapped_data[,i,]*wdata[dim(wdata)[1],
i+1],1,quantile,probs=c(0.05,0.25,0.5,0.75,0.95)))
stratification=bootstrapped_data[,i,sapply(seq(1,Nscenario,
Nscenario/10),function(x)which(rank(
bootstrapped_data[Nweeks,i,])==x)))*wdata[dim(wdata)[1],i+1]
MDD=bootstrapped_data[,i,MDD_location[i]]*
wdata[length(wdata[,1]),i+1]
matplot(dates,quantiles,lty=c(3,2,1,2,3),lwd=3,xaxt='n',
col='black',type='l',xlab='',ylab='',main=names(data)[i+1],
ylim=range(quantiles,stratification,MDD))
matlines(dates,stratification,lty=1,col='darkgrey',lwd=3)
lines(dates,MDD,col='red',lwd=3)
axis.Date(1,at=seq(as.Date("2012/1/1"),as.Date("2016/1/1"),
"years"))
```

```

}
dev.off()

#Distribution of maximum drawdowns
MDD=array(numeric(Nweeks*10*Nscenario),c(Nweeks,10,Nscenario))
for (i in 1:10){
for (j in 1:Nscenario){
peak=0
for (k in 1:Nweeks){
if (bootstrapped_data[k,i,j] > peak){
peak=bootstrapped_data[k,i,j]
} else {
DD=(peak-bootstrapped_data[k,i,j])/peak
if (DD > MDD[k,i,j]){
MDD[k:Nweeks,i,j]=DD
}}}}
}

pdf(file='Maximum drawdown boot.pdf',height=5*7,width=2*14.8)
par(mfrow=c(5,2),mar=c(1.5,1.5,2,1),cex=3.3,mgp=c(3,0.5,0),
cex.main=1)
for(i in 1:10){
matplot(dates,apply(MDD[,i,],1,max),type='l',lwd=9,lty=1,
col='red',xaxt='n',xlab='',ylab='',main=names(data)[i+1])
matlines(dates,t(apply(MDD[,i,],1,quantile,
probs=c(0.05,0.25,0.5,0.75,0.95))),lty=c(3,2,1,2,3),
lwd=9,col='black')
axis.Date(1,at=seq(as.Date("2012/1/1"),as.Date("2016/1/1"),
"years"))
}
dev.off()

pdf(file='Distributions of final values boot.pdf',height=5*7,
width=2*14.8)
par(mfrow=c(5,2),mar=c(1.5,1.5,2,1),cex=3.3,mgp=c(3,0.5,0),
cex.main=1)
for(i in 1:10){
xlim=quantile(bootstrapped_data[Nweeks,i,]*100,probs=c(0,0.99))
breaks=c(seq(xlim[1],xlim[2],length.out=50),
max(bootstrapped_data[Nweeks,i,]*100))
hist(bootstrapped_data[Nweeks,i,]*100,freq=F,
main=names(data)[i+1],breaks=breaks,xlim=xlim)
abline(v=quantile(bootstrapped_data[Nweeks,i,]*100,probs=
c(0.05,0.25,0.5,0.75,0.95)),col='blue',lwd=9,lty=c(3,2,1,2,3))
}

```



```
dev.off()

pdf(file='QQ final values.pdf',height=5*7,width=2*14.8)
par(mfrow=c(5,2),mar=c(1.5,1.5,2,1),cex=3.3,mgp=c(3,0.5,0),
    cex.main=1)
for(i in 1:10){
  quantiles_sim=quantile(simulated_data[Nweeks,i,]*100,
    probs=(1:990)/1000)
  quantiles_boot=quantile(bootstrapped_data[Nweeks,i,]*100,
    probs=(1:990)/1000)
  plot(quantiles_boot,quantiles_sim,main=names(data)[i+1],
    xlab='',ylab='')
  abline(lm(quantiles_sim[c(250,750)]~quantiles_boot[c(250,750)]),
    lwd=9,col='red')
}
dev.off()
```