Delft University of Technology
Faculty Electrical Engineering, Mathematics & Computer Science
Delft Institute of Applied Mathematics

Financial stock market modeling and the COVID-19 crisis

(Dutch title: Het modelleren van de aandelenmarkt en de coronacrisis)

Report on behalf of the
Delft Institute of Applied Mathematics
submitted in connection with

the degree of

BACHELOR OF SCIENCE
in
APPLIED MATHEMATICS

by

ISABEL KLEIN

Delft, The Netherlands
July 2021

BSc report APPLIED MATHEMATICS

"Financial stock market modeling and the COVID-19 crisis"
(Dutch title: "Het modelleren van de aandelenmarkt en de coronacrisis")

ISABEL KLEIN

Delft University of Technology

Supervisor

Dr. A.F.F. Derumigny

Other committee members

Dr. B. van den Dries

Dr. A. Papapantoleon

July 2021

Delft

Abstract

The COVID-19 crisis heavily affected financial stock markets. In March 2020 stock prices dropped immensely and markets became extremely volatile. In this report we model three European stock markets before and during the COVID-19 crisis to determine whether the dynamics of financial markets changed structurally compared to previous periods. Stock markets play a big role in our economy and can cause economic disruption when they crash. Therefore, it can be very useful to understand the dynamics of the stock market. Econometricians are nowadays often asked to model the non-constant volatility (conditional heteroscedasticity) of financial time series. This report uses Generalised AutoRegressive Conditionally Heteroscedastic (GARCH) models that are known for their precise modeling of conditional heteroscedasticity. They are also known for their ability to capture the key stylised facts, the common empirical properties applicable to all types of stock markets. This report includes general definitions and characteristics of the GARCH(p, q) process and covers topics such as autocorrelation of returns, kurtosis, leptokurticity and volatility clustering. For estimating the GARCH models this report uses the statistical program R which estimates the parameters by the Quasi-Maximum Likelihood Estimation (QMLE) method. We modeled different periods before and during COVID-19 and compared the estimated parameters with the corresponding 95% confidence intervals. To test the accuracy of the estimations we performed parametric bootstrapping. Throughout the report, models for the Dutch Amsterdam Exchange (AEX) index, the French Cotation Assistée Continu (CAC 40) index and the German Deutsche Aktien (DAX) index are analysed and compared. It seems that European markets may experience the impact of stock market crashes differently. The DAX index shows significant changes in the dynamics of the stock market due to the COVID-19 crisis whereas the AEX and CAC 40 index do not.

Layman Abstract

The COVID-19 crisis heavily affected financial stock markets. In March 2020 stock prices dropped immensely and markets became extremely volatile. In this report we model three European stock markets before and during the COVID-19 crisis to determine whether the dynamics of financial markets changed structurally compared to previous periods. Stock markets play a big role in our economy and can cause economic disruption when they crash. Therefore, it can be very useful to understand the dynamics of the stock market. For example, statisticians use stock market modeling for risk management and predicting future values. Empirical studies on financial stock markets have shown that several common properties are applicable to all types of stock markets, independent of location, time and details of the market structure. These properties are defined as the stylised facts of financial markets. For modeling the stock markets this report uses a type of model, known as a Generalised AutoRegressive Conditionally Heteroscedastic (GARCH) model, that captures the key stylised facts. Throughout the report models for the Dutch Amsterdam Exchange (AEX) index, the French Cotation Assistée Continu (CAC 40) index and the German Deutsche Aktien (DAX) index are analysed and compared. It seems that European markets may experience the impact of stock market crashes differently. The DAX index shows significant changes in the dynamics of the stock market due to the COVID-19 crisis whereas the AEX and CAC 40 index do not.

Contents

1 Introduction

The COVID-19 pandemic has had a tremendous impact on the entire world. As a consequence of the rapid spread of the virus, countries initiated lockdowns which caused economic damage. Many businesses went bankrupt and people lost their jobs, but also some sectors such as food delivery thrived. It was a scary and uncertain time that also heavily affected financial stock markets. In March 2020 stock prices dropped immensely and markets became extremely volatile. On February 28, 2020, stock markets announced their greatest single-week decline since the Global Financial Crisis in 2008 [25]. By March volatility had gone through the roof and markets showed an overall decline of 30% [13].

The 2020 stock market crash was clearly part of the economic damage caused by the pandemic and shows that the stock markets play a big role in our economy. While we can demonstrably measure the decline in stock prices during the crisis, a more challenging question is whether COVID-19 structurally changed the dynamics of the stock market. This report sets out to answer this important question.

To answer this question we model three European stock markets on different periods before and during the COVID-19 crisis. We estimate models fitted on data from the Amsterdam Exchange (AEX), the Cotation Assistée Continu (CAC 40) and the Deutsche Aktien (DAX) indices.

Chapter 2 presents a general understanding of financial stock markets and explains common topics including returns, volatility and Value at Risk. Additionally, we introduce several empirical properties applicable to all types of stock markets, defined as stylised facts. Key stylised facts such as *volatility clustering* and *fat-tailed distribution* are discussed in Section 2.2.

In the next chapter we introduce Generalised AutoRegressive Conditionally Heteroscedastic (GARCH) models, one of the most popular families of models for financial time series. These models are known for their precise modeling of conditional heteroscedasticity, which is the way volatility alternates between high and low periods. These models are also able to capture the main stylised facts. We look specifically at univariate GARCH models and present several definitions and properties of the GARCH(p, q) process. We discuss the simple case of the GARCH(1,1) model since this GARCH(p, q) specification gives the best fit for our dataset. We use the statistical program R to estimate the models. The parameters are estimated by the Quasi-Maximum Likelihood Estimation (QMLE) method, which is discussed in Section 3.4.2.

Lastly, Chapter 4 sets out the methodology of our research. We start with the Dutch AEX index. We split the data in different periods before and during the COVID-19 crisis and perform several estimations in R using the 'fgarch' package. GARCH processes are designed to model high and low volatility, but we are curious to see how it models the period during COVID-19. We compare the different models by presenting error bar plots of the estimated parameters and the corresponding 95% confidence intervals. In Section 4.4 we introduce the concept of parametric bootstrapping, a method to test the accuracy of the estimations. After the analysis of the AEX index, this same research is presented for the French CAC 40 and the German DAX indices.

2 Financial Stock Markets

To be able to answer our research question it is useful to have a general understanding of financial stock markets. The first section of this chapter describes how stocks are bought and sold in an order-driven marketplace. Common topics such as return values, price movement, volatility and Value at Risk (VaR) are explained and illustrated using data from the Amsterdam Exchange (AEX) index. Figures showing the decline of stock prices due to crises, such as the Global Financial Crisis in 2008 and the current COVID-19 crisis, will be presented. In these crisis periods we have felt how stock markets play a big role in our economy. Therefore it is very useful to understand the dynamics of the stock market. Section 2.2 presents several characteristics of financial time series that are necessary for this understanding. Empirical studies have shown that there exist common properties applicable to all types of stock markets that are defined as stylised facts. Eight key stylised facts will be thoroughly discussed in Section 2.2.2.

2.1 Concepts of Financial Stock Markets

The main object of study in finance are financial assets. Such financial assets are called securities or financial instruments and the most common ones are stocks (equity securities), bonds (debt securities) and exchange-traded funds (ETFs). These assets can be exchanged on organised markets which we call financial markets. Financial markets play an important role in free market economies and can cause great economic disruption when they crash. The financial stock market, where stocks are exchanged, belongs to this larger family of financial markets.

A stock represents a fraction of the ownership of a company and investors put in orders in the stock market to either buy or sell stocks. The market places provide an efficient environment and guarantee market transparency and liquidity. Liquidity defines the ease of converting a security into cash. Throughout the day when the marketplace is open, buyers and sellers send in their orders and the market directly matches as many as possible to execute trades, without intermediaries [8]. We therefore speak of an order-driven market. One of the functions of the stock market is determining the price of the stocks. All order-driven markets rank and match the incoming orders but have their own rules for determining the price at which contracts are executed [8]. The price determination process is based on finding the intersection of supply and demand.

An important Dutch stock market is the Amsterdam Exchange (AEX), home to some of the biggest companies in the Netherlands. There are many other markets including the well-known New York Stock Exchange and NASDAQ in the United States which are two of the largest in the world. The majority of this research project uses data from the *AEX index*, the aggregated data from all the companies in the AEX. Figure 2.1 shows a chart of the AEX index values from January 2007 until May 2021. The constant fluctuation of the price values is visible with peaks and some moments with deep lows. In the graph two large draw-downs are clearly visible, in 2008 and in 2020. These correspond to the Global Financial Crisis of 2008 and the COVID-19 crisis in 2020.

Figure 2.1: Time series of the AEX index values January 2007 - May 2021.

2.1.1 Returns

Most of the time in financial studies return values are used instead of price values. One reason for this is that the return value is independent of the unit of currency. Secondly, return series have more attractive statistical properties which make them easier to model [26]. These properties will be discussed later in Section 2.2. Now we will introduce two different types of returns, the *simple return* and *continuous compounded return*.

We consider observations to be at a certain moment in time and define the time-step between observations as Δt. The AEX index values are daily observations ($\Delta t = 1$ day) but other possibilities are for example per second, minute, hour or month. For a given time t, we denote P_t as the price of a given asset at time t. The simple return r_t is then defined as the relative price variation, and calculated as follows:

$$r_t = \frac{P_t - P_{t-1}}{P_{t-1}}. \tag{2.1}$$

Figure 2.2 shows the simple returns of the AEX index daily stock prices from Figure 2.1. On the other hand, the continuous compounded return R_t, also known as the log return (or simply called the return) is defined as:

$$R_t = \ln\Big(\frac{P_t}{P_{t-1}}\Big). \tag{2.2}$$

For small values of x it is known that $\ln(1+x) = x + O(x^2)$, where $O(x^2)$ describes the error term of this approximation. Therefore since return values are small and $R_t = \ln(1 + r_t)$, we see that the two return series are close to each other. The log return, however, is slightly preferred since its statistical properties are easier dealt with. Moreover, the logarithmic function makes it easier to change the time-scale, for examples when wanting to calculate the weekly return value from daily observations.

The large negative and positive values in Figure 2.2, called extreme returns, are of interest for investors observing the price movements. In Figure 2.2 the seven largest absolute returns are

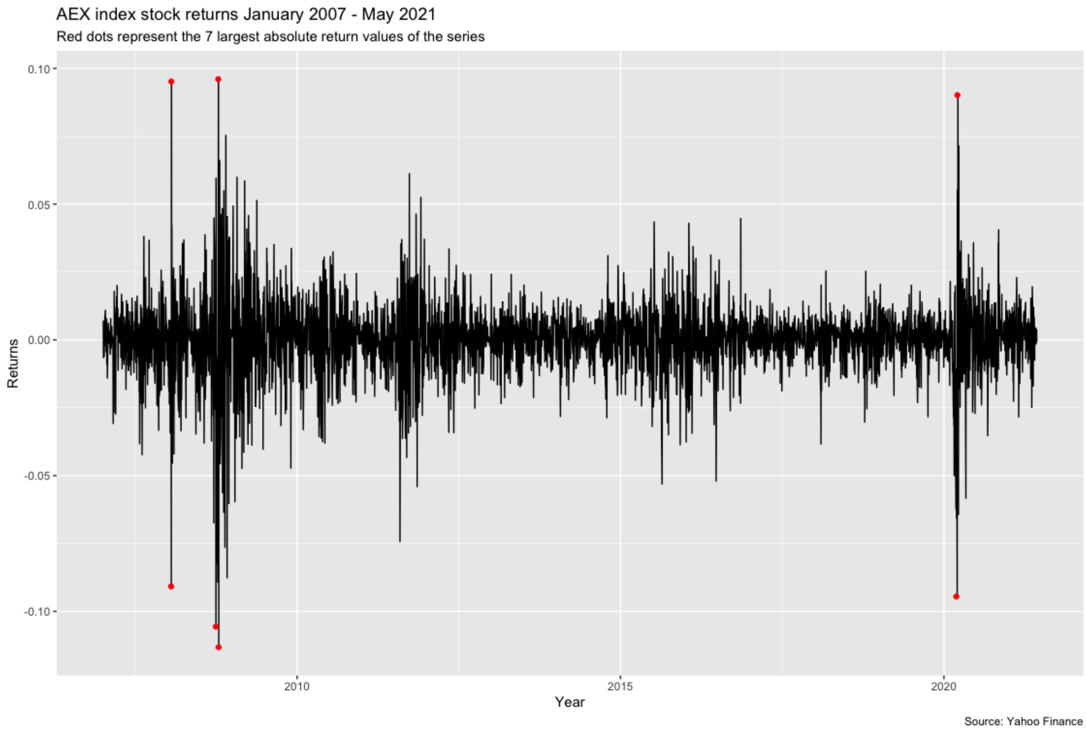

Figure 2.2: The AEX index stock returns January 2007 - May 2021. The red dots mark a few of the largest absolute returns values of the series that occurred during the Global Financial Crisis in 2008 and the current COVID-19 crisis that started in 2020.

marked by a red dot. The largest price drops occurred on 01/21/2008 and 03/11/2020, both during crisis periods. Looking at the graph we see that most of the returns deviate around the mean value zero. Finally, we define the volatility of market prices by the standard deviation of the returns. This concept will be further discussed in the next section.

2.1.2 Volatility

Investors like to have stability because they do not want the prices to drop or fluctuate very much, and therefore we study the volatility. To find a model that is coherent with the empirical properties of a stock market, it must be able to reflect its volatility. Volatility is measured by the standard deviation of the returns, or its conditional equivalent given the amount of information known at a certain point in time. It reflects the amplitude of the returns and is often seen as a good representation of the risk of a security. When the volatility is high it means that the standard deviation of the returns is high. The returns then have a large positive or negative value, which correspond to big positive movement or big negative movement. Note, however, that volatility does not measure the direction of the price changes but only the dispersion. The stock market alternates between periods of high and low volatility. Figure 2.2 shows a clear example of the consecutive high- and low-volatility periods.

The concept of volatility is important because we observe high volatility during crisis periods. Past crises can be easily seen in Figure 2.2 by the periods with large returns. Notice for example the high volatility periods in 2008 and 2020 that correspond to the stock market crashes during the Global Financial Crisis and the COVID-19 crisis, respectively.

With an understanding of the concept of volatility and its relation to world crises such as the COVID-19 pandemic, we can slightly reformulate the initial research question from the Introduction:

> Is the COVID-19 crisis just a period of high volatility or are there structural changes in the dynamics of financial markets compared to the previous period?

2.1.3 Value at Risk

"Past performance is no guarantee of future results."

This phrase is often used as a warning label for investors seeking advice on how to invest their money. There simply is no guarantee that an investment will do well just because it has done so in the past. Investors therefore try to estimate the risk of an investment. A low risk is generally seen as favorable. The most popular measure for the risk is the Value at Risk (VaR). This measure is always specified with a given confidence level α and denoted by α-VaR [17]. The α-VaR represents the maximum amount of cash you could lose with probability α (usually $\alpha = 95\%$ or $\alpha = 99\%$). Moreover, volatility and risk go together. A high-volatility stock is considered a riskier investment than a low-volatility stock. This is why investment managers are constantly watching the volatility trends. If they see the volatility go up, this might be a reason for them to retract their investment. However, stock market volatility is not necessarily a bad thing. It is not by definition bad to invest in a risky stock because there is always a chance that the prices will go up in the future. On the contrary it can actually pay off to invest in risky stocks because the profit is usually big when the stock starts doing well. This concept is also referred to as *the risk premium.* Riskier stocks with higher variances often have greater return rates [16].

2.2 Empirical Properties of Financial Time Series

Stock market observations are categorized as financial time series. The general characteristic of time series is that the observations are ordered in time, such as the daily AEX index data where Δt is one day. This allows one to see what influences the values over time and possibly predict future values by finding trends.

As explained for example by Thomas Lux (2009), we define the martingale property of financial prices as follows:

$$E[P_{t+1}|I_t] = P_t, \tag{2.3}$$

where I_t denotes the information set available at time t and P_t the stock price at time t. In other words, the expected value of the price the next day is equal to the present value, regardless of the information available. An example of a series with this property is a random walk. As a consequence we have that the realised price change (ϵ_t) is a random variable driven by the arrival of new information:

$$P_{t+1} - P_t = P_{t+1} - E[P_{t+1}|I_t] = \epsilon_t \tag{2.4}$$

with $\mathrm{E}[\epsilon_t] = 0$, for small time scales, due to the randomness of new information arrivals. On the contrary, for larger time-scales such as months and years $\mathrm{E}[\epsilon_t] = 0$ is not reasonable. In reality the mean value is greater than 0 with as a result an increase of the prices each year which can also be noticed in Figure 2.1. Moreover, $\mathrm{E}[\epsilon_t] = 0$ and Equation 2.3 are only true if the market is efficient and under the assumption that the investors are rational. For example, if the expected price tomorrow is 50 euros, you will not want to buy it for more nor sell it for less.

2.2.1 Stationarity and Ergodicity

According to Tsay (2010), the foundation of financial time series analysis is stationarity [26]. The assumption of stationarity is necessary to be able to mix data from different periods for estimates of return values. It is therefore common to assume that the stock return series are weakly stationary, also known as second-order stationary. Weakly stationary means that the time series have a mean and autocovariance that does not change with time. This is for instance visible in Figure 2.3 where the returns of the AEX index oscillate around zero. Note however that the series of stock prices, visible in the upper graph of Figure 2.3, is non-stationary. Therefore financial studies usually analyse the returns instead of the price values.

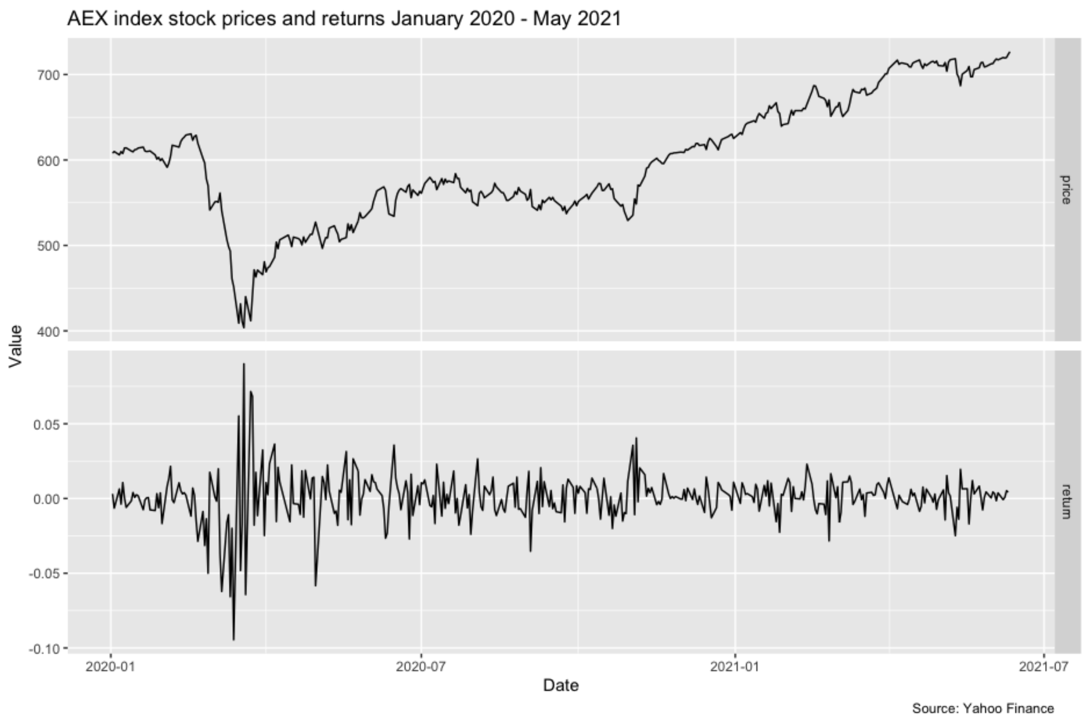

Figure 2.3: Price and return series of the AEX index data from January 2020 to May 2021. The large drawdown in the price series at the start of the COVID-19 crisis is coherent with the large return values at that time.

Is this stationarity property enough to ensure that the constant sample mean is a good representation of the actual mean? The answer is no. For this we need one more assumption:

ergodicity [6]. Ergodicity embodies the fact that a point in a stochastic process will eventually end up in every part of the space it moves in. This prevents the possibility of values getting stuck somewhere and not leaving an area of the outcome space. More observations help towards getting a more accurate estimation.

2.2.2 Stylised Facts

Market analysis is often an event-based approach for trying to explain market movements. But different stocks are not necessarily influenced by the same event and one would think they have very different properties. Does it make sense that the AEX index market has the same characteristics as a market selling completely different securities such as a commodity market? A long history of studies on financial time series has shown that there are statistical regularities that are traceable in a large number of financial time series, independent of the type of market. These statistical regularities are called *stylised facts* and are seen as the behavior roots of the trading process of a stock market.

The term stylised facts was introduced in 1961 by Nicholas Kaldor, an economist that criticised the economic growth theory at that time [5]. He constructed six statements, referred to as Kaldor's facts, summarising the statistical properties of long-term economic growth. He made the sensible argument that, before constructing a theory, one should begin from a summary of relevant facts. The term stylised fact was invented because he used scaling theory to concentrate on broad tendencies instead of individual details [5].

Up to this day stylised facts are still used widely in economics, for instance as a validation of a constructed model. Many articles and books on financial time series modeling also include research on stylised facts and the majority reference both the work of mathematician Benoît Mandelbrot and of econometrician Tim Bollerslev. Mandelbrot was important for analysing the distribution of price series and Bollerslev is known for his ideas on measuring and forecasting the stock market's volatility. A clear overview is presented in the article *Empirical properties of asset returns: stylized facts and statistical issues* by Rama Cont and the book *GARCH-models: Structure, statistical inference and financial applications* by Christian Francq and Jean-Michel Zakoian. They complement the findings well and we will often refer to both when elaborating on the stylised facts.

The following is a list of a few of the most important stylised facts.

1. *Tiny autocorrelation of returns:* Figure 2.4 (a) shows the autocorrelation of the returns of January 2007 to May 2021 for the AEX index. It can be seen that the autocorrelation is tiny and this is not only true for the AEX index but is something that is widely seen in stock markets. In general the series of price returns show very small autocorrelations. Francq and Zakoian (2010) compare it to white noise [14]. A time series is white noise if the variables are independent and identically distributed with mean equal to zero [4]. Note, however, series that hold data from smaller time steps, say in minutes or seconds (intraday time scales) instead of daily, potentially show significant autocorrelation due to microstructure effects [6].

2. *Large autocorrelations of the squared (and absolute) returns:* Squared returns (r_t^2) and absolute returns ($|r_t|$) are usually strongly correlated and this is illustrated for the AEX

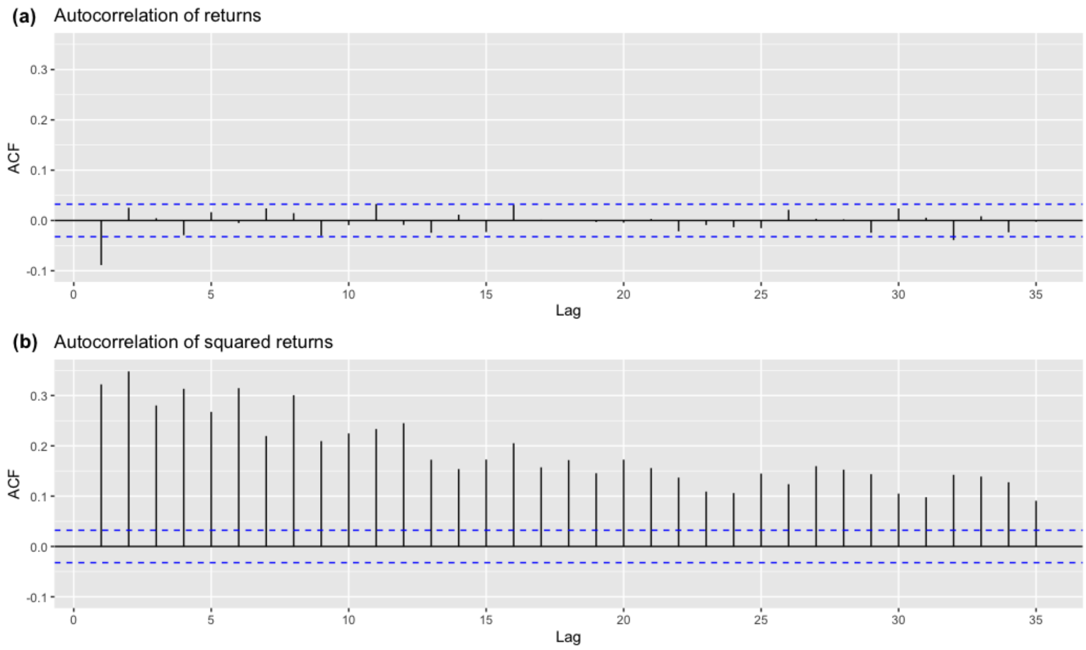

Figure 2.4: Autocorrelation of daily price returns (a) and squared price returns (b) of AEX index (January 2007 to May 2021).

index in Figure 2.4 (b). This implies that the amplitude of the returns is predictable a few days ahead, which contradicts the white noise comparison. In Figure 2.4 (b) one can additionally observe a slow decay of autocorrelation of the squared returns. Cont (2000) [6] indicates that this can be seen as a sign of long-range dependence. This property is also called *long memory* and refers to the fact that volatility changes slowly, taking a long time to go down.

3. *Volatility clustering:* Volatility clustering refers to the fact that large price variations tend to be followed by large price variations, as well as smaller variations following each other. Figure 2.2 also illustrates the clustering of high- and low-volatility periods. Notice that the largest values (denoted by the red dots) are in two periods, the crises in 2008 and 2020. There is a clear positive autocorrelation between the volatility at different moments in time.

4. *Fat-tailed distributions:* The empirical distribution of daily return values does not resemble a Gaussian distribution as is illustrated in Figure 2.5. Compared to a normal distribution there is a larger chance of extreme return values. The density distribution has fat tails, and this is clearly visible on a logarithmic scale (2.5(b)), where it decreases to zero more slowly than the normal distribution. Additionally the density function has a larger peak at the mean with respect to a normal distribution (2.5(a)). The returns are therefore nonnormal and their density is called *leptokurtic*. A measure for leptokurticity is the kurtosis coefficient. Section 3.3.3 will elaborate on the concept of kurtosis. The kurtosis of a normal distribution equals 3, but for daily return series this is much higher. For example, the AEX index returns have a kurtosis higher than 12 (see Appendix A.2.2). However, when the time interval grows, for example with monthly returns, then the leptokurticity fades away and it looks more and more like a normal distribution. This is a consequence of the central limit theorem. In the other direction, intraday return values have sharper peaks and heavier tails.

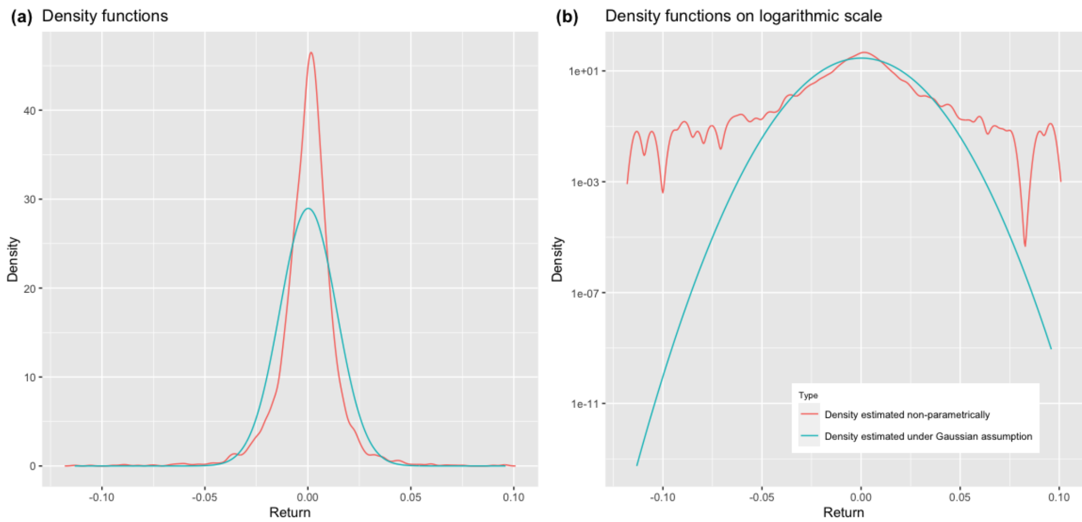

Figure 2.5: The empirical distribution of the AEX index returns (red) density function and a Gaussian (blue) distribution density function to illustrate the fat tails of leptokurtic densities. For (a) a normal scale is used and for (b) a logarithmic scale for a clear vision of the fatness of the tails.

5. *Gain/loss asymmetry:* For stock price series there is asymmetry in the up and down movement. Figure 2.1 and Figure 2.3 illustrate this fact well. One can observe the large drawdowns in prices, like the drawdown due to COVID-19 in 2020, but the upward movement is not equally big. The upward trends are less steep, reaching the peak values after a longer period of time.

6. *Leverage effect:* The leverage effect refers to the observation that most measures of volatility are negatively correlated with the corresponding returns. During crisis periods, prices go down and the volatility goes up. On the contrary, during a quiet period the volatility goes down and the price values go up. Price decreases cause higher volatility periods than price increases of the same magnitude. Francq and Zakoian (2010) note this as a clear asymmetry of the impact of past positive and negative returns on the volatility [14].

7. *Seasonality and nontrading periods:* The stock market is not open for trading all the time. It is only open on weekdays (exceptions for holidays) from approximately 9 am until 5 pm. This differs slightly per market. Seasonalities such as the day of the week or periods close to a holiday can very well influence the behaviour of the returns. This is not only true for daily data but also for intraday series. Additionally, nontrading periods can affect the volatility. When the market is closed, information on price variations accumulates slower than when they are open. The result can only be reflected when the markets reopen causing the volatility to then increase. This is why there is more variability in price returns after the weekends and holidays.

8. *Spillover effect:* The spillover effect refers to information spilling over within and across stock markets. This is also called *co-movement in volatility.* Nowadays, volatility of a certain stock in a market can influence other securities of that same market as well as other markets in general. This can even happen between different countries, where the state of a globally important stock market can affect the markets in other countries. When an event occurs, such as a natural disaster or a stock market crisis, this can influence many different markets and countries. The Global Financial Crisis in 2008 is a good example of this, where the crisis originated in the United States and influenced many other countries.

In conclusion, modeling stock market volatility is quite complicated since there is a lot to capture in the model. There are even more factors than mentioned above and, because there are so many things influencing volatility, it can be difficult to trace their exact impact. Even when they can be identified, artificially reproducing them in a model can be a complex task. Luckily, through many studies on asset returns several volatility estimation models have been introduced and the next step is finding one that incorporates most of the empirical properties.

3 GARCH Models

This chapter sets out several definitions and properties of a family of estimation models called Generalised AutoRegressive Conditionally Heteroscedastic (GARCH) models. The first section introduces the definition of conditional heteroscedasticity which accounts for the time-varying volatility property of stock price series. The ARCH and Generalised ARCH (GARCH) are two models famous for modeling the conditional heteroscedasticity. Sections 3.2 and 3.3 provide detailed definitions and property descriptions of the ARCH and GARCH models, and also check their ability to reflect the important stylised facts covered in Section 2.2.2. Specifically the simple case of GARCH(1,1) is thoroughly discussed and shown to fit our data the best compared to all GARCH(p, q) models. Moreover, in Section 3.4 the procedure of estimating GARCH models is presented. For the estimation we use the R environment since it offers built in packages that make it easy to estimate the parameters. R estimates the parameters by the Quasi-Maximum Likelihood Estimation (QLME) method, which we will present in Section 3.4.2. Lastly we will briefly discuss some limitations of GARCH models but explain why we use them for our research nonetheless.

3.1 Conditional Heteroscedasticity Models

Econometricians are nowadays being called in to build a model of financial returns, for example for estimating the precise Value at Risk [23], and thus to model the volatility of financial time series [11]. In the previous chapter we described properties that explain the complexity of financial time series modeling. Francq and Zakoian (2010) stress that the main stylised facts must be taken into account for any satisfactory statistical model for daily stock returns. Similarly Tsay (2010) stresses that a satisfactory estimation model must especially capture the volatility properties of asset returns such as volatility clustering, stationarity and continuity of volatility. The latter two describe the fact that volatility does not diverge to infinity and that today's volatility is strongly related to the volatility yesterday, respectively.

It seems logical to first consider some of the traditional time series and econometric models such as linear models and ARMA but these do not work in our situation. These models operate under the assumption of constant volatility whereas the conditional variance of returns changes over time [2]. This property of non-constant conditional variance is called *conditional heteroscedasticity*:

$$\mathrm{Var}(\epsilon_t|\epsilon_{t-1}, \epsilon_{t-2}, \dots) \not\equiv \text{constant}.$$

By assuming that the volatility is constant, the statistician would not be adapting to the situation of the stock markets. When allowing it to be conditional the statistician is using the information to have a more precise measurement.

3.2 The ARCH(q) Process

The first model that dealt with heteroscedasticity was the ARCH (AutoRegressive Conditional Heteroscedasticity) model introduced by Engle in 1982. This ARCH process acknowledges the time-varying volatility by letting the conditional variance be a function of past errors [2]. This makes it an autoregressive model and allows the uncertainty to be a dynamic process [22]. Using the notation of Francq & Zakoian (2010), an ARCH(q) model looks as follows:

$$r_t = \mu + \epsilon_t \quad \text{(returns)} \tag{3.1}$$

$$\epsilon_t = \sigma_t \eta_t \quad \text{(conditional average)} \tag{3.2}$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 \epsilon_{t-1}^2 + \ldots + \alpha_q \epsilon_{t-q}^2 \quad \text{(conditional variance)} \tag{3.3}$$

The r_t denotes the returns of a financial asset at time t and is the sum of the mean μ and the conditional average ϵ_t. Next σ_t^2 represents the conditional variance of the returns at time t and (η_t) is a sequence of identically distributed (iid) random variables with mean zero and variance equal to 1. Often a standard normal distribution is assumed as the distribution of η [26].

To show that this process indeed models the volatility of returns, we look at the simple ARCH(1) process. By substituting Equation 3.3 into Equation 3.2 we rewrite the equation for ϵ_t a bit and obtain the following:

$$\epsilon_t = \eta_t \sqrt{\alpha_0 + \alpha_1 \epsilon_{t-1}^2}.$$

Now by squaring both sides of the equation and taking the variance of ϵ_t we see that:

$$\begin{aligned} \mathrm{Var}(\epsilon_t) &= \mathrm{E}[\epsilon_t^2] - (\mathrm{E}[\epsilon_t])^2 \\ &= \mathrm{E}[\epsilon_t^2] \\ &= \mathrm{E}[\eta_t^2]\mathrm{E}[\alpha_0 + \alpha_1 \epsilon_{t-1}^2] \\ &= \mathrm{E}[\alpha_0 + \alpha_1 \epsilon_{t-1}^2] \\ &= \alpha_0 + \alpha_1 \mathrm{Var}(\epsilon_{t-1}). \end{aligned} \tag{3.4}$$

Here we used that η_t has zero mean and unit variance, and the fact that $\mathrm{E}[\epsilon_t] = 0$ (Equation 2.4). Equation 3.4 shows that the variance of ϵ_t is a linear function of the variance of the return value of the previous time-step. We see that large past returns ϵ_{t-1} result in large conditional variance for ϵ_t, which means it is more likely to have large returns than small returns. Hence this is the reason why this model can reproduce the volatility clustering we observe in asset returns.

3.3 The GARCH(p, q) Process

The ARCH model represents a key invention in the field of financial modeling and its creator Robert Engle indeed won the Nobel Prize for economics in 2003. Even though the model is quite simple, on the other hand it often requires many lags to properly describe the volatility process of an asset return [26]. As a solution to this flaw, Bollerslev (1986) introduced an extension to the ARCH model: a new and more general class of processes named Generalised AutoRegressive Conditionally Heteroscedastic (GARCH) [2]. Bollerslev explains that the Generalised ARCH models allow for both longer memory and a much more flexible lag structure. He extended the model in a way that resembles the extension of a standard time series AR (AutoRegressive) process to a general ARMA (AutoRegressive Moving-Average) process, letting the conditional variance be an ARMA process. GARCH models provide a slightly better fit of the time-varying volatility, as well as most of the other stylised facts mentioned in Section 2.2.2.

3.3.1 Definitions

For the definition of the GARCH(p,q) process we will follow the book on GARCH Models by Francq & Zakoian (2010). They start with a definition based on the first two conditional moments, the *semi-strong* GARCH(p,q) process, and after define the *strong* GARCH(p, q) process from which the GARCH model is constructed.

Definition 1 (Semi-strong GARCH(p, q) process)**.** *Let I_t denote the information set available at time t. A process (ϵ_t) is called a GARCH(p, q) process if its first two conditional moments exist and satisfy:*

(i) $\epsilon_t|I_t \sim \mathcal{N}(0, \sigma_t^2)$, $t \in \mathbb{Z}$.

(ii) There exist constants ω, α_i and β_j, with $i = 1, \ldots, q$ and $j = 1, \ldots, p$, such that

$$\sigma_t^2 = \omega + \sum_{i=1}^{q} \alpha_i \epsilon_{t-i}^2 + \sum_{j=1}^{p} \beta_j \sigma_{t-j}^2, \quad t \in \mathbb{Z}. \tag{3.5}$$

In Definition 1 ϵ_t once again denotes the real-valued, discrete time stochastic process of the asset returns [2] and σ_t^2 the corresponding conditional variance. The constant parameters must satisfy a few constraints. We have $\omega > 0$, $\alpha_i \geq 0$, $\beta_j \geq 0$ and, additionally, they must satisfy:

$$\sum_{i=1}^{r} (\alpha_i + \beta_i) < 1. \tag{3.6}$$

Here we define $r := \max(p, q)$ and, moreover, $\alpha_i := 0$ if $i > q$ and $\beta_j := 0$ if $j > p$. The constraint in Equation 3.6 ensures that the variance does not eventually explode. To illustrate what would happen if this constraint is not satisfied, look at Equation 3.4 for the ARCH case. If we would have $\alpha_1 > 1$, then with every time step $\mathrm{Var}(\epsilon_t)$ increases in a geometric way, causing it to eventually explode.

Francq & Zakoian define the innovation of the process ϵ_t^2 as the variable $\nu_t = \epsilon_t^2 - \sigma_t^2$. Rewriting this slightly such that $\sigma_{t-j}^2 = \epsilon_{t-j}^2 - \nu_{t-j}$ and substituting in Equation 3.5 we get the following representation:

$$\epsilon_t^2 = \omega + \sum_{i=1}^{r} (\alpha_i + \beta_i) \epsilon_{t-i}^2 + \nu_t - \sum_{j=1}^{p} \beta_j \nu_{t-j}. \tag{3.7}$$

Equation 3.7 illustrates the linear structure of an ARMA model and the fact that if (ϵ_t) is GARCH(p, q), then (ϵ_t^2) is an ARMA(r, q) process [14].

Remark (Autocorrelation characteristics)*.* In Section 2.2.2 we observed some characteristic features with respect to the autocorrelation of both the returns and the squared returns. Firstly we saw that the returns show very small autocorrelation. This empirical fact is captured by Definition 1, since (i) shows that:

$$\mathrm{E}[\epsilon_t|\epsilon_u, u < t] = 0,$$

which implies that the returns ϵ_t are uncorrelated. Secondly, contrary to the absence of correlation between return values, there is a clear autocorrelation between the squared returns. We see that this empirical characteristic is also captured in the GARCH process by looking at Equation 3.7. Consider the simple GARCH(1,1) case. If the kurtosis is finite, then the squared process (ϵ_t^2) is ARMA(1,1) and thus we have the following equation for the h-order autocorrelations:

$$\mathrm{Corr}(\epsilon_t^2, \epsilon_{t-h}^2) = K(\alpha_1 + \beta_1)^h,$$

where K is a constant independent of h [14].

In spite of reflecting desired properties of asset returns, Francq & Zakoian note that Definition 1 does not offer a direct solution that satisfies the above conditions. For this reason they define a slightly more restrictive class of processes that, however, enables the discovery of explicit solutions. In the same manner as for the ARCH model, we again define (η_t) as a sequence of iid random variables and represent the GARCH model in the following way:

Definition 2 (Strong GARCH(p, q) process)**.** *Let (η_t) be an iid sequence with distribution η. The process (ϵ_t) is called a strong GARCH(p, q) (with respect to the sequence (η_t)) if*

$$\begin{cases} \epsilon_t &= \sigma_t \eta_t \\ \sigma_t^2 &= \omega + \sum_{i=1}^q \alpha_i \epsilon_{t-i}^2 + \sum_{j=1}^p \beta_j \sigma_{t-j}^2 \end{cases} \tag{3.8}$$

where the α_i and β_j are non-negative constants, ω is a (strictly) positive constant and $\sigma_t > 0$.

Notice that if we have $\beta_j = 0$ for all values of $j = 1, \dots, p$, the process is reduced to ARCH(q) process (3.3). Accordingly α is often referred to as the ARCH coefficient and β as the GARCH coefficient. Moreover, substituting $\epsilon_{t-i} = \sigma_{t-i}\eta_{t-i}$ in (3.8) we get the following:

$$\sigma_t^2 = \omega + \sum_{i=1}^r (\alpha_i \eta_{t-i}^2 + \beta_i)\sigma_{t-i}^2,$$

where again $r := \max(p, q)$. From this representation the volatility process of a strong GARCH process is shown to be the solution of an autoregressive equation with random coefficients [14]. Furthermore this process indeed models the volatility of returns, which is easily seen by rewriting the variance in the same way as in Section 3.2 for the ARCH process. Next we will consider the simple case of GARCH(1,1) and highlight some properties of the parameters and variables in the model.

3.3.2 GARCH(1,1)

From Definition 2 we obtain the GARCH(1,1) model by setting p and q equal to 1 in Equation 3.8:

$$\begin{cases} \epsilon_t = & \sigma_t \eta_t \\ \sigma_t^2 = & \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \end{cases} \tag{3.9}$$

with $\omega > 0$, $\alpha \geq 0$, $\beta \geq 0$ and (η_t) iid (0,1). Let μ denote the mean of the returns, and r_t is then modeled as follows:

$$r_t = \mu + \epsilon_t. \tag{3.10}$$

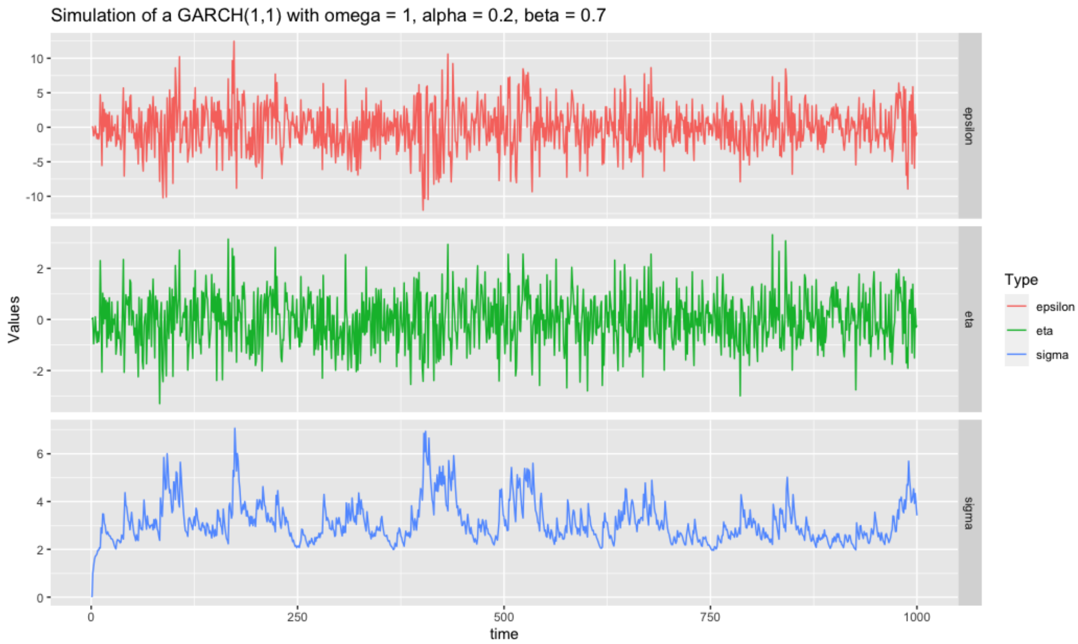

Figure 3.1: Simulation of the GARCH(1,1) process during 1000 timesteps with $\omega = 1$, $\alpha = 0.2$ and $\beta = 0.7$ and $\eta \sim \mathcal{N}(0, 1)$.

To illustrate how ϵ_t is obtained by multiplying σ_t by η_t we present Figure 3.1. This figure corresponds to a simulation of the GARCH(1,1) process during 1000 timesteps with parameters $\omega = 1$, $\alpha = 0.2$ and $\beta = 0.7$. To model the unexpected shock we let $\eta_t \sim \mathcal{N}(0, 1)$. A clear representation of how the different variables and parameters influence one another is given in Figure 3.2.

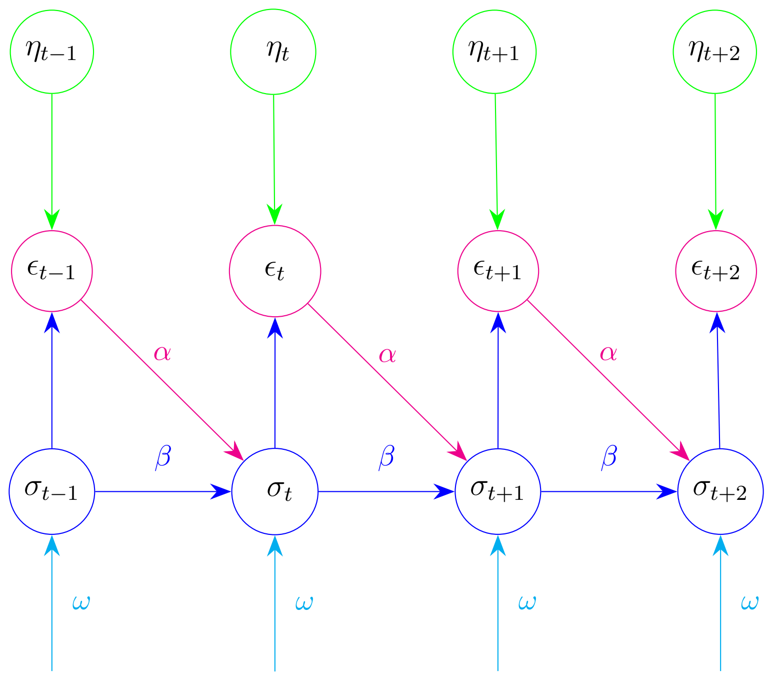

Figure 3.2: Directed Acyclic Graph (DAG) of the GARCH(1,1) model.

The parameter ω represents the variance during the *calm* periods. That is, when $\epsilon_{t-1} \approx 0$ and $\sigma_{t-1} \approx 0$, then we have $\sigma_t^2 \approx \omega$. As mentioned before, α is the ARCH parameter and since $\epsilon_t = \sigma_t \eta_t$, it models the influence of the unexpected shock from the random variable η_t. The GARCH parameter β represents the influence of the variance the day before. Furthermore, to not be explosive we must have:

$$\alpha + \beta < 1. \tag{3.11}$$

By Theorem 2.1 and 2.2 from the book on GARCH models by Francq & Zakoian it holds that if the model described by (3.9) satisfies (3.11), then we have a strict and second order stationary solution of the GARCH(1,1). Note that this is also presented by the simulations in Figure 3.1. Moreover, the solution we obtain for ϵ_t is ergodic.

Section 3.3.1 showed how the two stylised facts on autocorrelation are captured in the GARCH models. Another empirical fact that is embedded in the model is volatility clustering. The GARCH structure lets the amplitude of ϵ_t be a function of its past values [14] and, therefore, high (respectively low) volatility periods will be followed by high volatility periods. This is also illustrated in Figure 3.1 where we see that the extreme values are not independently distributed on the entire time line.

3.3.3 Kurtosis

In Section 2.2.2 we introduced the kurtosis coefficient as a measure for leptokurticity, the heaviness of the tails of a distribution. Recall that a normal distribution has kurtosis coefficient equal to 3, but for distributions with fat tails this value is higher. The kurtosis coefficient κ is based on the fourth order moment of a distribution (μ_4) which is defined as:

$$\mu_4 = \mathrm{E}((r_t - \mu)^4) = \mathrm{E}(\epsilon_t^4)$$

since we have $r_t - \mu = \epsilon_t$ (3.10). The kurtosis coefficient κ is defined as follows:

$$\kappa = \frac{\mu_4}{\sigma^4}.$$

From this definition we can compute the kurtosis coefficient of the stationary marginal distribution of a GARCH-process by Defintion 2:

$$\kappa_\epsilon = \frac{\mathrm{E}(\epsilon_t^4)}{\Big(\mathrm{E}(\epsilon_t^2)\Big)^2} = \frac{\mathrm{E}\Big(\mathrm{E}(\epsilon_t^4|\epsilon_u, u<t)\Big)}{\Big(\mathrm{E}\Big(\mathrm{E}(\epsilon_t^2|\epsilon_u, u<t)\Big)\Big)^2} = \frac{\mathrm{E}(\sigma_t^4)}{\Big(\mathrm{E}(\sigma_t^2)\Big)^2}\,\kappa_\eta. \tag{3.12}$$

The first equality shows that κ is the ratio of the fourth-order moment to the squared second-order moment which follows from the fact that $\mathrm{E}(\epsilon_t)^2 = 0$. From the last equality the relation to κ_η, the kurtosis coefficient of (η_t), can been seen. Recall that (η_t) is independently identically distributed with zero mean and unit variance, and therefore has $\kappa_\eta = \mathrm{E}(\eta_t^4)$. Equation 3.12 shows that κ_ϵ is greater than or equal to κ_η. The distribution of (ϵ_t) has fatter tails when the ratio of the variance of σ_t^2 to the squared expectation is large [14].

Using the conditional fourth-order moment we obtain that the conditional kurtosis of ϵ_t ($\kappa_{\epsilon|t}$) is constant and equal to κ_η:

$$\kappa_{\epsilon|t} = \frac{\mathrm{E}(\epsilon_t^4|\epsilon_u, u<t)}{\sigma_t^4} = \frac{\sigma_t^4 \mathrm{E}(\eta_t^4)}{\sigma_t^4} = \mathrm{E}(\eta_t^4) = \kappa_\eta.$$

This shows the difference between the tails of the marginal and conditional distributions [14].

Ultimately, from equation 3.12 we can calculate κ_ϵ for the GARCH(1,1) case:

$$\kappa_\epsilon = \frac{1-(\alpha+\beta)^2}{1-(\alpha-\beta)^2-\alpha^2(\kappa_\eta-1)}\ \kappa_\eta \tag{3.13}$$

and if $\eta_t \sim \mathcal{N}(0,1)$ then we have $\kappa_\eta = 3$.

3.4 Estimating GARCH Models

Now that we have defined GARCH-type processes and discussed several properties, the next step is estimating such a model. For our research we want to investigate European stock markets and therefore the first model we estimate is based on data from the AEX index. The statistical programming language R consists of tools that make the computation of the estimation quite simple. It uses Quasi-Maximum Likelihood Estimation (QLME) for the coefficients. Additionally we will discuss which GARCH specification is the best fit for our estimation and what criterion we use for this determination.

3.4.1 Importing the data in R

To estimate the parameter we use the R analytics environment. R is an open source statistical programming language and is widely used in academia and the financial industry [23] [22]. R grants access to packages through an open library created by users. The package we use for modeling volatility is the 'fGarch' package, offering tools for simulating, estimating and forecasting GARCH models. We focus on the estimation of the model and therefore use the function `garchFit` with the following function description:

`garchFit: fits the parameters of a GARCH process` [28].

The first stock price series we wish to model is the one already frequently mentioned in this report, the AEX index. We extracted the daily stock prices of the AEX index from the Yahoo! Finance database using the 'quantmod' package. The Yahoo! Finance source also holds data from different market indices such as the French CAC 40 (Cotation Assistée en Continu) and the German DAX (De Deutscher Aktien) that will be discussed in the next chapter. We downloaded the AEX index data from January 2007 through June 2021, including the 2008 Global Financial Crisis and the current 2020 COVID-19 crisis. We use the available daily opening prices as P_t and use Equation 2.1 to compute r_t. We call the `garchFit` function on the returns (r_t), which gives us the desired estimated parameters corresponding to a GARCH model based on the AEX index data.

Summary output of the `garchFit` function can be found in the Appendix (A.1). It displays multiple statistics, such as the estimated parameters with their corresponding standard error values, t-values and p-values. Additionally it displays several tests concerning the standardised residuals as well as several goodness-of-fit measures, for example the log-likelihood value and the information criterion values: the Aikaike Information Criterion (AIC), the Bayesian Information Criterion (BIC), the Schwarz Information Criterion (SIC) and the Hannan-Quinn Information Criterion (HQIC).

3.4.2 Quasi-Maximum Likelihood Estimation

The function `garchFit` in the package 'fgarch' estimates the parameters by the Quasi-Maximum Likelihood Estimation (QMLE). We will present the procedure of the conditional (on initial values) QML method following Chapter 7 of the book by Francq & Zakoian, but specify for the simple case of GARCH(1,1).

Let $\epsilon_1, \ldots, \epsilon_n$ be a realisation of the GARCH(1,1) process satisfying Equations 3.9 and 3.10. Let θ denote the vector of parameters, belonging to the parameter space $\Theta \subset (0, \infty) \times [0, \infty)^2$, and define it as follows:

$$\theta = (\theta_1, \theta_2, \theta_3)' := (\omega, \alpha, \beta)'. \tag{3.14}$$

The true value of the parameters is unknown and denoted by

$$\theta_0 = (\omega_0, \alpha_0, \beta_0)'.$$

We need to specify a distribution for η_t in order to compute the likelihood. We therefore work with a function called the *quasi-likelihood*, which coincides with the likelihood when η_t has a standard Gaussian distribution [14]. Given initial values ϵ_0 and $\tilde{\sigma}_0^2$ and under the second-order stationarity assumption it is reasonable to let

$$\epsilon_0^2 = \tilde{\sigma}_0^2 = \frac{\omega}{1 - \alpha - \beta}. \tag{3.15}$$

The Gaussian quasi-likelihood is then defined as

$$L_n(\theta) = L_n(\theta; \epsilon_1, \ldots, \epsilon_n) = \prod_{t=1}^{n} \frac{1}{\sqrt{2\pi\tilde{\sigma}_t^2}} \exp\left(-\frac{\epsilon_t^2}{2\tilde{\sigma}_t^2}\right),$$

where we have

$$\tilde{\sigma}_t^2 = \tilde{\sigma}_t^2(\theta) = \omega + \alpha\epsilon_{t-1}^2 + \beta\tilde{\sigma}_{t-1}^2.$$

Finally the QLME of θ is then defined as any measurable solution $\hat{\theta}_n$ of

$$\hat{\theta}_n \in \arg\min_{\theta \in \Theta} L_n(\theta).$$

There exists a whole theory about the QMLE estimators and how they satisfy the usual convergence and asymptotic properties. For the proof of these theorems we refer to the book of Francq & Zakoian (pages 177 - 187). The theorems on the consistency and asymptotic normality of the estimators are particularly interesting. Francq & Zakoian prove that the QMLE of the coefficients of a GARCH model is asymptotically normal even if the distribution of the variables η_t is not Gaussian. Therefore the inference is valid even if the assumption of a Gaussian distribution is not true, for example if η_t follows a student-t distribution.

3.4.3 Finding the Best GARCH Specification

The next step is to find the most appropriate GARCH(p, q) model for our dataset. We use a model selection procedure based on the information criterion AIC to determine for which values p and q the model shows the best fit. By adding lags, so increasing p and q, we are adding a direct influence from past values. For example for a GARCH(1,2) model, the schematic structure

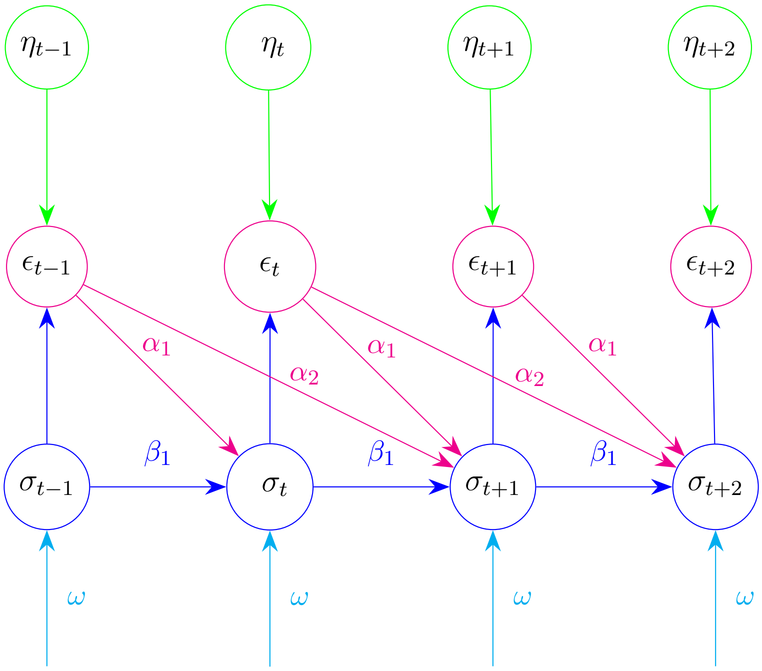

Figure 3.3: Directed Acyclic Graph (DAG) of the GARCH(1,2) model.

of variables can be viewed in Figure 3.3.

Note that the difference with the GARCH(1,1) structure in Figure 3.2 is the addition of α_2 which includes a direct influence of ϵ_t on σ_{t+2}. Furthermore, for a GARCH(2,2) model we would add the parameter β_2, adding a direct influence from the variance further in the past, and so on. It is not a strange thought that adding more variables from the past can improve the approximation. It could very well be that the volatility of today is dependent on the volatility of two days ago and therefore must be accounted for in the model. Nevertheless, a model with a low number of coefficients is preferred. Both Engle (2001) and Tsay (2010) argue that the GARCH(1,1) model is the most robust and that higher-order models do not succeed in outperforming the GARCH(1,1) model [23]. We can test if these findings by Engle and Tsay are also true for our dataset by measuring which fits our data the best.

To find the best GARCH(p,q) specification we compare their AIC (Aikaike Informatoin Criterion) values which are provided by the R summary outputs of `garchFit`. The AIC is computed by the following formula:

$$AIC = \frac{-2}{N_{obs}}(loglik - N_{par}). \tag{3.16}$$

Here *loglik* is the resulting log-likelihood, N_{obs} the number of observations and N_{par} the number of parameters of the model. This criterion penalizes the number of coefficients in the model by subtracting N_{par} from the log-likelihood. Preferring the model with the smallest number of parameters when choosing between two models with equal likelihood comes from an epistemological concept known as *Occam's Razor*. Occam's Razor is a problem solving principle and states that *Entities should not be multiplied without necessity* [24]. The best model is therefore defined as the one with the smallest AIC, since it then has a high loglikelihood and a low number of parameters.

We let p and q take on values 1 through 4 and compared the AIC values of all possible combinations of p and q. The GARCH(1,1) was a clear winner with the lowest AIC value and for this reason we chose this model for our estimation. This test does not show us if GARCH models in general are good for our situation, but we do know now that of all GARCH(p, q) models, GARCH(1,1) gives the best fit.

3.5 Limitations and Extensions of GARCH Models

In this report we will not go into detail on the goodness-of-fit test for GARCH models in general. No model can perfectly portray every empirical property of financial time series and we therefore follow the in statistics well-known expression by George Box [3]:

"All models are wrong, but some are useful."

We are not saying that GARCH perfectly models our stock price data, but they do present useful findings for our research question. Even if GARCH models would be rejected by a goodness-of-fit test it does not mean that they are bad. They can still be very useful and informative regardless. In order to better capture the stylised facts many extensions of the classical GARCH models have been designed. Nevertheless, many studies have also shown that these extensions do not necessarily greatly outperform the classical model. We therefore choose to stick with the simple GARCH model. There are, however, three drawbacks of the classical GARCH models that we find are worth mentioning.

The first one is the most obvious one, namely the symmetric form of a classical GARCH model. The conditional variance depends on the modulus of past variables and therefore the impact from past positive and negative innovations is the same [14]. This contradicts the empirical stylised fact called *gain/loss asymmetry* that occurs in reality. In past data we observe large drawdowns in prices, for instance in crisis periods, but the following upward movement is not equally big. To allow this asymmetry, an option is to use the EGARCH extension (Exponential GARCH) model introduced by Nelson (1991).

Secondly, Tsay (2010) highlights that certain empirical studies on financial time series show that the tail behavior of GARCH models is still too short in comparison to actual observations. This is shown by looking at the kurtosis coefficients. Assuming a normal distribution for (η_t) and therefore letting $\kappa_\eta = 3$ we find that the kurtosis coefficient for the GARCH model with estimated coefficients is still too small to portray the kurtosis of the empirical findings. To improve the kurtosis coefficient we can let η have a student-t distribution, but Tsay (2010) states that even then the tails are still too short.

Lastly, the leverage effect property of stock price series is not fully captured by the GARCH model definition. It is known that price decrease causes higher volatility than price increase of the same magnitude, and therefore does not respond equally to positive and negative shocks [26]. Classical GARCH models fail to reproduce this due to the dependence of the conditional variance on the magnitude of the lagged residuals apposed to their signs. The TGARCH extension (Threshold GARCH) model is known to adequately capture this fact as it uses zero as a threshold to separate impacts from past shocks [26].

4 Modeling Stock Indices and the COVID-19 Crisis

This chapter sets out the methodology of our research for finding the answer to the research question, whether COVID-19 has structurally changed the dynamics of the stock market. We are interested in European markets and therefore estimate GARCH(1,1) models on stock market indices from three different countries in Europe. First we continue with the Dutch AEX index. GARCH processes are designed to model high and low volatility, but we are curious to see how it models the period during COVID-19. Throughout this chapter we estimate models based on different subsets of the AEX index dataset and compare the estimated parameter values by their 95% confidence intervals. This way we can compare models fitted on data from before and during COVID-19 periods. Section 4.4 introduces the concept of parametric bootstrapping, a method that tests the accuracy of estimations, and illustrates in three clear steps how it is used in our research. Afterwards we apply the same computations to two other European indices, the French Cotation Assistée en Continu (CAC 40) and the German Deutscher Aktien (DAX). In Section 4.5 we compare the estimated results of the different markets. All estimations and computations discussed in this chapter were performed in R and several relevant codes and results have been presented in the Appendix.

4.1 Splitting the data

The AEX index is the first index we model and the available data ranges from January 2007 until June 2021. The first step is finding the start date of the COVID-19 stock market crisis. In multiple figures in Chapter 2 we have seen the deep drawdown of the stock prices in March 2020. The largest absolute return values occurred between March 11 and 17, the same week that the Netherlands closed all shops, restaurants and other public areas. But actually, the crisis had influenced the market earlier than this date. After investigating the returns we define the COVID start date as January 30, 2020. This is when the large fluctuations of the returns started, which is visible in Figure 2.3. With this COVID cut-off date we distinguish the periods Before COVID (BC) and Post COVID (PC). We define all dates before 30/01/2020 as BC and all dates after as PC, even though the PC period is technically *during* the COVID-19 crisis.

We want to estimate different models based on data from BC and PC periods. Since we have data starting in January 2007, we take different BC periods to compare with the PC period that is only a little over a year long. We distinguish three different BC periods and base them on past periods of high volatility to have diversity between the datasets. Each BC period will end 30/01/2020 but their start date will differ. We recognise two of the past phases with high volatility in Figure 2.2: the Global Financial Crisis in 2008 and the sudden market crash in 2015/16 (from now on referred to as the 2015 crash). The first BC period is chosen not too far back in time and begins right after the high volatility period of the 2015 crash, on 12/01/2016. The second period is chosen to start before this high volatility phase and has start date: 10/10/2014. For the third period we want to include the Global Financial Crisis in 2008 and therefore set its start date on 01/01/2008. There are now five datasets on which we estimate a model and each period is denoted by the following names:

$$\begin{array}{lllr}
\text{Period 1}: & [12/01/2016 - 01/30/2020] & \text{End 2015 Crash} & (\text{BC1}) \\
\text{Period 2}: & [10/10/2014 - 01/30/2020] & \text{Begin 2015 Crash} & (\text{BC2}) \\
\text{Period 3}: & [01/01/2008 - 01/30/2020] & \text{Begin 2008 Crash} & (\text{BC3}) \\
\text{Period 4}: & [01/31/2020 - 04/30/2021] & \text{Post COVID} & (\text{PC}) \\
\text{Period 5}: & [01/01/2007 - 04/30/2021] & \text{All data} & (\text{All})
\end{array}$$

4.2 First Estimation

The code used for the estimation of the five defined periods can be found in the Appendix (A.2.3). The estimation of the coefficients looks as follows:

$$(r_t) \xrightarrow[(QLME)]{\text{Estimation}} \begin{pmatrix} \hat{\mu}_{BC1}, & \hat{\mu}_{BC2}, & \hat{\mu}_{BC3}, & \hat{\mu}_{PC}, & \hat{\mu}_{all} \\ \hat{\omega}_{BC1}, & \hat{\omega}_{BC2}, & \hat{\omega}_{BC3}, & \hat{\omega}_{PC}, & \hat{\omega}_{all} \\ \hat{\alpha}_{BC1}, & \hat{\alpha}_{BC2}, & \hat{\alpha}_{BC3}, & \hat{\alpha}_{PC}, & \hat{\alpha}_{all} \\ \hat{\beta}_{BC1}, & \hat{\beta}_{BC2}, & \hat{\beta}_{BC3}, & \hat{\beta}_{PC}, & \hat{\beta}_{all} \end{pmatrix}$$

where $\hat{\mu}$, $\hat{\omega}$, $\hat{\alpha}$ and $\hat{\beta}$ are called the *pseudo-true values*. To compare the estimations we use error bar plots that portray the different confidence intervals, which can be viewed in Figure 4.1.

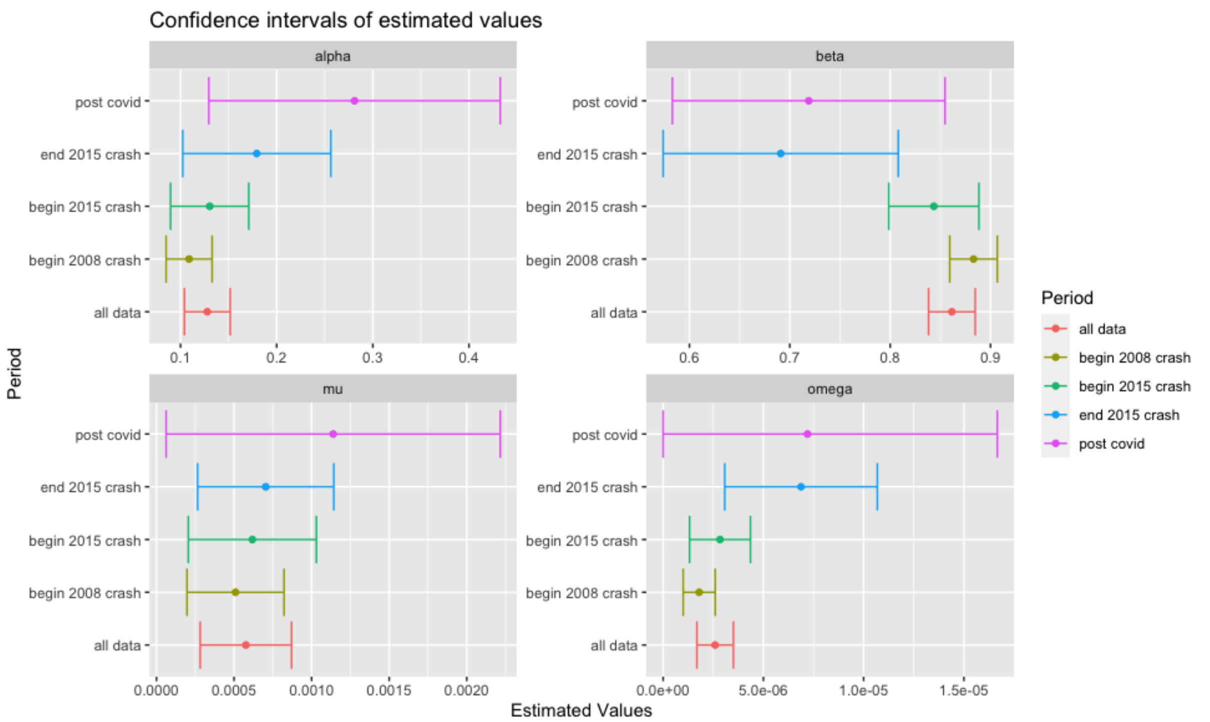

Figure 4.1: Estimating the GARCH(1,1) model using the AEX index data. The estimated parameter values and the corresponding 95% confidence intervals for five different periods (three BC periods, the PC period and the entire dataset) are presented with error bars.

To start, we look at the graph with the estimated values of the mean μ. Recall that $r_t = \mu + \epsilon_t$.

We see that for each period $\mu > 0$ and, moreover, that μ is increasing over time, implying that there is an overall upward trend in the stock market. This seems to be coherent with reality if we look at Figure 2.1. Even though the prices are going up and down constantly, the mean value has been increasing since the Global Financial Crisis.

We are especially interested in the cases where the confidence intervals of different periods are disjoint and we see this happen for the parameters α and β. This implies that the model has changed in time, however, these changes do no seem a consequence of COVID-19. In fact we can see that the period right before COVID, the End 2015 Crash period, looks a lot like the PC period. The models estimated on those periods are not significantly different at the level of 95%. We therefore cannot reject the assumption that they are equal. Looking for example at the confidence intervals for β, we notice that the intervals corresponding to the PC and the End 2015 Crash are very similar. Even though the End 2015 Crash period did not contain data with extreme return values, the model still seems very similar to that of the PC model, supporting the fact that there have not been changes in the model due to COVID-19.

On the contrary, the interval for Begin 2008 Crash is completely disjoint with the PC interval for α and β. We therefore have strong evidence that these two models are different. This implies that there indeed have been changes in the dynamics of the stock market since the 2008 market crash. However, it seems that these changes are due to other changes in the world that occurred between 2008 and 2015, and not because of COVID-19.

Lastly, we want to point out the differences in the sizes of the confidence intervals. Obviously the interval for the PC is the largest, which is due to the fact that it is estimated on the smallest dataset. When the dataset is smaller the estimated values are usually less accurate, resulting in a larger confidence interval. Comparatively the All data period has the smallest confidence interval. Because of this fact there might be some limitations to this comparison and to tackle this we proceed with new estimations that will be discussed in the next section.

4.3 Second Estimation

Based on the similar confidence intervals in Figure 4.1 of the estimated parameters for the End 2015 crash and PC periods it seems that COVID-19 has not caused for much change in the model. However, there might be some limitations to this comparison concerning the sizes of the two datasets. The PC period is only a little over a year, while the End 2015 Crash consists of four years of data. For the other BC periods this difference is even larger. We therefore proceed with more estimations on smaller periods. Moreover, this also gives us a more precise view of what is really happening in the price series.

We split the timeline into many new periods that all have the same size. Starting in January 2007 and ending in April 2021, we have 158 periods that all contain 14 months of data. Each new period is constructed by shifting the previous period up one month, which results in the 158 periods in total. We now have:

$$(r_t) \xrightarrow[(QLME)]{\text{Estimation}} \begin{pmatrix} \hat{\mu}_1, & \dots, & \hat{\mu}_{158} \\ \hat{\omega}_1, & \dots, & \hat{\omega}_{158} \\ \hat{\alpha}_1, & \dots, & \hat{\alpha}_{158} \\ \hat{\beta}_1, & \dots, & \hat{\beta}_{158} \end{pmatrix}$$

and the results are presented in Figure 4.2. The code corresponding to these estimation can be found in the Appendix (A.2.5).

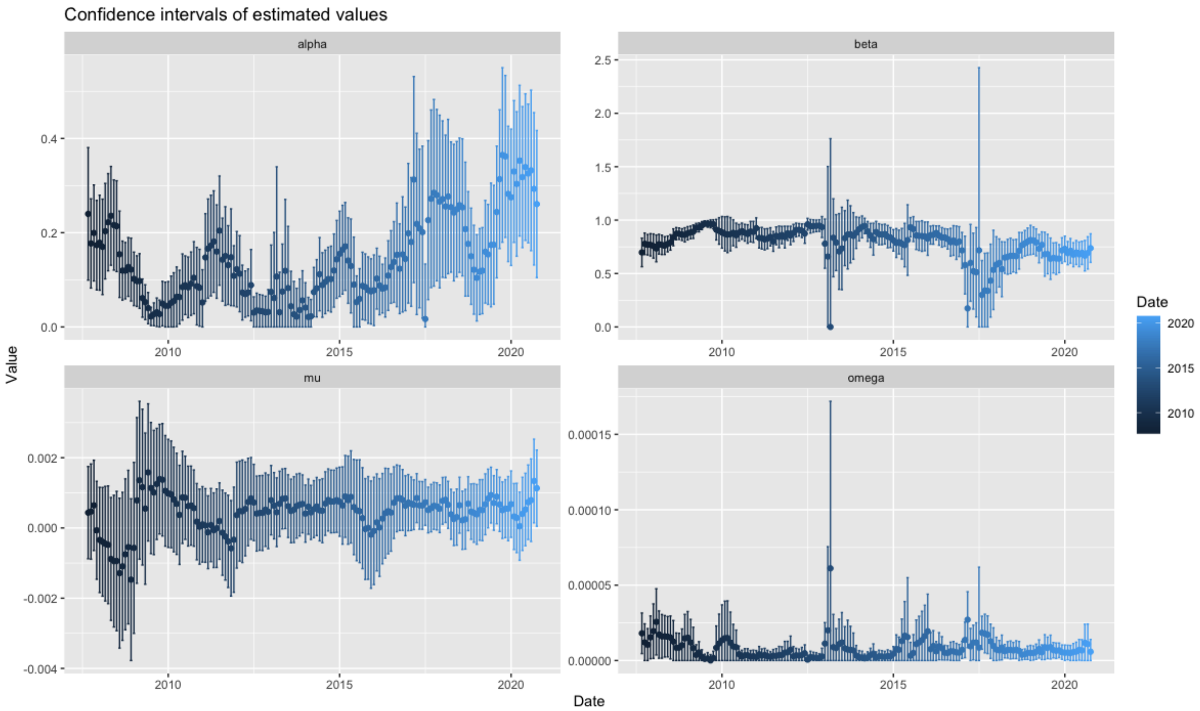

Figure 4.2: Error bar plots of the estimated parameter values and the corresponding 95% confidence intervals. They present the estimated models based on 158 different subsets of the AEX index dataset. All subsets consist of 14 months of stock price data and thus have the same size.

In Figure 4.1 we saw only positive estimations for the parameter μ but now the periods occurring during heavy crisis times have a negative estimated value. The estimations for ω appear to have stayed fairly constant over time with the exception of one large outlier in 2013, but we have not been able to track down what happened on this day in 2013 that could have caused it.

For β we find that the estimated value does fluctuate in time. But comparing the last BC periods with the PC periods we do not find any disjoint intervals and therefore cannot reject the assumption that the parameter values are equal. The same holds for α. Even though α clearly shows to have increased in time, this change does not seem to have been a consequence of the COVID-19 crisis.

4.4 Parametric Bootstrap

We have estimated the parameters μ, ω, α and β, but how accurate are these estimations? To test the accuracy we use *parametric bootstrapping*. This method belongs to the larger bootstrapping family that was introduced by Bradley Efron (1979) as a computer-based method for estimating the standard error. It is seen as a substitute for testing statistical inference. Bootstrapping assigns measures of accuracy, for instance confidence intervals, to sample estimates [10]. Bootstrapping depends on the concept of a *bootstrap sample*, which in the parametric bootstrap case is carried out by a parametric model such as the GARCH(1,1). A simulation study by Adriana Cornea (2014) shows that with heavy-tailed distributions, the quality of inference based on a parametric bootstrap is sufficient.

Let θ denote the vector of parameters in the model and $\hat{\theta}$ the pseudo-true values obtained from the QLM estimation. We are interested in the accuracy of $\hat{\theta}$, and approximate $\theta - \hat{\theta}$ by $\hat{\theta}^* - \hat{\theta}$. Here $\hat{\theta}^*$ denotes the estimated parameters based on the simulated bootstrap sample. Now let X denote the observations from the available data and P_{θ_0} the set of actual values for μ, ω, α and β, then:

$$X \sim P_{\theta_0}, \quad (P_\theta)_{\theta \in \Theta}.$$

We test the null hypothesis:

$$H_0 : P_{\theta_0} \stackrel{?}{=} P_\theta \text{ for some } \theta \in \Theta,$$

and thus apply parametric bootstrapping as follows:

$$X \xrightarrow[(QLME)]{\text{Estimation}} \hat{\theta} \xrightarrow[GARCH(1,1)]{\text{Simulation}} X^* \xrightarrow[(QLME)]{\text{Estimation}} \hat{\theta}^*$$

The series X^* is called the bootstrap sample and is created by simulating a sample of returns by the GARCH(1,1) process using the mean value of $\hat{\mu}_i, \hat{\omega}_i, \hat{\alpha}_i$ and $\hat{\beta}_i$ for $i = 1, \ldots 158$. The bootstrap sample X^* has the same size as X, the original observations. As a final step we estimate new parameters with the bootstrap sample X^*.

The code corresponding to this process can be found in the Appendix (A.2.6) where we followed the following three steps:

Step 1: Estimate the parameters for all 158 subsets of the AEX index dataset as explained in Section 4.3. Next, for each parameter take the mean value of the pseudo-true values.

$$(r_t) \xrightarrow[(QLME)]{\text{Estimation}} \begin{pmatrix} \hat{\mu}_1, & \ldots, & \hat{\mu}_{158} \\ \hat{\omega}_1, & \ldots, & \hat{\omega}_{158} \\ \hat{\alpha}_1, & \ldots, & \hat{\alpha}_{158} \\ \hat{\beta}_1, & \ldots, & \hat{\beta}_{158} \end{pmatrix} \begin{matrix} \longrightarrow & \bar{\mu} \\ \longrightarrow & \bar{\omega} \\ \longrightarrow & \bar{\alpha} \\ \longrightarrow & \bar{\beta} \end{matrix}$$

Step 2: Let the mean values $\bar{\mu}, \bar{\omega}, \bar{\alpha}$ and $\bar{\beta}$ be the constant coefficients in the GARCH(1,1) process and simulate a sample (r_t^*) that has the same size as (r_t) with the following model:

$$\begin{cases} r_t^* = \bar{\mu} + \epsilon_t^* \\ \epsilon_t^* = \sigma_t^* \eta_t^*, \quad \eta_t^* \sim \mathcal{N}(0,1) \\ {\sigma_t^*}^2 = \bar{\omega} + \bar{\alpha} {\epsilon_{t-1}^*}^2 + \bar{\beta} {\sigma_{t-1}^*}^2 \end{cases}$$

Step 3: Lastly, estimate the parameter values again to obtain $\hat{\theta}^*$:

$$(r_t^*) \xrightarrow[(QLME)]{\text{Estimation}} \begin{pmatrix} \hat{\mu}_1^*, & \ldots, & \hat{\mu}_{158}^* \\ \hat{\omega}_1^*, & \ldots, & \hat{\omega}_{158}^* \\ \hat{\alpha}_1^*, & \ldots, & \hat{\alpha}_{158}^* \\ \hat{\beta}_1^*, & \ldots, & \hat{\beta}_{158}^* \end{pmatrix}$$

By applying this method we can check whether the assumption of constant coefficients is satisfied. If the estimated values of the parametric bootstrap are coherent with the results in Section 4.3, then we cannot reject the null assumption of constant coefficients. This implies that the model has stayed more or less the same over time and thus no significant changes due to COVID-19 have occurred.

The results of the third step are portrayed in the same way as in Figure 4.2 with error bar plots with the 95% confidence intervals for each parameter. Comparing the parametric bootstrap results with the original estimation in Figure 4.2 we find that there in fact are differences. The parametric bootstrap sample does not seem to portray reality, which might be due to the assumption of constant parameters.

We are curious to see if the PC estimated parameters portray the data well and therefore apply parametric bootstrapping once again. Now instead of taking the mean of all pseudo-true values we let the constant coefficient be the PC estimated parameters $\hat{\mu}_{158}$, $\hat{\omega}_{158}$, $\hat{\alpha}_{158}$ and $\hat{\beta}_{158}$. The results can be found in the Appendix (A.2.6). Similarly to the first bootstrap process, these results are also not coherent with the majority of the original estimation. However, as expected they indeed give a good representation of the COVID crisis period.

This lets us believe that the constant parameter assumption is incorrect and that there have been changes in the model. However, note that these simulations are done with $\eta_t \sim \mathcal{N}(0,1)$ and because of its randomness, every simulation will give a different result. Therefore we execute several simulations to test if our previous findings are still satisfied. We run 100 simulations and estimate parameters on the same five periods discussed in Section 4.2. We count how many confidence intervals of the BC periods are disjoint with the PC period confidence interval. How we executed this process can be found in the Appendix (A.2.7). From these simulations we see that it is rare to have disjoint intervals, but this is not reasonable with reality. We observed that the intervals of Begin 2008 Crash and PC are disjoint in practice in Figure 4.1. The simulations do not reflect the usual behaviour of the model and this justifies our thought of rejecting the constant coefficient assumption.

4.5 Other European Markets

Obviously the AEX marketplace is only one example and might not represent all European marketplaces. For this reason we did the same investigation for two other big European market indices, namely the Cotation Assistée en Continu (CAC 40) in France and the Deutscher Aktien (DAX) in Germany.

4.5.1 Cotation Assistée en Continu (CAC 40)

We first have a look at the CAC 40 index, a stock market index that represents 40 of the largest companies on the Euronext Paris stock exchange market. The opening prices and returns of the CAC 40 index are presented in Figure 4.3. Notice that the trends of the price and return series are very similar to the AEX index data. They show the same periods of extreme return values corresponding to high volatility periods. There is a difference in the price plots, namely that the values of the CAC 40 index are much larger. However, this difference is not relevant since this is only due to the size of the market and does not effect the dynamics.

This is one of the reasons why we use the returns of the opening prices for our calculations and these are very similar to AEX index values. In the Appendix (A.3) we illustrate how the CAC

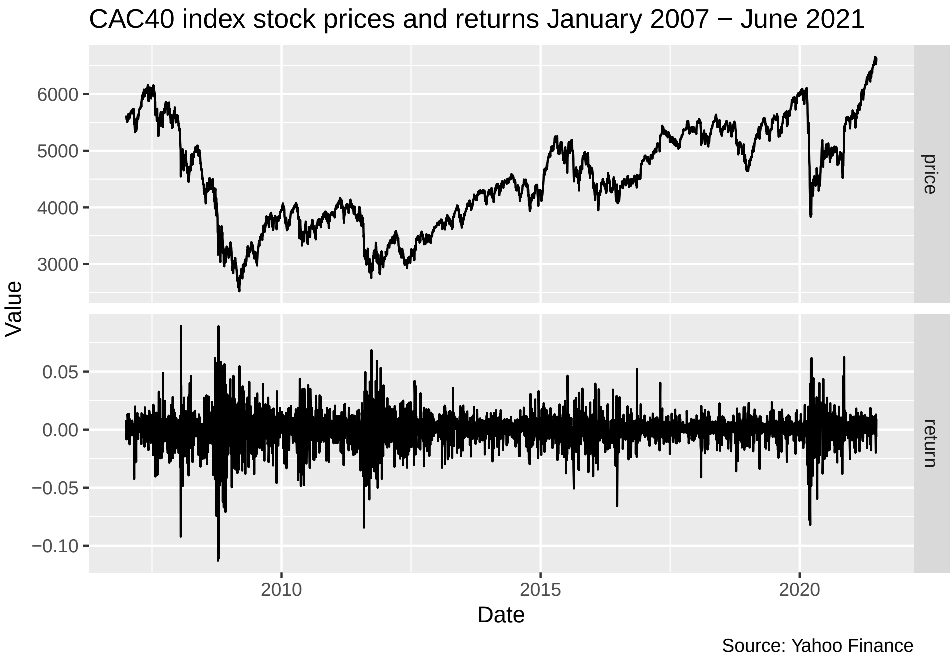

Figure 4.3: Time series of the CAC 40 index daily stock prices and return values from January 2007 to June 2021.

40 index satisfies several key stylised facts such as the autocorrelation properties of the returns and squared returns and the fat-tailed (leptokurtic) distribution.

The CAC 40 index estimations are very similar to those of the AEX index and can be found in the Appendix (A.3). COVID-19 does not seem to have caused noticeable changes in the structure of the model. Moreover, we again see that the model has changed since 2007 as there is strong evidence to reject the assumption of constant parameters, but this change seems due to the Global Financial Crisis in 2008.

4.5.2 Deutscher Aktien (DAX)

The second European index we investigate is the German DAX, consisting of the 30 of the largest companies trading on the Frankfurt Stock Exchange. Figure 4.4 presents the opening prices and returns of the DAX index. The trends of the prices and return values are fairly similar to those of the AEX and CAC 40 data. The DAX price values are even higher than the CAC 40 but this is again irrelevant for the dynamics of the market. Nevertheless the price values of the DAX index appear to have increased more since the Global Financial Crisis in 2008 than the AEX and CAC 40 index prices. The recovery period from this crisis is much shorter than for the other markets. The index prices have greatly surpassed the peak reached before the 2008 crash. In the Appendix (A.4) the autocorrelation and the fat-tailed properties of the DAX data are illustrated. Other results can be found in A.4 as well.

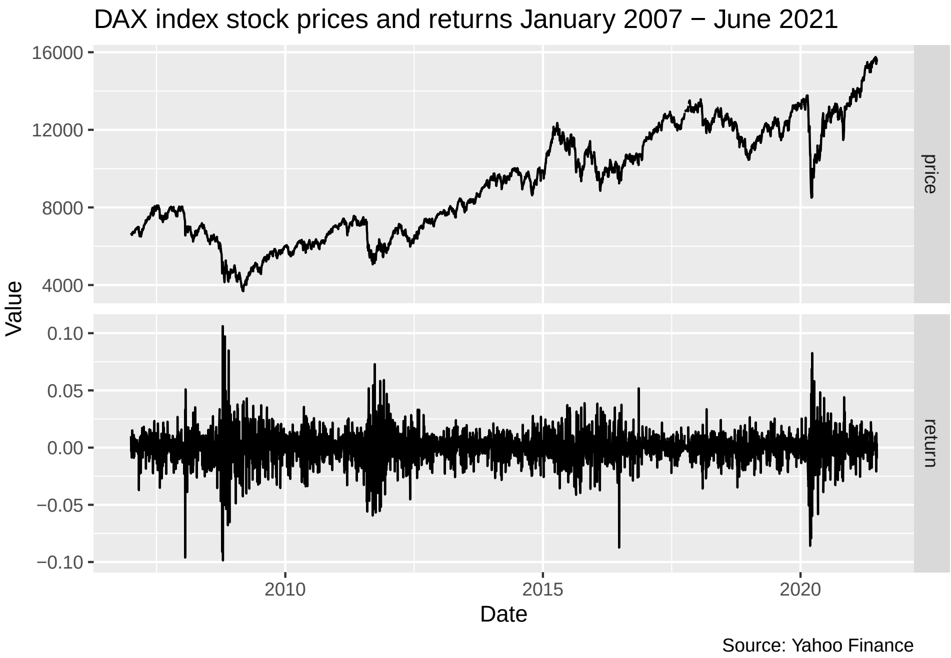

Figure 4.4: Time series of the DAX index daily stock prices and return values from January 2007 to June 2021.

The conclusions based on the CAC 40 index data were very similar to those for the AEX index, but the DAX index on the other hand presents interesting alternative findings. In Figure 4.5 we see that for α all of the confidence intervals of the BC periods are disjoint with the PC period. These results are different from the other indices since now we find disjointness with the PC interval for *all* the BC periods, and not only for the Begin 2008 Crash period. The fact that the confidence intervals for End 2015 Crash and PC are disjoint is specifically interesting. This constitutes strong evidence that the models fitted on these periods are different. If the models are different, this implies that there has been a change in the structure of the model somewhere between 2016 and 2021. We are guessing that this change is caused by COVID-19, however, there is no way to conclude this for sure. A surprising finding is that the Global Financial Crisis does not seem to have changed the model significantly, in contrast to the conclusions for the CAC 40 and AEX indices. We also performed the *Second Estimation* process from Section 4.3. These results are presented in Figure 4.6 and are coherent with our conclusion.

To test the accuracy of these findings we applied parametric bootstrapping and computed the number of disjoint intervals we experienced for the bootstrap sample estimation. In Figure 4.1 we observed that for the original data the BC confidence intervals are all disjoint with the PC confidence interval. But after running 100 simulations with constant parameters we see that it is rare to have disjoint intervals which is not reasonable in practice. This does not reflect the usual behaviour of the model and justifies our thought of rejecting the constant coefficient assumption. We can state with 95% confidence that the model has changed in time. We cannot prove what model is right for the DAX index data but we can reject the wrong ones, and a model with constant parameters is wrong.

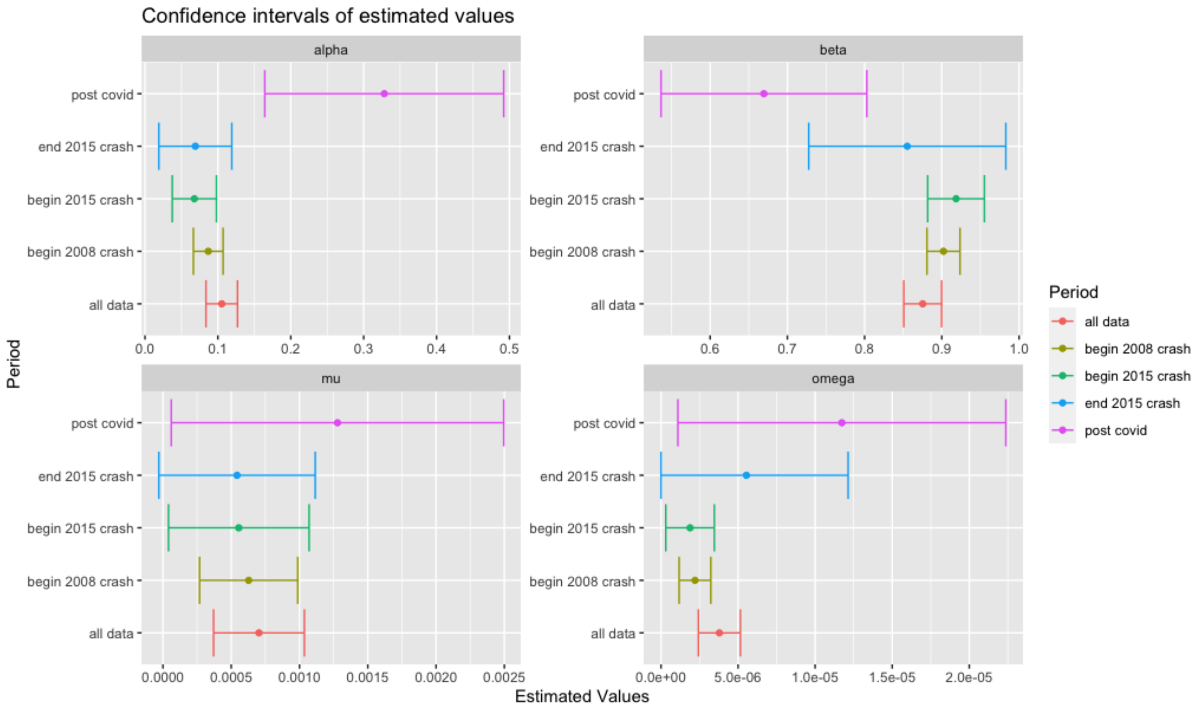

Figure 4.5: Estimating the GARCH(1,1) model using the DAX index data. The estimated parameter values and the corresponding 95% confidence intervals for five different periods (three BC periods, the PC period and the entire dataset) are presented with error bar plots.

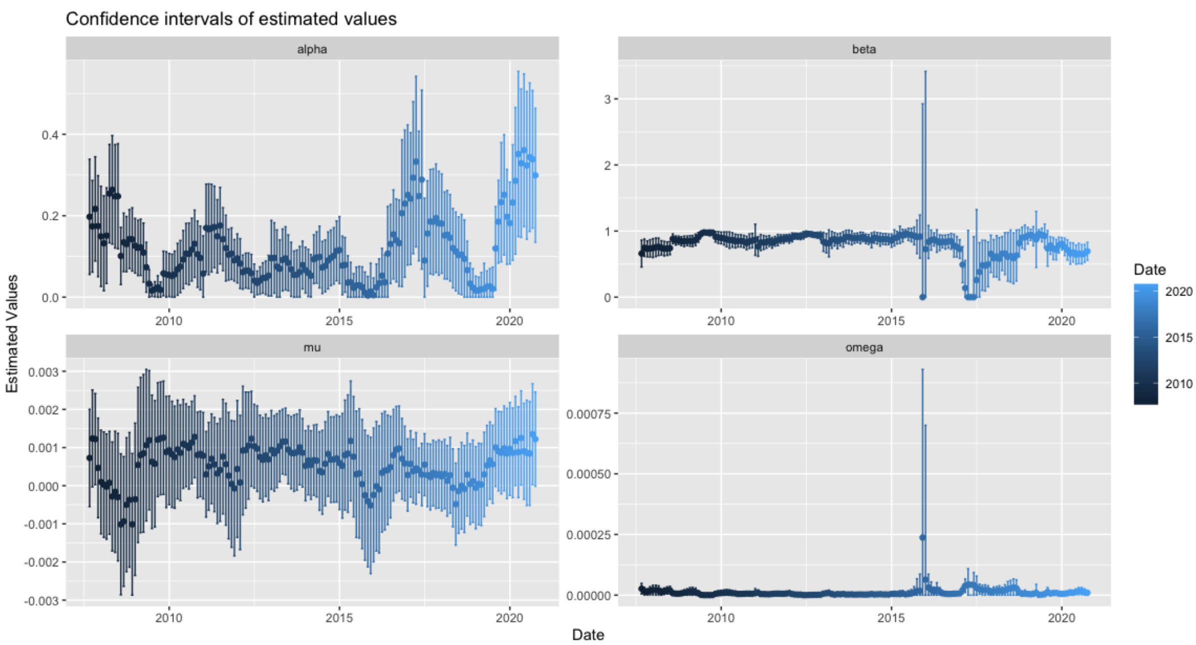

Figure 4.6: Error bar plots of the estimated parameter values and the corresponding 95% confidence intervals. They present the estimated models based on 158 different subsets of the DAX index dataset. All subsets consist of 14 months of stock price data and thus have the same size.

5 Conclusion

The impact of COVID-19 on the dynamics of the stock market does not appear to be the same for all European markets. We modeled three different stock indices, the Dutch Amsterdam Exchange (AEX), the French Cotation Assistée Continu (CAC 40) and the German Deutscher Aktien (DAX). The conclusions for the CAC 40 and AEX were very similar, whereas the DAX on the other hand presented interesting alternative findings.

Of the three market indices, only the DAX showed significant differences between the estimated models fitted on data from before and during the COVID-19 crisis. This implies that COVID-19 has indeed changed the dynamics of this market. This justifies our decision to compare different markets, because there is no reason for them to actually be similar. For example, the German economy is widely based on export so the belief that the dynamics of the stock market have changed because of COVID-19 is quite reasonable. The dynamics of the market have changed in the sense that, when assuming that the stock returns satisfy a GARCH(1,1) process, the constant parameters of that process are different when modeling the periods before and during COVID-19.

On the contrary, for the CAC 40 and the AEX indices there do not seem to have been significant changes in the models due to the COVID-19 crisis. This conclusion is based on the result that the models fitted on the period before and during the crisis are not seemingly different. This affirms the ability of GARCH models to model both high and low periods of volatility, since the periods before and during the COVID-19 crisis show low volatility and high volatility, respectively. Moreover, this leaves us to believe that the COVID-19 crisis in fact just resembles a period of high volatility for the Dutch and French market.

What is surprising is that the model fitted on the period during COVID-19 looks more like the model fitted on low volatility data than the model fitted on data from a period which includes the Global Financial Crisis in 2008. The Global Financial Crisis clearly appears to have had a greater impact for the Dutch and French markets, showing significant changes in the model, in comparison to the COVID-19 crisis. Moreover, the markets seem to have recovered from the COVID-19 crisis much faster than they recovered from the Global Financial Crisis. It took less than a year for the prices to have surpassed the peak value reached before the COVID crisis, whereas after the Global Financial Crisis this took almost 10 years for both the AEX and CAC markets.

On the other hand, the Global Financial Crisis does not seem to have changed the DAX index model significantly. The recovery period from this crisis is also much smaller compared to the recovery period experienced in the other markets. This is another justification for the assumption that markets experience impacts of stock market crises differently.

6 Discussion

For our research we have modeled several European stock markets before and during the COVID-19 crisis. We used the GARCH(1,1) model, a simple case of the family of estimations models called GARCH (General AutoRegressive Conditinally Heteroscedastic) models. GARCH models are known for their ability to model the time-varying volatility property of stock returns. The classical GARCH model we used is popular because of its simplicity and ability to capture most of the empirical properties of financial time series but there are, however, a few restrictions.

In Section 3.5 we discussed how a few stylised facts are not fully captured by the classical GARCH models. The symmetric form of a classical GARCH model is the most obvious one. This does not portray the empirical properties defined as the gain/loss asymmetry and the leverage effect. Many extensions of the classical model have been introduced to account for this asymmetry property such as the Exponential GARCH (EGARCH), Threshold GARCH (TGARCH) and the Asymmetric Power GARCH (APARCH). These and several other extension models are discussed in Chapter 4 of the book on GARCH models by Francq & Zakoian [14].

Another limitation regards the heavy-tailed distribution property of financial time series. We realised that the kurtosis of the estimated model was too small in comparison to the kurtosis of the actual observations, showing that the tail behavior of GARCH models is still too short. We could have let η follow a student-t distribution to improve the kurtosis coefficient, but studies have shown that even then the tails are still too short. Additionally, during our research we experienced that the algorithm in R was unable to compute the standard error for a few estimated parameters because of numerical convergence problems. This either suggests that the classical GARCH model is not a good fit or that we should indeed replace $\mathcal{N}(0,1)$, the distribution of the innovation η, by a more appropriate distribution.

Moreover, besides using univariate models we could have also considered a multivariate extension of the GARCH models. We have fitted three one-dimensional models but we could have also used a three-dimensional model where there is interaction between the different stock markets. Chapter 10 of the book by Francq & Zakoian discusses several multivariate GARCH models.

"All models are wrong, but some are useful" (George Box, 1976).

No model can perfectly portray every empirical property of financial time series and, even though the GARCH model had some limitations, it has still provided us with useful findings. However, for further research it would be interesting to try several of the more complicated univariate or multivariate extensions of the GARCH models or let η follow a student-t distribution to possibly improve the model. Lastly, the conclusion that European markets experience impact of stock market crises differently piqued our curiosity. The COVID-19 crash seems to have changed the dynamics of the German Deutsche Aktien (DAX) index but for the Dutch Amsterdam Exchange (AEX) index and the French Cotation Assistée en Continu (CAC 40) index these changes were insignificant. It would be interesting to further investigate the differences between these market indices to learn more about why the DAX index was influenced more than the others.

A Appendix

A.1 Estimating GARCH Models in R

A.1.1 Importing the data

We imported the AEX index data from Yahoo! Finance through the 'getSymbols' function from the 'quantmod' package. We calculated the simple returns with the opening price values.

```
getSymbols(c("^AEX"), src = "yahoo")
```

```
## [1] "^AEX"
```

```
AEX_Open = as.vector(AEX$AEX.Open)
returns = na.exclude((AEX_Open[-1] - AEX_Open[-length(AEX$AEX.Open)] ) /
  AEX_Open[-length(AEX_Open)])
```

A.1.2 Fitting the Model

This table reports the result of fitting a GARCH(1,1) model with normally distributed innovation for the AEX index data using the 'fGarch' package. The data ranges from January 2007 through May 2021. The opening prices are obtained from Yahoo! Finance. This table presents the Quasi-maximum likelihood parameter estimates with their corresponding standard error values, t-values and p-values. Additionally the examination of the standardised residuals for serial correlation as well as several goodness-of-fit measures are reported. For example the log-likelihood value and the different information criterion values: AIC, BIC, SIC and HQIC.

```
fit_test = garchFit(~ garch(1,1), data = returns, trace = FALSE)
```

```
summary(fit_test)
```

```
##
## Title:
##  GARCH Modelling
##
## Call:
##  garchFit(formula = ~garch(1, 1), data = returns, trace = FALSE)
##
## Mean and Variance Equation:
##  data ~ garch(1, 1)
## <environment: 0x7f8f31a21f50>
##  [data = returns]
```

```
##
## Conditional Distribution:
##  norm
##
## Coefficient(s):
##         mu       omega      alpha1       beta1
## 5.8729e-04  2.6946e-06  1.3096e-01  8.5790e-01
##
## Std. Errors:
##  based on Hessian
##
## Error Analysis:
##         Estimate  Std. Error  t value Pr(>|t|)
## mu     5.873e-04   1.492e-04    3.937 8.27e-05 ***
## omega  2.695e-06   4.728e-07    5.699 1.21e-08 ***
## alpha1 1.310e-01   1.244e-02   10.532  < 2e-16 ***
## beta1  8.579e-01   1.209e-02   70.982  < 2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Log Likelihood:
##  11469.47    normalized:  3.104891
##
## Description:
##  Fri Jun 18 12:22:38 2021 by user:
##
##
## Standardised Residuals Tests:
##                                 Statistic p-Value
##  Jarque-Bera Test   R    Chi^2  510.0091  0
##  Shapiro-Wilk Test  R    W      0.9800977 0
##  Ljung-Box Test     R    Q(10)  12.05229  0.2815723
##  Ljung-Box Test     R    Q(15)  15.7447   0.3992189
##  Ljung-Box Test     R    Q(20)  26.61889  0.1463307
##  Ljung-Box Test     R^2  Q(10)  8.778673  0.5532293
##  Ljung-Box Test     R^2  Q(15)  17.55793  0.2866187
##  Ljung-Box Test     R^2  Q(20)  18.54497  0.5515554
##  LM Arch Test       R    TR^2   14.81047  0.2519656
##
## Information Criterion Statistics:
##       AIC       BIC       SIC      HQIC
## -6.207616 -6.200887 -6.207618 -6.205221
```

The estimated model looks as follows:

$$\begin{cases} r_t = 0.0005873 + \epsilon, \quad \epsilon_t = \sigma_t \eta_t, \quad \eta_t \sim \mathcal{N}(0,1) \\ \sigma_t^2 = 0.000002695 + 0.1310\epsilon_{t-1}^2 + 0.8579\sigma_{t-1}^2 \end{cases}$$

A.2 Modeling the AEX Index

A.2.1 Importing the data and R Packages

We imported the AEX index data from Yahoo! Finance with the 'getSymbols' function in the 'quantmod' package. All other packages needed for the rest of the research are listed below. The package 'ggplot2' is used for all plots in the report.

We printed `head(AEX)` to find out the range of the dataset: January 2007 through June 2021. The last value changes every day since new values keep coming in. The AEX index data provides us with six columns of data: opening prices, closing prices, highest price of the day, lowest price of the day, volume and the adjusted value (equal to the closing price value). We use the opening prices to compute the returns which will be used for the model estimation.

```
library(quantmod)
library(fGarch)
library(ggplot2)
library(ggpubr)
library(tidyverse)
library(forecast)
library(moments)
library(lubridate)
```

```
getSymbols(c("^AEX"), src = "yahoo") #source is Yahoo! Finance
```

```
## [1] "^AEX"
```

```
head(AEX) #first values
```

```
##            AEX.Open AEX.High AEX.Low AEX.Close AEX.Volume AEX.Adjusted
## 2007-01-02   497.91   501.39  497.86    501.00   68535200       501.00
## 2007-01-03   501.48   501.80  499.85    501.35   96484400       501.35
## 2007-01-04   498.20   500.96  497.93    500.33  117535300       500.33
## 2007-01-05   498.39   499.21  494.04    494.05  125123300       494.05
## 2007-01-08   495.37   495.88  491.42    492.81  113168100       492.81
## 2007-01-09   495.04   497.97  494.36    495.87  139352700       495.87
```

```
AEX_Open = as.vector(AEX$AEX.Open)
returns = (AEX_Open[-1] - AEX_Open[-length(AEX$AEX.Open)] ) /
  AEX_Open[-length(AEX_Open)]
```

A.2.2 Kurtosis

The kurtosis coefficient of the returns, defined in Section 3.3.3, is calculated by the function `kurtosis` from the 'moments' package. The kurtosis coefficient of the AEX index returns equals

12.16677. This high value is expected since we know financial time series often have a leptokurtic distribution with fat tails and thus a high kurtosis coefficient. Note that the kurtosis coefficient of a normal distribution with zero mean and unit variance equals 3. Figure 2.5 illustrates the fatness of the tails for the AEX index data.

```
kurtosis_coef = kurtosis(na.exclude(returns))
kurtosis_coef
```

```
## [1] 12.15176
```

A.2.3 First Estimation

As discussed in Section 4.1 we divided the dataset into five periods:

$$\begin{array}{llrl}
\text{Period 1}: & [12/01/2016 - 01/30/2020] & \text{End 2015 Crash} & (\text{BC1}) \\
\text{Period 2}: & [10/10/2014 - 01/30/2020] & \text{Begin 2015 Crash} & (\text{BC2}) \\
\text{Period 3}: & [01/01/2008 - 01/30/2020] & \text{Begin 2008 Crash} & (\text{BC3}) \\
\text{Period 4}: & [01/31/2020 - 04/30/2021] & \text{Post COVID} & (\text{PC}) \\
\text{Period 5}: & [01/01/2007 - 04/30/2021] & \text{All data} & (\text{All})
\end{array}$$

We estimated the GARCH(1,1) model on all five periods with the `garchFit` function from the 'fGarch' package. For the summary output of `garchFit` we refer to the Appendix (A.1.2). The values of interest are the parameter estimates $\mu, \omega, \alpha, \beta$ and the corresponding standard error values.

```
indexstartBC.1 = which(index(AEX)>=as.Date("2016-12-01"))[1]
indexstartBC.2 = which(index(AEX)>=as.Date("2014-10-10"))[1]
indexstartBC.3 = which(index(AEX)>=as.Date("2008-01-01"))[1]
indexendBC = which(index(AEX)>=as.Date("2020-01-30"))[1]
indexstartPC = which(index(AEX)>=as.Date("2020-01-31"))[1]
indexendPC = which(index(AEX)>=as.Date("2021-04-30"))[1]

AEX_OpenBC.1 = AEX_Open[indexstartBC.1:indexendBC]
AEX_OpenBC.2 = AEX_Open[indexstartBC.2:indexendBC]
AEX_OpenBC.3 = AEX_Open[indexstartBC.3:indexendBC]
AEX_OpenPC = AEX_Open[indexstartPC:indexendPC]

#Period  1
returnsBC.1 = (AEX_OpenBC.1[-1] - AEX_OpenBC.1[-length(AEX_OpenBC.1)] ) /
  AEX_OpenBC.1[-length(AEX_OpenBC.1)]
fitBC.1 = garchFit(data=na.exclude(returnsBC.1), trace = F)

#Period 2
returnsBC.2 = (AEX_OpenBC.2[-1] - AEX_OpenBC.2[-length(AEX_OpenBC.2)] ) /
  AEX_OpenBC.2[-length(AEX_OpenBC.2)]
```

```
fitBC.2 = garchFit(data=na.exclude(returnsBC.2), trace = F)

#Period 3
returnsBC.3 = (AEX_OpenBC.3[-1] - AEX_OpenBC.3[-length(AEX_OpenBC.3)] ) /
  AEX_OpenBC.3[-length(AEX_OpenBC.3)]
fitBC.3 = garchFit(data = na.exclude(returnsBC.3), trace = F)

#Period 4 : Post Covid
returnsPC = (AEX_OpenPC[-1] - AEX_OpenPC[-length(AEX_OpenPC)] ) /
  AEX_OpenPC[-length(AEX_OpenPC)]
fitPC = garchFit(data = na.exclude(returnsPC), trace = F)

#All data
returnsBC.all= na.exclude(returns[index(AEX)<as.Date("2021-04-30")])
fit.al = garchFit(data=returnsBC.all, trace = F)
```

A.2.4 Comparing the Estimated Parameters

To compare the estimated models of the five periods we put all the necessary information into a data frame. This data frame has the following columns: parameter, situation, estimate, standard error, lower and upper. Situation refers to the different periods. The lower and upper values are the left and right boundary values, respectively, of the 95% confidence intervals of the estimated parameters. All computations are done with the 95% confidence interval, so all results are given with 95% confidence. The boundaries for each parameter are calculated by:

$$lower = estimate - 1.96 \times st.error$$

$$upper = estimate + 1.96 \times st.error$$

where 1.96 is the quantile of the normal distribution corresponding to the 95% confidence interval. We may use this quantile because of the asymptotic normality discussed in Section 3.4.2. We have $\alpha = 1 - 0.95$ and therefore $q_N^{1-\frac{\alpha}{2}} = q_N^{0.975} = 1.96$.

```
par <- c("mu", "omega","alpha","beta")
parameter <- rep(par, each=5)
sit <- c("end 2015 crash", "begin 2015 crash", "begin 2008 crash",
         "all data", "post covid")
situation <- rep(sit, times =4)
estimate <- c(fitBC.1@fit$par[1],fitBC.2@fit$par[1], fitBC.3@fit$par[1],
              fit.al@fit$par[1],fitPC@fit$par[1],
              fitBC.1@fit$par[2],fitBC.2@fit$par[2], fitBC.3@fit$par[2],
              fit.al@fit$par[2],fitPC@fit$par[2],
              fitBC.1@fit$par[3],fitBC.2@fit$par[3], fitBC.3@fit$par[3],
              fit.al@fit$par[3], fitPC@fit$par[3],
              fitBC.1@fit$par[4],fitBC.2@fit$par[4], fitBC.3@fit$par[4],
              fit.al@fit$par[4] ,fitPC@fit$par[4])
st.error <- c(fitBC.1@fit$se.coef[1], fitBC.2@fit$se.coef[1],
              fitBC.3@fit$se.coef[1],fit.al@fit$se.coef[1],fitPC@fit$se.coef[1],
```

```
              fitBC.1@fit$se.coef[2], fitBC.2@fit$se.coef[2],
              fitBC.3@fit$se.coef[2],fit.al@fit$se.coef[2],fitPC@fit$se.coef[2],
              fitBC.1@fit$se.coef[3], fitBC.2@fit$se.coef[3],
              fitBC.3@fit$se.coef[3],fit.al@fit$se.coef[3],fitPC@fit$se.coef[3],
              fitBC.1@fit$se.coef[4], fitBC.2@fit$se.coef[4],
              fitBC.3@fit$se.coef[4],fit.al@fit$se.coef[4],fitPC@fit$se.coef[4])
lower <- estimate - 1.96*st.error
upper <- estimate + 1.96*st.error

#Adding the constraint for omega, alpha and beta: cannot have negative values
for (i in 6:20) {
  lower[i]= max(0, lower[i], na.rm = TRUE)
}
estimates_df<- data.frame(parameter, situation,
                                estimate, st.error, lower, upper)
estimates_df
```

```
##    parameter        situation     estimate     st.error        lower        upper
## 1         mu   end 2015 crash 7.041256e-04 2.239327e-04 2.652175e-04 1.143034e-03
## 2         mu begin 2015 crash 6.181023e-04 2.105592e-04 2.054063e-04 1.030798e-03
## 3         mu begin 2008 crash 5.095245e-04 1.596174e-04 1.966744e-04 8.223746e-04
## 4         mu         all data 5.762947e-04 1.501491e-04 2.820026e-04 8.705869e-04
## 5         mu       post covid 1.139188e-03 5.494966e-04 6.217491e-05 2.216202e-03
## 6      omega   end 2015 crash 6.876278e-06 1.942107e-06 3.069748e-06 1.068281e-05
## 7      omega begin 2015 crash 2.835913e-06 7.752309e-07 1.316461e-06 4.355365e-06
## 8      omega begin 2008 crash 1.797922e-06 4.074060e-07 9.994065e-07 2.596438e-06
## 9      omega         all data 2.594074e-06 4.644445e-07 1.683762e-06 3.504385e-06
## 10     omega       post covid 7.200222e-06 4.829794e-06 0.000000e+00 1.666662e-05
## 11     alpha   end 2015 crash 1.792714e-01 3.930481e-02 1.022340e-01 2.563088e-01
## 12     alpha begin 2015 crash 1.303144e-01 2.075707e-02 8.963049e-02 1.709982e-01
## 13     alpha begin 2008 crash 1.088985e-01 1.220025e-02 8.498598e-02 1.328110e-01
## 14     alpha         all data 1.278157e-01 1.213936e-02 1.040225e-01 1.516088e-01
## 15     alpha       post covid 2.809891e-01 7.734594e-02 1.293911e-01 4.325872e-01
## 16      beta   end 2015 crash 6.909641e-01 5.976549e-02 5.738238e-01 8.081045e-01
## 17      beta begin 2015 crash 8.435192e-01 2.289614e-02 7.986427e-01 8.883956e-01
## 18      beta begin 2008 crash 8.830475e-01 1.208715e-02 8.593567e-01 9.067383e-01
## 19      beta         all data 8.615018e-01 1.184857e-02 8.382786e-01 8.847250e-01
## 20      beta       post covid 7.188314e-01 6.929469e-02 5.830138e-01 8.546490e-01
```

Error bar plots of the estimated parameters with their 95% confidence intervals are presented in Section 4.2 in Figure 4.1. We are interested in the cases where the confidence intervals of the different periods are disjoint.

A.2.5 Second Estimation

We split the timeline into many new periods that all have the same size. Starting in January 2007 and ending in April 2021, we have 158 periods that all contain 14 months of data. Each new period is constructed by shifting the previous period up one month, which results in 158 periods in total. We estimate a GARCH(1,1) model for each of these periods and construct a data frame similar to the one in Section A.2.4. The results are visualised in Figure 4.2 with error bar plots of the estimated parameters and the corresponding 95% confidence intervals.

```
estiMu = rep(NA, 158)
estiOmega = rep(NA, 158)
estiAlpha = rep(NA, 158)
estiBeta = rep(NA, 158)

st.e.Mu = rep(NA, 158)
st.e.Omega = rep(NA, 158)
st.e.Alpha = rep(NA, 158)
st.e.Beta = rep(NA, 158)

lmu = rep(NA, 158)
lomega = rep(NA, 158)
lalpha = rep(NA, 158)
lbeta = rep(NA, 158)

s = as.Date(rep(NA, 158))

for (t in 1:158){
  start= which(index(AEX)>= ymd("20070801") %m+% months(t-7))[1]
  end = which(index(AEX)>= ymd("20070801") %m+% months(t+7))[1]
  AEX_Open.r=AEX_Open[start:end]
  return = (AEX_Open.r[-1] - AEX_Open.r[-length(AEX_Open.r)] ) /
    AEX_Open.r[-length(AEX_Open.r)]
  fit = garchFit(data = na.exclude(return), trace = F)
  p <- c("mu", "omega", "alpha", "beta")
  parameters <- rep(p, each = 158)
  s[t] <- ymd(ymd("20070801") %m+% months(t))
  situations <- rep(s, times=4)
  estiMu[t] <- fit@fit$par[1]
  estiOmega[t] <- fit@fit$par[2]
  estiAlpha[t] <- fit@fit$par[3]
  estiBeta[t] <- fit@fit$par[4]
  estimates <- c(estiMu, estiOmega, estiAlpha, estiBeta)
  st.e.Mu[t] <- fit@fit$se.coef[1]
  st.e.Omega[t] <- fit@fit$se.coef[2]
  st.e.Alpha[t] <- fit@fit$se.coef[3]
  st.e.Beta[t] <- fit@fit$se.coef[4]
  st.errors <- c(st.e.Mu, st.e.Omega, st.e.Alpha, st.e.Beta)
  lmu[t] <- estiMu[t] - 1.96*st.e.Mu[t]
  lomega[t] <- estiOmega[t] - 1.96*st.e.Omega[t]
  lomega[t] <- max(0, lomega[t], na.rm = TRUE)
  lalpha[t] <- estiAlpha[t] - 1.96*st.e.Alpha[t]
  lalpha[t] <- max(0, lalpha[t], na.rm = TRUE)
  lbeta[t] <- estiBeta[t] - 1.96*st.e.Beta[t]
```

```
  lbeta[t] <- max(0, lbeta[t], na.rm = TRUE)
  lowers= c(lmu, lomega, lalpha, lbeta)
  uppers <- estimates + 1.96*st.errors
  df <- data.frame(parameters, situations, estimates, st.errors, lowers, uppers)
}
```

A.2.6 Parametric Bootstrap

We apply parametric bootstrapping to the estimated models and use the process explained in Section 4.4. We first apply the method where the coefficients in the bootstrap sample are the mean of the pseudo-true values and afterwards do the same procedure but with the PC estimated parameters as the constant coefficients.

Parametric sample with the mean of the pseudo-true values:

Step 1: Estimate the parameters for all 158 subsets of the AEX index dataset as explained in Section 4.3. Next, for each parameter take the mean value of the pseudo-true values:

```
mu_bar = mean(estiMu)
omega_bar= mean(estiOmega)
alpha_bar= mean(estiAlpha)
beta_bar = mean(estiBeta)
```

Step 2: Let the mean values $\bar{\mu}$, $\bar{\omega}$, $\bar{\alpha}$ and $\bar{\beta}$ be the constant coefficients in the GARCH(1,1) process and simulate the bootstrap sample (r_t^*) that has the same size as (r_t) with the following model:

```
Timelength = length(AEX_Open)
Time = 1:Timelength
eta_t = rep(NA, Timelength)
epsilon_t = rep(NA, Timelength)
sigma_t = rep(NA, Timelength)
Price_t = rep(NA, Timelength)
r_t = rep(NA, Timelength)

eta_t[1] = 0
epsilon_t[1] = 0
sigma_t[1] = 0
r_t[1]=0
Price_t[1] = 489

for (t in 2:Timelength){
  eta_t[t] = rnorm(1)
  sigma_t[t] = sqrt( omega_bar + alpha_bar * (epsilon_t[t-1])^2 +
```

```
  beta_bar * (sigma_t[t-1])^2 )
  epsilon_t[t] = eta_t[t] * sigma_t[t]
  Price_t[t] = Price_t[t-1] * exp(epsilon_t[t])
  r_t[t] = mu_bar + epsilon_t[t]
}
#plots of the simulated sample
```

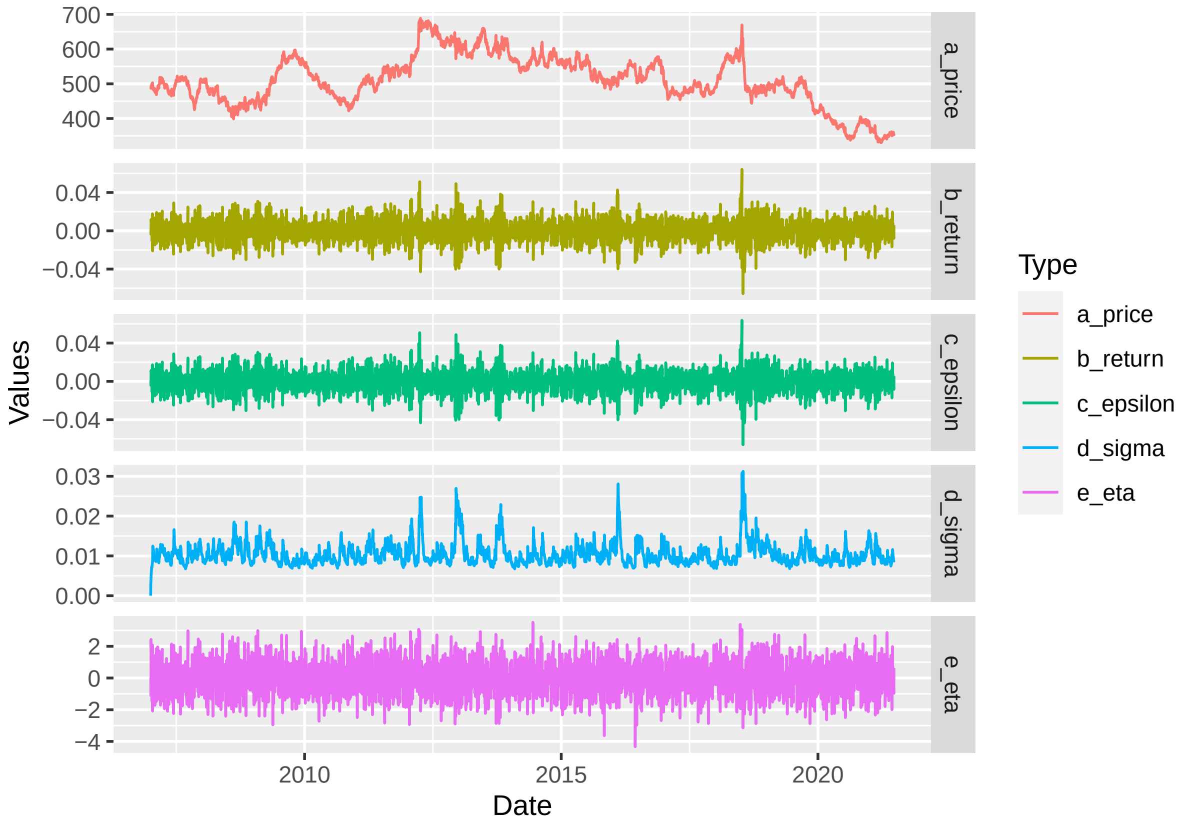

Note that this simulation is not entirely realistic. In the graph corresponding to the price values (a_price) we notice that the price is able to increase just as steeply as decrease. This is because $\eta_t \sim \mathcal{N}(0,1)$ and therefore it returns positive and negative values of the same magnitude. In reality, however, this is not the case. This asymmetry is defined as the *gain/loss asymmetry* stylised fact discussed in Section 2.2.2.

Step 3: Lastly, estimate the parameter values again to obtain $\hat{\theta}^*$:

```
estiMu2 = rep(NA, 158)
estiOmega2 = rep(NA, 158)
estiAlpha2 = rep(NA, 158)
estiBeta2 = rep(NA, 158)

st.e.Mu2 = rep(NA, 158)
st.e.Omega2 = rep(NA, 158)
st.e.Alpha2 = rep(NA, 158)
st.e.Beta2 = rep(NA, 158)
```

```
lmu2 = rep(NA, 158)
lomega2 = rep(NA, 158)
lalpha2 = rep(NA, 158)
lbeta2 = rep(NA, 158)

s = as.Date(rep(NA, 158))
for (t in 1:158){
  start= which(index(AEX)>= ymd("20070801") %m+% months(t-7))[1]
  end = which(index(AEX)>= ymd("20070801") %m+% months(t+7))[1]
  return2 = r_t[start:end]
  fit2 = garchFit(data = na.exclude(return2), trace = F)
  p <- c("mu", "omega", "alpha", "beta")
  parameters2 <- rep(p, each = 158)
  s[t] <- ymd(ymd("20070801") %m+% months(t))
  situations2 <- rep(s, times=4)
  estiMu2[t] <- fit2@fit$par[1]
  estiOmega2[t] <- fit2@fit$par[2]
  estiAlpha2[t] <- fit2@fit$par[3]
  estiBeta2[t] <- fit2@fit$par[4]
  estimates2 <- c(estiMu2, estiOmega2, estiAlpha2, estiBeta2)
  st.e.Mu2[t] <- fit2@fit$se.coef[1]
  st.e.Omega2[t] <- fit2@fit$se.coef[2]
  st.e.Alpha2[t] <- fit2@fit$se.coef[3]
  st.e.Beta2[t] <- fit2@fit$se.coef[4]
  st.errors2 <- c(st.e.Mu2, st.e.Omega2, st.e.Alpha2, st.e.Beta2)
  lmu2[t] <- estiMu2[t] - 1.96*st.e.Mu2[t]
  lomega2[t] <- estiOmega2[t] - 1.96*st.e.Omega2[t]
  lomega2[t] <- max(0, lomega2[t], na.rm = TRUE)
  lalpha2[t] <- estiAlpha2[t] - 1.96*st.e.Alpha2[t]
  lalpha2[t] <- max(0, lalpha2[t], na.rm = TRUE)
  lbeta2[t] <- estiBeta2[t] - 1.96*st.e.Beta2[t]
  lbeta2[t] <- max(0, lbeta2[t], na.rm = TRUE)
  lowers2= c(lmu2, lomega2, lalpha2, lbeta2)
  uppers2 <- estimates2 + 1.96*st.errors2
  df_2 <- data.frame(parameters2, situations2, estimates2, st.errors2,
     lowers2, uppers2)
}
```

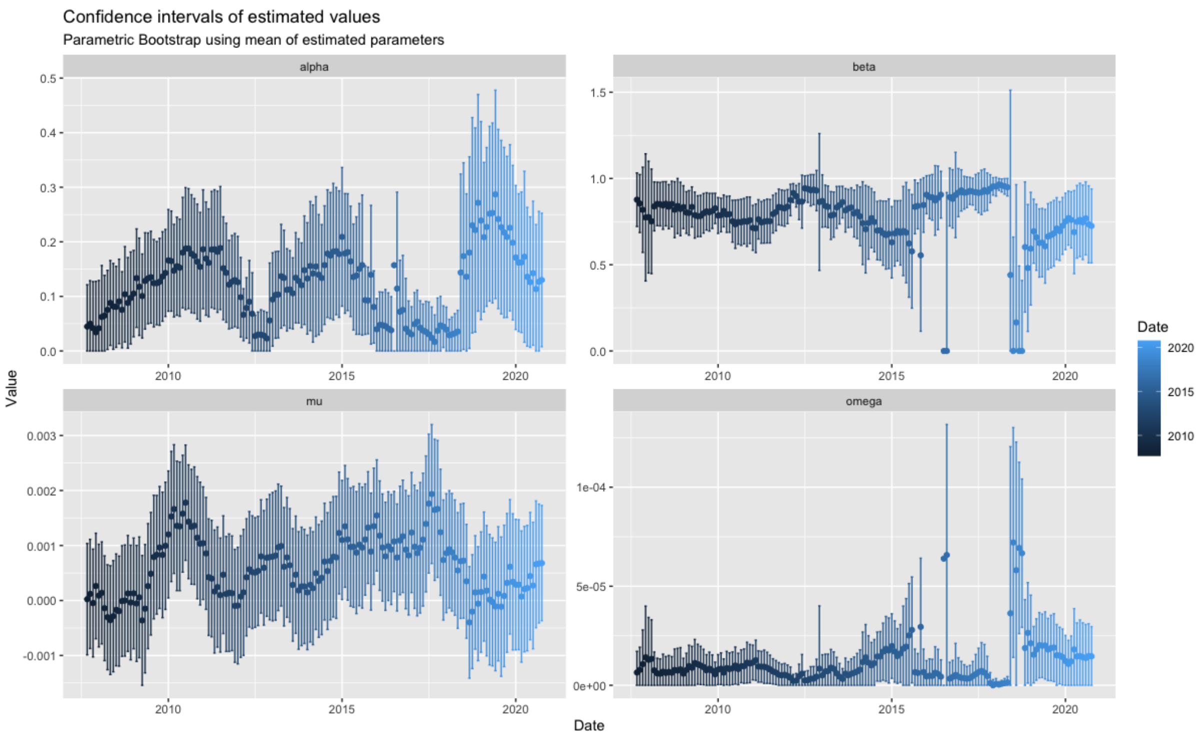

Parametric sample with the PC pseudo-true value:

Step 1: Estimate the parameters for all 158 subsets of the AEX index dataset as explained in Section 4.3. Next, for each parameter take the pseudo-true values of the last estimation:

$$(r_t) \xrightarrow[(QLME)]{\text{Estimation}} \begin{pmatrix} \hat{\mu}_1, & \dots, & \hat{\mu}_{158} \\ \hat{\omega}_1, & \dots, & \hat{\omega}_{158} \\ \hat{\alpha}_1, & \dots, & \hat{\alpha}_{158} \\ \hat{\beta}_1, & \dots, & \hat{\beta}_{158} \end{pmatrix} \begin{matrix} \longrightarrow & \hat{\mu}_{158} \\ \longrightarrow & \hat{\omega}_{158} \\ \longrightarrow & \hat{\alpha}_{158} \\ \longrightarrow & \hat{\beta}_{158} \end{matrix}$$

```
mu_hat= estiMu[length(estiMu)]
omega_hat = estiOmega[length(estiOmega)]
alpha_hat = estiAlpha[length(estiAlpha)]
beta_hat = estiBeta[length(estiBeta)]
```

Step 2: Let the values $\hat{\mu}_{158}$, $\hat{\omega}_{158}$, $\hat{\alpha}_{158}$ and $\hat{\beta}_{158}$ be the constant coefficients in the GARCH(1,1) process and simulate a sample (r_t^*) that has the same size as (r_t) with the following model:

$$\begin{cases} r_t^* = \hat{\mu}_{158} + \epsilon_t^* \\ \epsilon_t^* = \sigma_t^* \eta_t^*, \quad \eta_t^* \sim \mathcal{N}(0,1) \\ {\sigma_t^*}^2 = \hat{\omega}_{158} + \hat{\alpha}_{158} {\epsilon_{t-1}^*}^2 + \hat{\beta}_{158} {\sigma_{t-1}^*}^2 \end{cases}$$

```
Timelength = length(AEX_Open)
Time = 1:Timelength
eta_tn = rep(NA, Timelength)
```

```
epsilon_tn = rep(NA, Timelength)
sigma_tn = rep(NA, Timelength)
Price_tn = rep(NA, Timelength)
r_tn = rep(NA, Timelength)

eta_tn[1] = 0
epsilon_tn[1] = 0
sigma_tn[1] = 0
r_tn[1]=0
Price_tn[1] = 489

for (t in 2:Timelength){
  eta_tn[t] = rnorm(1)
  sigma_tn[t] = sqrt( omega_hat + alpha_hat * (epsilon_tn[t-1])^2 +
     beta_hat * (sigma_tn[t-1])^2 )
  epsilon_tn[t] = eta_tn[t] * sigma_tn[t]
  Price_tn[t] = Price_tn[t-1] * exp(epsilon_tn[t])
  r_tn[t] = mu_hat + epsilon_tn[t]
}
#plots of the simulated sample
```

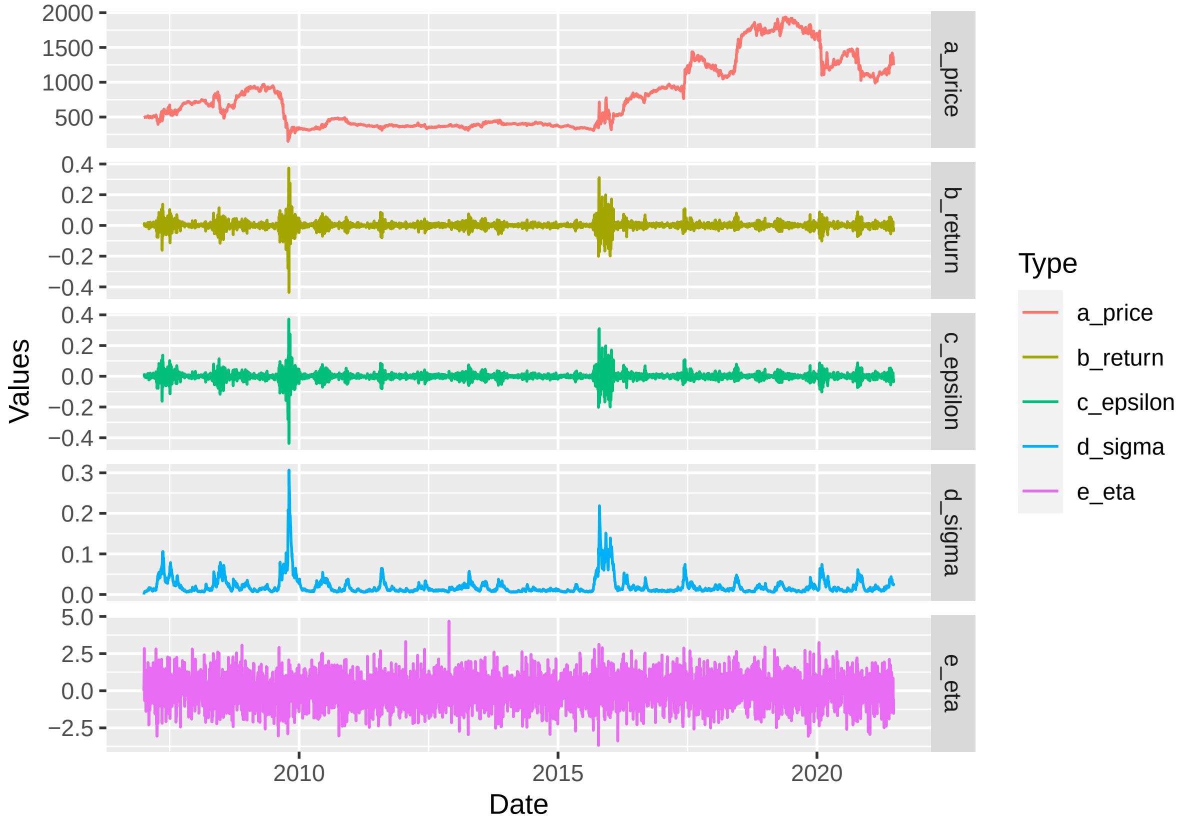

Step 3: Lastly, estimate the parameter values again to obtain $\hat{\theta}^*$:

$$(r_t^*) \xrightarrow[(QLME)]{\text{Estimation}} \begin{pmatrix} \hat{\mu}_1^*, & \dots, & \hat{\mu}_{158}^* \\ \hat{\omega}_1^*, & \dots, & \hat{\omega}_{158}^* \\ \hat{\alpha}_1^*, & \dots, & \hat{\alpha}_{158}^* \\ \hat{\beta}_1^*, & \dots, & \hat{\beta}_{158}^* \end{pmatrix}$$

```
estiMu3 = rep(NA, 158)
estiOmega3 = rep(NA, 158)
estiAlpha3 = rep(NA, 158)
estiBeta3 = rep(NA, 158)

st.e.Mu3 = rep(NA, 158)
st.e.Omega3 = rep(NA, 158)
st.e.Alpha3 = rep(NA, 158)
st.e.Beta3 = rep(NA, 158)

lmu3 = rep(NA, 158)
lomega3 = rep(NA, 158)
lalpha3 = rep(NA, 158)
lbeta3 = rep(NA, 158)

s = as.Date(rep(NA, 158))
for (t in 1:158){
  start= which(index(AEX)>= ymd("20070801") %m+% months(t-7))[1]
  end = which(index(AEX)>= ymd("20070801") %m+% months(t+7))[1]
  return3 = r_tn[start:end]
  fit3 = garchFit(data = na.exclude(return3), trace = F)
  p <- c("mu", "omega", "alpha", "beta")
  parameters3 <- rep(p, each = 158)
  s[t] <- ymd(ymd("20070801") %m+% months(t))
  situations3 <- rep(s, times=4)
  estiMu3[t] <- fit3@fit$par[1]
  estiOmega3[t] <- fit3@fit$par[2]
  estiAlpha3[t] <- fit3@fit$par[3]
  estiBeta3[t] <- fit3@fit$par[4]
  estimates3 <- c(estiMu3, estiOmega3, estiAlpha3, estiBeta3)
  st.e.Mu3[t] <- fit3@fit$se.coef[1]
  st.e.Omega3[t] <- fit3@fit$se.coef[2]
  st.e.Alpha3[t] <- fit3@fit$se.coef[3]
  st.e.Beta3[t] <- fit3@fit$se.coef[4]
  st.errors3 <- c(st.e.Mu3, st.e.Omega3, st.e.Alpha3, st.e.Beta3)
  lmu3[t] <- estiMu3[t] - 1.96*st.e.Mu3[t]
  lomega3[t] <- estiOmega3[t] - 1.96*st.e.Omega3[t]
  lomega3[t] <- max(0, lomega3[t], na.rm = TRUE)
  lalpha3[t] <- estiAlpha3[t] - 1.96*st.e.Alpha3[t]
  lalpha3[t] <- max(0, lalpha3[t], na.rm = TRUE)
  lbeta3[t] <- estiBeta3[t] - 1.96*st.e.Beta3[t]
  lbeta3[t] <- max(0, lbeta3[t], na.rm = TRUE)
  lowers3= c(lmu3, lomega3, lalpha3, lbeta3)
```

```
  uppers3 <- estimates3 + 1.96*st.errors3
  df_3 <- data.frame(parameters3, situations3, estimates3, st.errors3,
     lowers3, uppers3)
}
```

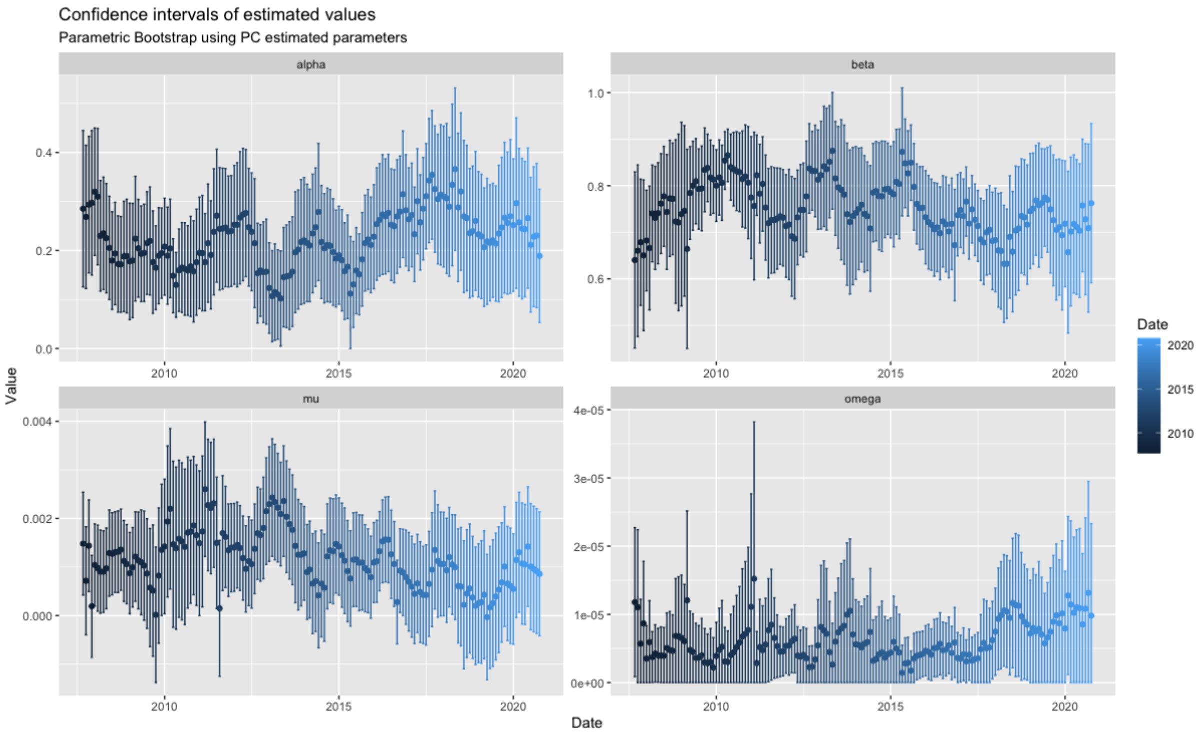

Section 4.4 discusses both the parametric bootstrap results presented here.

A.2.7 Validation with 100 Simulations

The parameter values are now again the mean of the pseudo-true values from the first estimation. We simulate 100 different samples and estimate new models for the five periods from Section 4.1. We calculate how many confidence intervals of the different BC periods are disjoint with the PC confidence interval. The result is presented in a data frame.

```
Timelength = length(AEX_Open)
Time = 1:Timelength
eta_tn = rep(NA, Timelength)
epsilon_tn = rep(NA, Timelength)
sigma_tn = rep(NA, Timelength)
Price_tn = rep(NA, Timelength)
r_tn = rep(NA, Timelength)

eta_tn[1] = 0
epsilon_tn[1] = 0
sigma_tn[1] = 0
```

```
r_tn[1]=0
Price_tn[1] = 489 #first opening price value of the original observations

indexstartBC.1 = which(index(AEX)>=as.Date("2016-12-01"))[1]
indexstartBC.2 = which(index(AEX)>=as.Date("2014-10-10"))[1]
indexstartBC.3 = which(index(AEX)>=as.Date("2008-01-01"))[1]
indexendBC = which(index(AEX)>=as.Date("2020-01-30"))[1]
indexstartPC = which(index(AEX)>=as.Date("2020-01-31"))[1]
indexendPC = which(index(AEX)>=as.Date("2021-04-30"))[1]

disjoint_mu_BC1 = 0
disjoint_mu_BC2 = 0
disjoint_mu_BC3= 0
disjoint_mu_BC.all= 0

disjoint_omega_BC1 = 0
disjoint_omega_BC2 = 0
disjoint_omega_BC3 = 0
disjoint_omega_BC.all = 0

disjoint_alpha_BC1 = 0
disjoint_alpha_BC2 = 0
disjoint_alpha_BC3= 0
disjoint_alpha_BC.all= 0

disjoint_beta_BC1 = 0
disjoint_beta_BC2 = 0
disjoint_beta_BC3 = 0
disjoint_beta_BC.all = 0

count = 0
while (count < 100){
  count = count +1
  for (t in 2:Timelength){
    eta_tn[t] = rnorm(1)
    sigma_tn[t] = sqrt( omega_bar + alpha_bar * (epsilon_tn[t-1])^2 +
     beta_bar * (sigma_tn[t-1])^2 )
    epsilon_tn[t] = eta_tn[t] * sigma_tn[t]
    Price_tn[t] = Price_tn[t-1] * exp(epsilon_tn[t])
    r_tn[t] = mu_bar + epsilon_tn[t]
  }

  sim_BC.1 = r_tn[indexstartBC.1:indexendBC]
  sim_BC.2 = r_tn[indexstartBC.2:indexendBC]
  sim_BC.3 = r_tn[indexstartBC.3:indexendBC]
  sim_all = r_tn[1:indexendPC]
  simPC = r_tn[indexstartPC:indexendPC]

  fitBC.1 = garchFit(data= sim_BC.1, trace = F)
```

```
fitBC.2 = garchFit(data= sim_BC.2, trace = F)
fitBC.3 = garchFit(data = sim_BC.3, trace = F)
fit.all = garchFit(data= sim_all, trace = F)
fitPC = garchFit(data = simPC, trace = F)

par <- c("mu", "omega","alpha","beta")
parameter <- rep(par, each=5)
sit <- c("end 2015 crash", "begin 2015 crash", "begin 2008 crash",
       "all data", "post covid")
situation <- rep(sit, times =4)
estimate <- c(fitBC.1@fit$par[1],fitBC.2@fit$par[1], fitBC.3@fit$par[1],
              fit.al@fit$par[1],fitPC@fit$par[1],
              fitBC.1@fit$par[2],fitBC.2@fit$par[2], fitBC.3@fit$par[2],
              fit.al@fit$par[2],fitPC@fit$par[2],
              fitBC.1@fit$par[3],fitBC.2@fit$par[3], fitBC.3@fit$par[3],
              fit.al@fit$par[3], fitPC@fit$par[3],
              fitBC.1@fit$par[4],fitBC.2@fit$par[4], fitBC.3@fit$par[4],
              fit.al@fit$par[4] ,fitPC@fit$par[4])
st.error <- c(fitBC.1@fit$se.coef[1], fitBC.2@fit$se.coef[1],
              fitBC.3@fit$se.coef[1],fit.al@fit$se.coef[1],fitPC@fit$se.coef[1],
              fitBC.1@fit$se.coef[2], fitBC.2@fit$se.coef[2],
              fitBC.3@fit$se.coef[2],fit.al@fit$se.coef[2],fitPC@fit$se.coef[2],
              fitBC.1@fit$se.coef[3], fitBC.2@fit$se.coef[3],
              fitBC.3@fit$se.coef[3],fit.al@fit$se.coef[3],fitPC@fit$se.coef[3],
              fitBC.1@fit$se.coef[4], fitBC.2@fit$se.coef[4],
              fitBC.3@fit$se.coef[4],fit.al@fit$se.coef[4],fitPC@fit$se.coef[4])
lower <- estimate - 1.96*st.error
for (i in 6:20) {
  lower[i] = max(0, lower[i], na.rm = TRUE)
}
upper <- estimate + 1.96*st.error
sim_estimates_df <- data.frame(parameter, situation, estimate,
                     st.error, lower, upper)
#mu
  #BC1 vs PC
if ((upper[1] < lower[5]) | upper[5] < lower[1] ){
  disjoint_mu_BC1 = disjoint_mu_BC1 + 1
}
  #BC2 vs PC
if ((upper[2] < lower[5]) | upper[5] < lower[2] ){
  disjoint_mu_BC2 = disjoint_mu_BC2 + 1
}
  #BC3 vs PC
if ((upper[3] < lower[5]) | upper[5] < lower[3] ){
  disjoint_mu_BC3 = disjoint_mu_BC3 + 1
}
  #BC.all vs PC
if ((upper[4] < lower[5]) | upper[5] < lower[4] ){
  disjoint_mu_BC.all = disjoint_mu_BC.all + 1
```

```
}

#omega
  #BC1 vs PC
if ((upper[6] < lower[10]) | upper[10] < lower[6] ){
  disjoint_omega_BC1 = disjoint_omega_BC1 + 1
}
  #BC2 vs PC
if ((upper[7] < lower[10]) | upper[10] < lower[7] ){
  disjoint_omega_BC2 = disjoint_omega_BC2 + 1
}
  #BC3 vs PC
if ((upper[8] < lower[10]) | upper[10] < lower[8] ){
  disjoint_omega_BC3 = disjoint_omega_BC3 + 1
}
  #BC.all vs PC
if ((upper[9] < lower[10]) | upper[10] < lower[9] ){
  disjoint_omega_BC.all = disjoint_omega_BC.all + 1
}
#alpha
  #BC1 vs PC
if ((upper[11] < lower[15]) | upper[15] < lower[11] ){
  disjoint_alpha_BC1 = disjoint_alpha_BC1 + 1
}
  #BC2 vs PC
if ((upper[12] < lower[15]) | upper[15] < lower[12] ){
  disjoint_alpha_BC2 = disjoint_alpha_BC2 + 1
}
#BC3 vs PC
if ((upper[13] < lower[15]) | upper[15] < lower[13] ){
  disjoint_alpha_BC3 = disjoint_alpha_BC3 + 1
}
#BC.all vs PC
if ((upper[14] < lower[15]) | upper[15] < lower[14] ){
  disjoint_alpha_BC.all = disjoint_alpha_BC.all + 1
}

#Beta
#BC1 vs PC
if ((upper[16] < lower[20]) | upper[20] < lower[16] ){
  disjoint_beta_BC1 = disjoint_beta_BC1 + 1
}
#BC2 vs PC
if ((upper[17] < lower[20]) | upper[20] < lower[17] ){
  disjoint_beta_BC2 = disjoint_beta_BC2 + 1
}
#BC3 vs PC
if ((upper[18] < lower[20]) | upper[20] < lower[18] ){
  disjoint_beta_BC3 = disjoint_beta_BC3 + 1
```

```
  }
  #BC.all vs PC
  if ((upper[19] < lower[20]) | upper[20] < lower[19] ){
    disjoint_beta_BC.all = disjoint_beta_BC.all + 1
  }
}
```

```
par <- c("mu", "omega", "alpha", "beta")
parameter <- rep(par, each = 4)
per<- c("all data", "BC1", "BC2", "BC3")
period <- rep(per, times=4)
num_disjoint <- c(disjoint_mu_BC.all,disjoint_mu_BC1, disjoint_mu_BC2,
disjoint_mu_BC3, disjoint_omega_BC.all, disjoint_omega_BC1, disjoint_omega_BC2,
disjoint_omega_BC3, disjoint_alpha_BC.all, disjoint_alpha_BC1, disjoint_beta_BC2,
disjoint_alpha_BC3, disjoint_beta_BC.all, disjoint_beta_BC1, disjoint_beta_BC2,
disjoint_beta_BC3)

data.frame(parameter, period, num_disjoint)
```

```
##    parameter   period num_disjoint
## 1         mu all data            0
## 2         mu      BC1            0
## 3         mu      BC2            0
## 4         mu      BC3            0
## 5      omega all data            2
## 6      omega      BC1            1
## 7      omega      BC2            2
## 8      omega      BC3            2
## 9      alpha all data            2
## 10     alpha      BC1            0
## 11     alpha      BC2            6
## 12     alpha      BC3            2
## 13      beta all data            5
## 14      beta      BC1            1
## 15      beta      BC2            6
## 16      beta      BC3            5
```

A.3 Modeling the CAC 40 Index

Importing the data

```
getSymbols( # CAC40
  c("^FCHI"),
  src = "yahoo")
head(FCHI) #first values
FCHI_Open = as.vector(FCHI$FCHI.Open)
returns.FCHI = (FCHI_Open[-1] - FCHI_Open[-length(FCHI_Open)] ) /
  FCHI_Open[-length(FCHI_Open)]
```

```
## [1] "^FCHI"
##            FCHI.Open FCHI.High FCHI.Low FCHI.Close FCHI.Volume FCHI.Adjusted
## 2007-01-02   5575.76   5621.65  5575.63    5617.71    85910000       5617.71
## 2007-01-03   5621.00   5623.67  5596.82    5610.92   118580700       5610.92
## 2007-01-04   5573.73   5585.54  5547.17    5574.56   130465700       5574.56
## 2007-01-05   5552.64   5566.24  5517.35    5517.35   126420500       5517.35
## 2007-01-08   5532.57   5555.67  5509.06    5518.59   115053800       5518.59
## 2007-01-09   5549.01   5563.48  5532.62    5533.03   151688200       5533.03
```

Autocorrelation

Here we show the autocorrelation of the returns (a) and the squared returns (b) of the CAC 40 index. We notice the tiny autocorrelation between the returns while the squared returns on the other hand show larger autocorrelation. These two properties are general characteristics of stock returns and thus support the use of GARCH models.

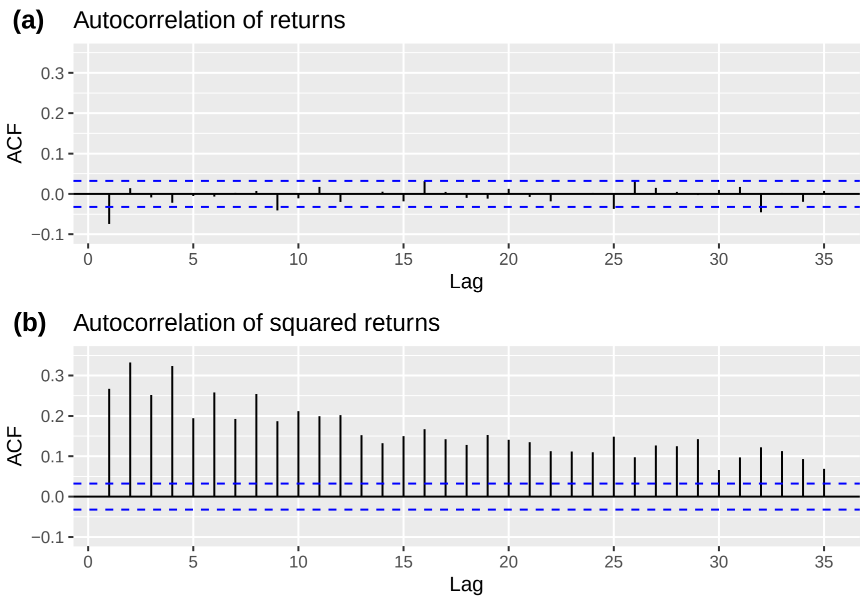

Kurtosis

Next the kurtosis coefficient of the returns is calculated. This value equals 9.66256 which is smaller than for AEX index but still high. This high value is expected since we know financial time series often have a leptokurtic distribution with fat tails and thus a high kurtosis coefficient. The plot below presents the distribution density function of the returns and its deviance from a normal distribution. The logarithmic scale is used in (b) for a clear view on the fatness of the tails.

```
#kurtosis coefficient
kurtosis_coef = kurtosis(na.exclude(returns.FCHI))
kurtosis_coef
```

```
## [1] 9.66256
```

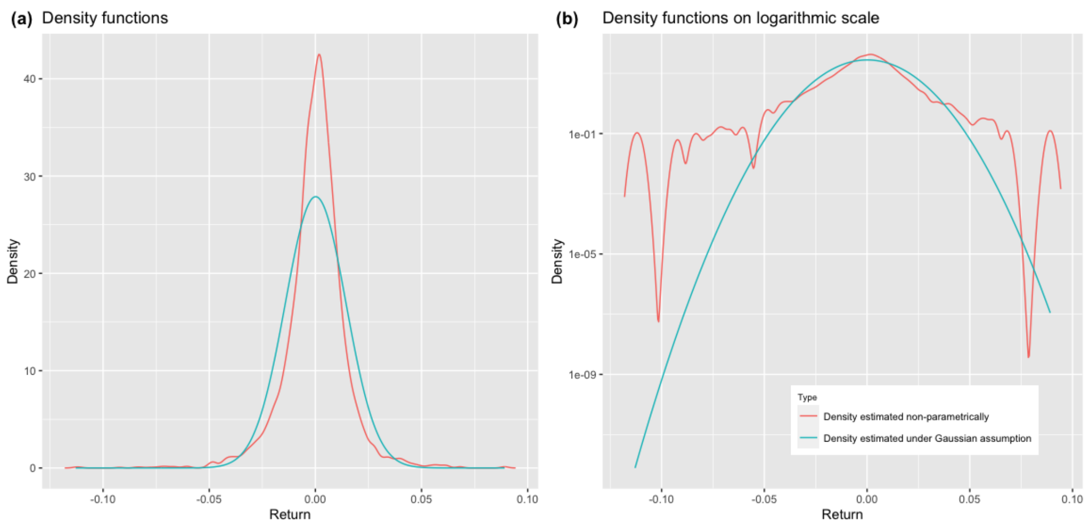

First Estimation

We decided to split the dataset into the same five periods as we did for the AEX index in Section 4.1 since the CAC 40 data shows very similar trends for both the price and return series. Moreover, in the past the whole of Europe has been affected by crises around the same time and therefore it seems reasonable to stick with these periods.

The results are presented in the same way as in Section A.2.4 with error bar plots of the confidence intervals. These results are very similar to the AEX index results in Figure 4.1. Therefore the same remarks discussed in Section 4.2 are relevant. Notice that for the CAC 40 index the disjointness between the confidence intervals Begin 2008 Crash and PC is even stronger for α and β, than for the AEX index.

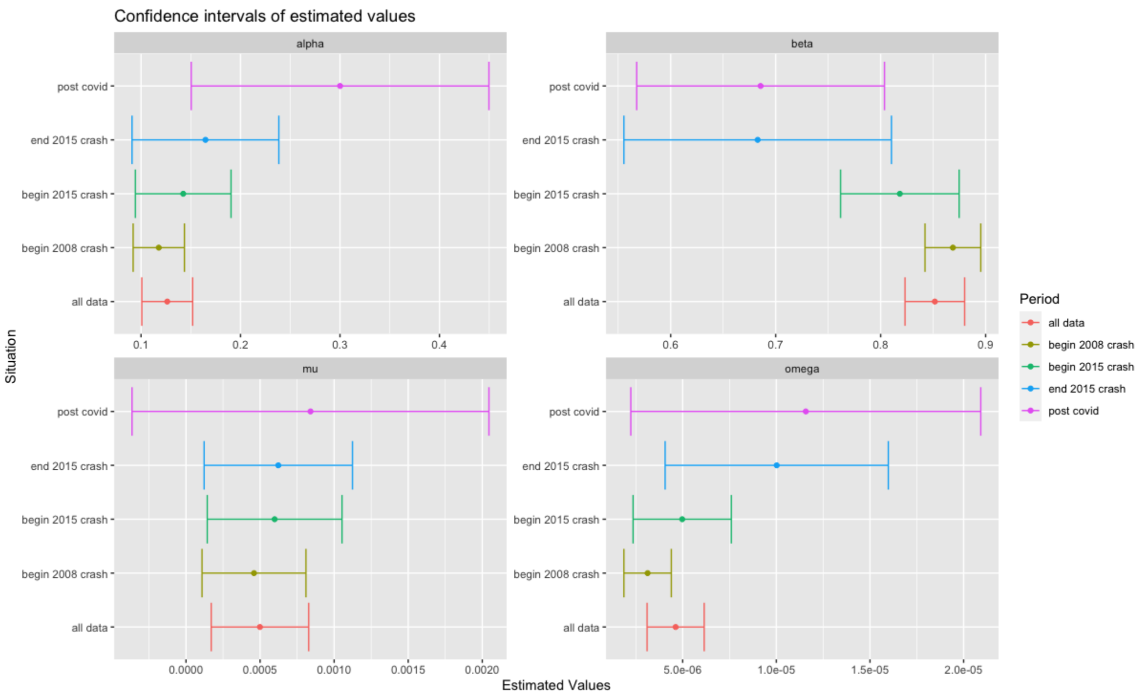

Second Estimation

We apply the same Second Estimation process from Section A.2.5 for the CAC 40 index. Because of numerical convergence problems the algorithm is unable to compute the standard error for a few points and therefore we do not represent the error bar there.

Parametric Bootstrap

We apply parametric bootstrapping to the estimated models and use the same 3 step process explained in Section 4.4. We first let the coefficients in the bootstrap sample be the mean of the pseudo-true values and afterwards do the same procedure but with the PC estimated parameters as the constant coefficients. These computations were done in the exact same manner as presented in Section A.2.6. Afterwards we apply parametric bootstrapping once more with different values for the constant coefficients of the bootstrap sample, namely the estimated coefficients from the final PC period. The results of both of the processes, the estimations of the bootstrap samples, are illustrated in the following figures with the bar plots of the 95% confidence intervals for all periods.

Parametric sample with the mean of the pseudo-true values:

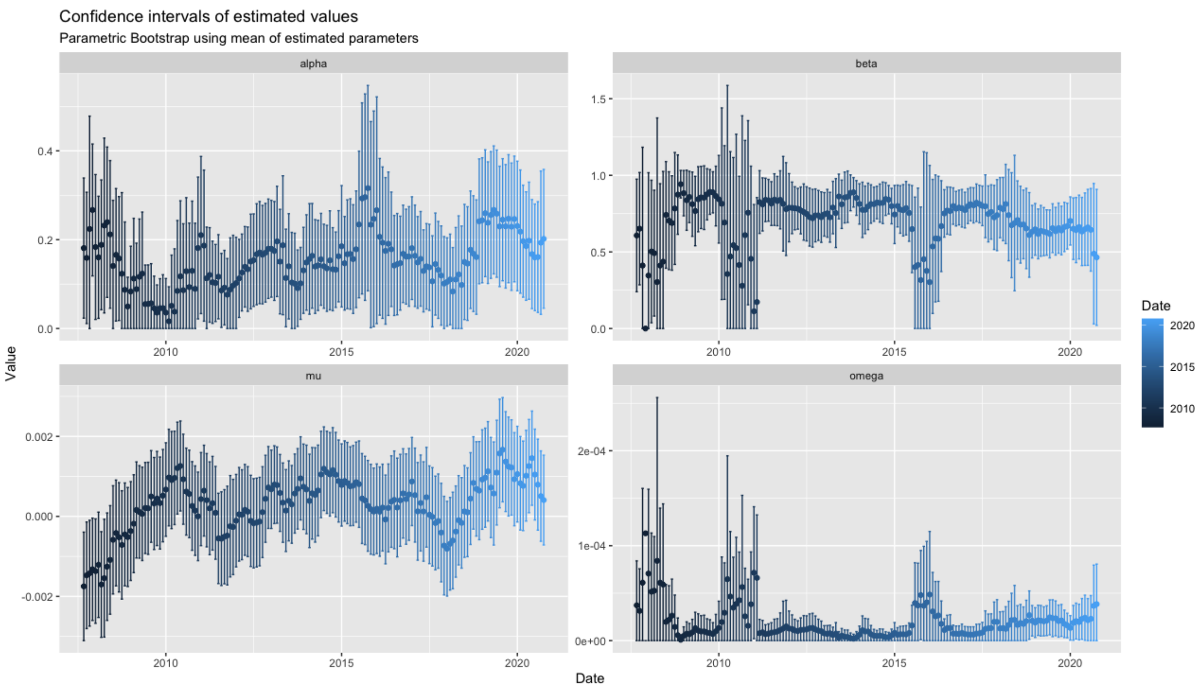

Parametric sample with the PC pseudo-true value:

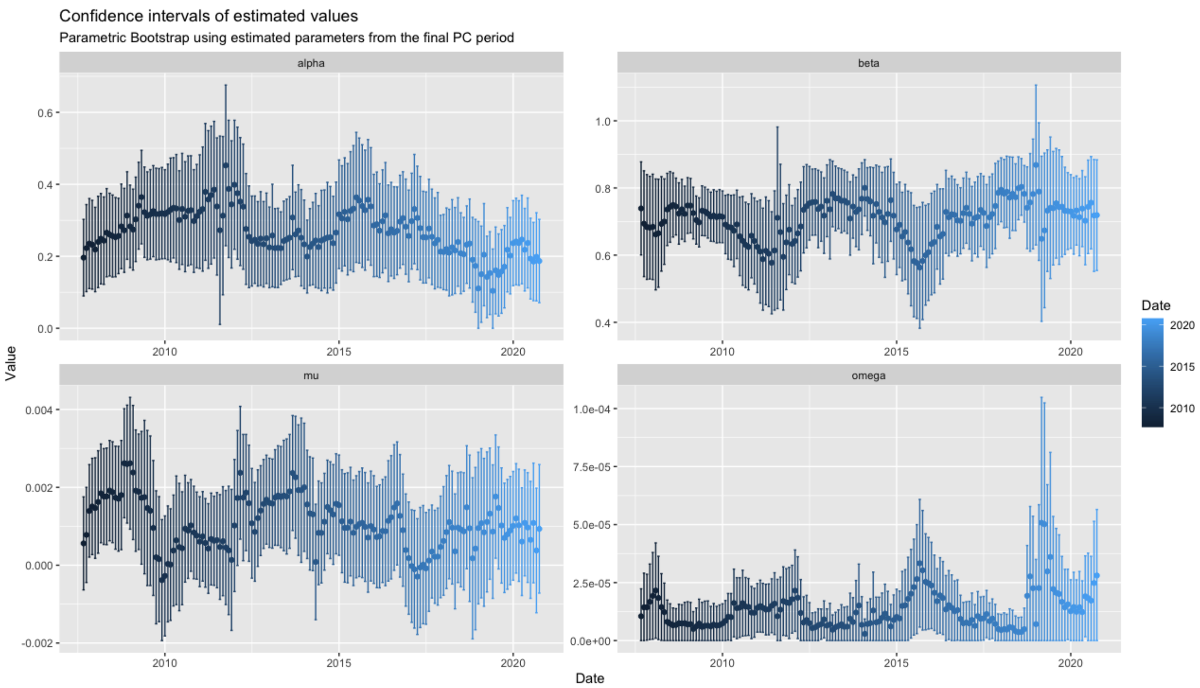

Validation 100 simulations

The parameter values are now again the mean of the pseudo-true values from the first estimation. We simulate 100 different samples and estimate new models for the five periods from section 4.1. We calculate how many confidence intervals of the BC periods are disjoint with the PC confidence interval. The result is presented in a data frame. We see that it is rare to have disjoint intervals which is not reasonable with reality. This justifies our thought of rejecting the constant coefficient assumption.

```
##    parameter   period num_disjoint
## 1         mu all data            0
## 2         mu      BC1            0
## 3         mu      BC2            0
## 4         mu      BC3            1
## 5      omega all data            1
## 6      omega      BC1            1
## 7      omega      BC2            1
## 8      omega      BC3            1
## 9      alpha all data            2
## 10     alpha      BC1            0
## 11     alpha      BC2            1
## 12     alpha      BC3            2
## 13      beta all data            2
## 14      beta      BC1            1
## 15      beta      BC2            1
## 16      beta      BC3            2
```

A.4 Modeling the DAX Index

Importing the data

```
getSymbols(
  # DAX Index
  c("^GDAXI"),
  src = "yahoo")
head(GDAXI) #first values
GDAXI_Open = as.vector(GDAXI$GDAXI.Open)
returns.GDAXI = (GDAXI_Open[-1] - GDAXI_Open[-length(GDAXI_Open)] ) /
  GDAXI_Open[-length(GDAXI_Open)]
```

```
## [1] "^GDAXI"
##            GDAXI.Open GDAXI.High GDAXI.Low GDAXI.Close GDAXI.Volume  GDAXI.Adjusted
## 2007-01-02    6614.73    6682.81   6614.73     6681.13     70947800         6681.13
## 2007-01-03    6681.23    6704.04   6667.58     6691.32    105913800         6691.32
## 2007-01-04    6662.24    6684.11   6634.74     6674.40    143461600         6674.40
## 2007-01-05    6661.90    6665.51   6588.03     6593.09    142562500         6593.09
## 2007-01-08    6603.55    6628.36   6579.90     6607.59    116670200         6607.59
## 2007-01-09    6625.28    6665.91   6606.87     6614.37    136073800         6614.37
```

Autocorrelation

Here we show the autocorrelation of the returns (a) and the squared returns (b) of the DAX index. We notice the tiny autocorrelation between the returns, even smaller than those of the AEX and CAC 40 indices, while the squared returns on the other hand show larger autocorrelation.

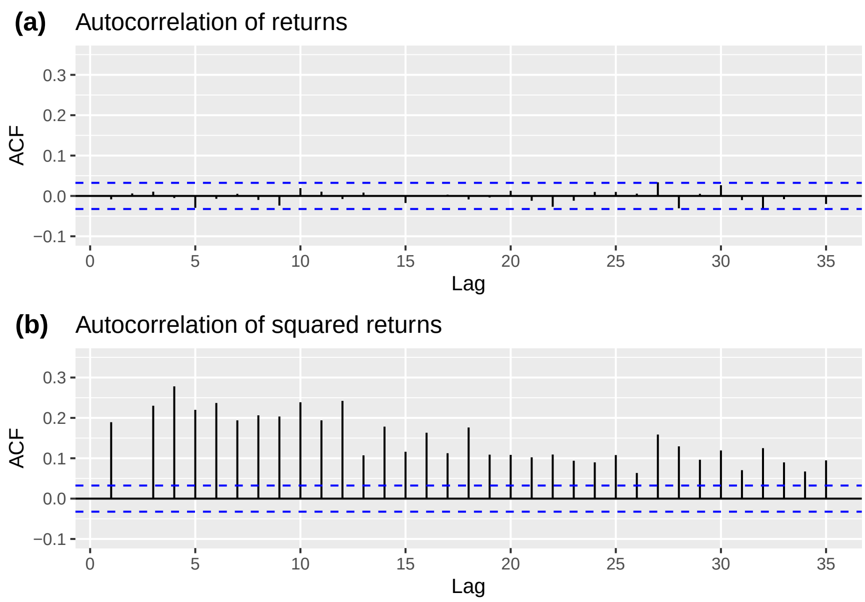

Kurtosis

Next the kurtosis coefficient of the returns is calculated. This value equals 10.32092, smaller than for AEX index but still high. This high value is expected since we know financial time series often have a leptokurtic distribution with fat tails and thus a high kurtosis coefficient. The plot below presents the distribution density function of the returns and its deviance from a normal distribution. The logarithmic scale is used in (b) for a clear view on the fatness of the tails.

```
#kurtosis coefficient
kurtosis_coef = kurtosis(na.exclude(returns.GDAXI))
kurtosis_coef
```

```
## [1] 10.32092
```

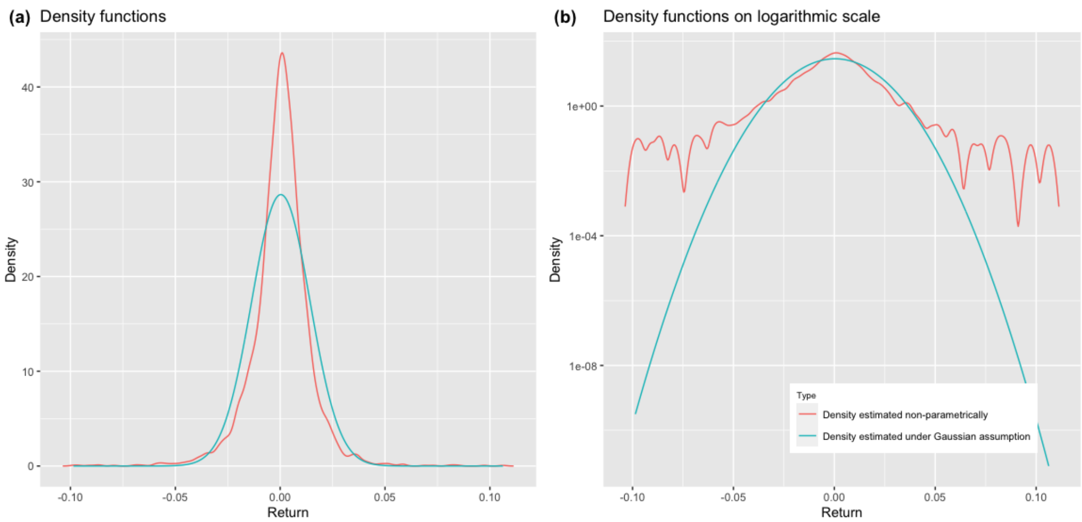

Parametric Bootstrap

We apply parametric bootstrapping to the estimated models and use the same 3 step process explained in Section 4.4. We first let the coefficients in the bootstrap sample be the mean of the pseudo-true values and afterwards do the same procedure but with the PC estimated parameters as the constant coefficients. These computations were done in the exact same manner as presented in Section A.2.6. Afterwards we apply parametric bootstrapping once more with different values for the constant coefficients of the bootstrap sample, namely the estimated coefficients from the final PC period. The results of both of the processes, the estimations of the bootstrap samples, are illustrated in the following figures with the bar plots of the 95% confidence intervals for all periods.

Parametric sample with the mean of the pseudo-true values:

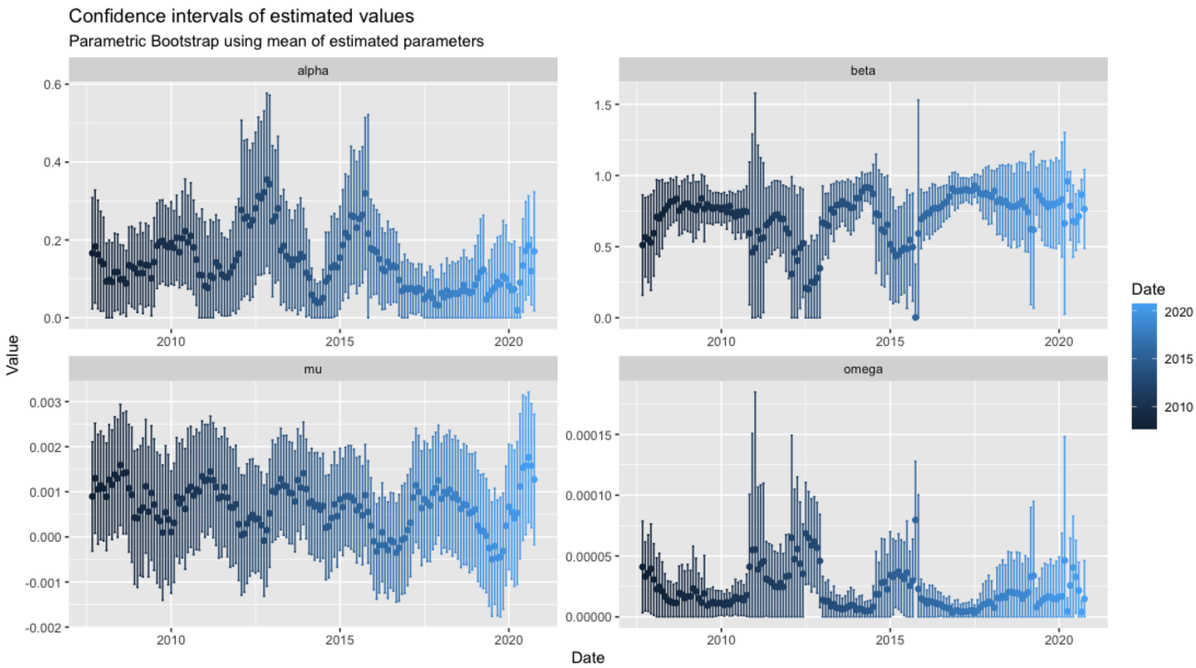

Parametric sample with the PC pseudo-true value:

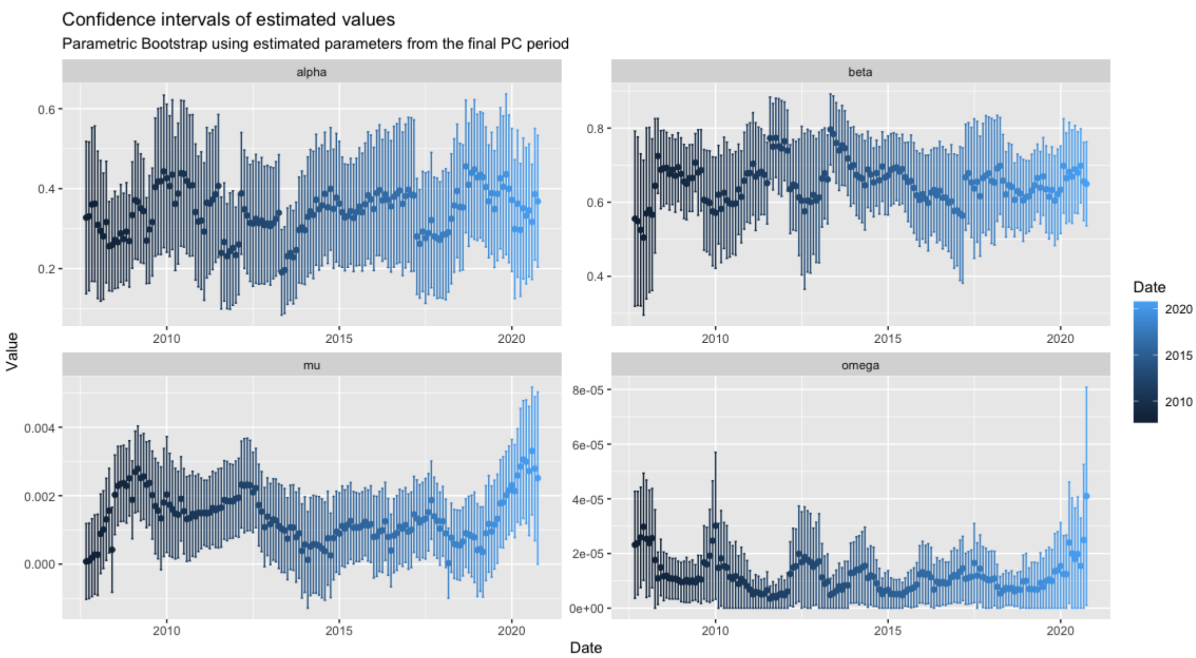

Validation 100 simulations

The parameter values are now again the mean of the estimated values from the first estimation. We simulate 100 different samples and estimate new models for the five periods from Section 4.1. We calculate how many confidence intervals of the BC periods are disjoint with the PC confidence interval. The result is presented in a data frame.

```
##    parameter   period num_disjoint
## 1         mu all data            1
## 2         mu      BC1            1
## 3         mu      BC2            1
## 4         mu      BC3            1
## 5      omega all data            0
## 6      omega      BC1            0
## 7      omega      BC2            1
## 8      omega      BC3            1
## 9      alpha all data            0
## 10     alpha      BC1            0
## 11     alpha      BC2            1
## 12     alpha      BC3            1
## 13      beta all data            2
## 14      beta      BC1            1
## 15      beta      BC2            1
## 16      beta      BC3            2
```

We observed in Figure 4.5 that the intervals for α of the BC periods were disjoint with the PC period intervals in practice. But from the results we see that it is rare to have disjoint intervals which is not reasonable with reality. This does not reflect the usual behaviour of the model and and justifies our thought of rejecting the constant coefficient assumption. We can not prove what model is right for this data but we can reject the wrong ones, and this model with constant parameters is wrong.

Bibliography

[1] V. Alto. Time Series: Why Do We Need Stationarity and Ergodicity.

[2] T. Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327, 1986.

[3] G. Box. Science and statistics. *Journal of the American Statistical Association*, 71(356):791–799, 1976.

[4] J. Brownlee. White Noise Time Series with Python.

[5] M. Buchanan. It's a (stylized) fact! *Nature Phys*, 8(3), 2012.

[6] R. Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236, 2000.

[7] A. Cornea-Madeira and R. Davidson. A parametric bootstrap for heavy-tailed distributions. *Econometric Theory*, 31(3):449–470, 2015.

[8] F. de Jong and B. Rindi. *The Microstructure of Financial Markets.* Cambridge University Press, 2009.

[9] Y. Dodge. *Coefficient of Kurtosis. In: The Concise Encyclopedia of Statistics*, pages 91–92. Springer New York, New York, NY, 2008.

[10] B. Efron and R. J. Tibshirani. *An Introduction to the Bootstrap.* Monographs on statistics and applied probability 57. Chapman and Hall, 1993.

[11] R. Engle. GARCH101: The use of ARCH/GARCH models in applied econometrics. *Journal of Economic Perspectives*, 15(4):157–168, 2001.

[12] F. Fabozzi and P. P. Drake. *Finance: Capital Markets, Financial Management, and Investment Management.* Wiley, 2009.

[13] Organisation for Economic Co-operation and Development (OECD). Global financial markets policy responses to covid-19, 2020.

[14] C. Francq and JM. Zakoian. *GARCH models: structure, statistical inference and financial applications.* John Wiley & Sons, 2010.

[15] A. Ghalanosm. *rmgarch: Multivariate GARCH models*, 2019. R package version 1.3-6.

[16] G. N. Gregoriou. *Stock Market volatility.* CRC Press, 2009.

[17] M. R. Hardy. An introduction to risk measures for actuarial applications. Technical report, University of Waterloo, 2006.

[18] C. He and T. Teräsvirta. Properties of moments of a family of garch processes. *Journal of Econometrics*, 92(1):173–192, 1999.

[19] T. Lux. *Stochastic Behavioral Asset-Pricing Models and the Stylized facts*, chapter 3. Elsevier, 2009.

[20] B. Mandelbrot. The Variation of Certain Speculative Prices. *The Journal of Business*, 36(4):394–419, 1963.

[21] D. B. Nelson. Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, 59(2):347–370, 1991.

[22] M. S. Perlin, M. Mastella, D. F. Vancin, and H. P. Ramos. A garch tutorial with r. *Revista de Administração Contemporânea*, 25(1), 2020.

[23] O. Saabbaghi. Modeling Conditional Volatility in R. *Journal for Economic educators*, 20(1):28–37, 2020.

[24] J. Schaffer. What not to multiply without necessity. *Australasian Journal of Philosophy*, 93(4):644–664, 2015.

[25] E. Smith. European stocks fall 12 % on the week as coronavirus grips markets, 2020. CNBC.

[26] R. S. Tsay. *Analysis of Financial Time Series.* Wiley, 2010.

[27] R. Wehrens, H. Putter, and L. M. C. Buydens. The bootstrap: a tutorial. *Chemometrics and intelligent laboratory systems*, 54:35–52, 2000.

[28] D. Wuertz. *Package 'fGarch'*, 2020. R package version 3042.83.2.

[29] E. Zivot and J. Wang. *Univariate GARCH Modeling.* Springer, 2006.