ISCTE Business School
Instituto Universitário de Lisboa

MODELLING AND FORECASTING BRENT PRICES

Tânia Cristina Dinis Marques e Silva

Dissertation submitted for the degree of
Master in Finance

Supervisor:
Prof. Doutor José Joaquim Dias Curto

November 2010

Resumo

Desde o início dos tempos que o Brent, mais conhecido por petróleo, tem sido utilizado em diversas aplicações, devido à sua elevada densidade energética, facilidade de transporte e relativa abundância. Nos últimos anos, o Brent tornou-se na fonte de energia mais importante, desempenhando um papel preponderante na manutenção da nossa actual sociedade. Neste contexto, o objectivo principal deste trabalho é modelar e prever os preços mensais e diários do Brent, de forma a melhor compreender e antever o seu comportamento. Na modelação e previsão dos preços utilizaram-se duas abordagens diferentes.

A primeira baseia-se na análise de séries temporais com memória longa. A presença de memória longa é verificada na média condicional e modelada a partir de modelos ARFIMA. Esta característica é também analisada na volatilidade da série e modelada através de modelos FIGARCH, FIAPARCH ou FIEGARCH. A outra abordagem considera modelos estocásticos de mudança de regime, nomeadamente modelos STAR, SETAR e MS-AR.

A modelação dos preços diários de Brent é feita com base em modelos de séries temporais considerando memória longa, uma vez que esta característica foi identificada na volatilidade da série. Modelos de mudança de regime foram também aplicados, no entanto a hipótese de não linearidade foi rejeitada. Relativamente aos resultados obtidos para a série mensal de preços, não foi detectada a presença de heteroscedasticidade condicional nem de memória longa. Os modelos de mudança de regime foram também considerados e, neste caso, foi identificado um modelo de dois estados, verificando-se diferenças significativas entre os regimes identificados.

Palavras-chave: Preço do Brent, Modelos de séries temporais com memória longa, Modelos de mudança de regime, Previsões.

Classificação JEL: C23; Q43.

Abstract

Since early times, the Brent, usually know as crude oil or petroleum, has been used in several fields due to its high energy density, easy transportability, and relative abundance. During the past years, it has become the most important source of energy in the world, and it plays a prominent part in the maintenance of our modernised and industrialised civilization. Therefore, the aim of this work is to analyse, understand and forecast the Brent prices behaviour, on a daily and monthly basis.

To model and forecast Brent prices, two different approaches have been considered. In the first approach, the presence of long memory is tested in the conditional mean and modelled by using long memory time series models, namely ARFIMA models. The long memory characteristic is also checked in the volatility, and it is modelled with FIGARCH models, or some of their variants, namely FIAPARCH or FIEGARCH models. The alternative approach considers stochastic regime-switching models, namely STAR, SETAR and MS-AR models.

The results obtained from the daily Brent prices analysis suggest that the data under investigation should be modelled with long memory time series, as there was evidence of long-range dependence on the volatility of the Brent prices. Regime-switching models were also applied to the daily data, but the hypothesis of non-linearity was rejected. Regarding the monthly Brent prices, neither conditional heteroskedasticity nor long-range dependence were identified. Regime-switching models were also considered, and in this case a two state model was identified, showing clear differences between the achieved regimes.

Keywords: Brent prices, Long memory time series models, Regime-switching models, Forecasting.

JEL Classification: C23; Q43.

Sumário Executivo

Nos últimos anos, o Brent assumiu um papel preponderante na nossa sociedade, tornando-se na fonte de energia mais importante para a sobrevivência da nossa actual civilização. O seu forte impacto na sociedade contribui para a necessidade de melhor compreender o comportamento dos seus preços, de forma a permitir, com maior exactidão, antever a sua tendência futura. Assim, o objectivo desta análise é modelar e prever os preços diários e mensais do Brent. Para o efeito, foram consideradas duas metodologias distintas: modelos de séries temporais com memória longa e modelos de mudança de regime.

A primeira metodologia consiste na verificação da presença de memória longa na média condicional da série em análise, procedendo-se à sua modelação com base em modelos ARFIMA (*Autoregressive fractionally integrated moving average*). Esta característica é também analisada na variância condicional dos dados, procedendo-se neste caso à aplicação de modelos FIGARCH (*Fractionally Integrated Generalized Autoregressive conditional heteroskedastic*), FIAPARCH (*Fractionally Integrated Asymmetric Power ARCH*) ou FIEGARCH (*Fractionally Integrated Exponential GARCH*). A outra abordagem permite analisar se o comportamento dos preços do Brent pode ser descrito por diferentes regimes, sendo considerados dois mecanismos de mudança de regime. O primeiro assume que a transição entre regimes é determinada por uma variável observável e, neste caso, consideram-se modelos STAR (*Smooth transition autoregressive*) ou SETAR (*Self-exciting threshold autoregressive*). Quanto ao segundo, em que a transição está dependente de uma variável não observável, aplicam-se os modelos MS-AR (*Markov-switching autoregressive*).

Os resultados diários sugerem que a série de preços de Brent deve ser modelada através de modelos de séries temporais com memória longa, uma vez que esta característica foi identificada na variância condicional dos dados. Como modelo final escolheu-se o FIGARCH(1,1) aplicado aos retornos da série diária de preços, dado ter sido o modelo que apresentou maior capacidade preditiva quando aplicado o teste de Diebold-Mariano. Os modelos de mudança de regime foram também aplicados à série.

No entanto a rejeição da hipótese de não-linearidade permite concluir que os preços diários do Brent devem ser analisados com base em modelos lineares.

As conclusões obtidas pela análise dos preços mensais de Brent são bastante dife–rentes das alcançadas no caso da análise diária. A aplicação de modelos de séries temporais permitiu a identificação de um modelo AR(6) para os retornos mensais, não se verificando nem heteroscedasticidade condicional nos resíduos do modelo, nem memória longa na média e variância condicional da série em análise. Os modelos de mudança de regime foram também aplicados e optou-se pelo modelo SETAR(2) de dois estados, uma vez que foi o que apresentou maior capacidade preditiva quando aplicado o teste de Diebold-Mariano. Este modelo é bastante interessante pois permite uma clara distinção entre os dois regimes identificados. Enquanto o primeiro regime, constituído pela maioria das observações, apresenta preços mensais menos elevados e com maior variação, o segundo regime é composto por preços mais elevados e com uma variação menor. Verifica-se, também, que o primeiro regime tem uma duração de quase dois anos e o segundo dura apenas poucos meses.

Nesta análise não foi possível identificar uma relação entre as mudanças de regime e presença de memória longa. No entanto, num estudo futuro, poderá ser bastante inte–ressante estudar-se a aplicação de modelos de mudança de regime à variância condicional dos preços diários do Brent.

Em seguida, apresentamos as contribuições desta tese para a literatura econométrica e financeira no âmbito da análise e previsão de preços de Brent. O primeiro ponto a referir é a análise e comparação das previsões obtidas pelos modelos de séries temporais com memória longa e modelos de mudança de regime.

Relativamente aos modelos de séries temporais, foram utilizados diversos modelos na modelação dos retornos dos preços do Brent, alguns dos quais (por exemplo, o FIAPARCH e FIEGARCH) pouco usuais no âmbito da literatura econométrica. É de referir, também, que a memória longa é testada simultaneamente na média e variância condicional. Quanto aos modelos de mudança de regime, investigamos se os preços do Brent são melhor caracterizados pelo seu comportamento passado ou por uma variável desconhecida e não observável.

Outro aspecto interessante é a análise paralela efectuada para os dados diários e mensais, que leva a conclusões bastante interessantes e distintas. Por fim, e de forma a comparar a capacidade preditiva dos modelos, utiliza-se, como complemento às tradicionais medidas de análise de previsão, o teste de Diebold-Mariano para testar se as diferenças de previsão entre os vários modelos identificados são estatisticamente significativas.

Acknowledgements

First of all I would like to express my gratitude to Professor Dias Curto for his guidance and support which were extremely helpful. My thanks to Manuel Garcia who guided me in the long memory topic and that indicated me the software that made this analysis possible.

A special thanks to EDP, in particular to the Risk department that gave me the opportunity to develop this dissertation, showing interest throughout the writing process of the present text. I hope my work corresponds to their needs. Thanks to my director Eng. Allen Lima, and to all my "risk" colleagues. A special word to Gonçalo Ferreira, for our many talks on the topic, and his pertinent suggestions along the work.

Last but not the least; I cannot finish without thanking my family and friends for all their support, and for giving me confidence and energy to undertake this work.

Contents

List of Tables

List of Figures

Basic Notation

General

L	Lag operator
$\nabla^d = (1-L)^d$	Lag difference operator of order d
$N(\mu, \sigma^2)$	Univariate normal distribution with expected value μ and variance σ^2
μ	Expected value
σ	Standard deviation

Acronyms

ACF	Autocorrelation function
ADF	Augmented Dickey-Fuller (unit root test)
AIC	Akaike's information criterion
AR	Autoregressive
ARCH	Autoregressive conditional heteroskedastic
ARFIMA	Autoregressive fractionally integrated moving average
ARIMA	Autoregressive integrated moving average
ARMA	Autoregressive moving average
BIC	Bayesian information criterion
CI	Confidence interval
DM	Diebold-Mariano (test)
EGARCH	Exponential GARCH
FIAPARCH	Fractionally Integrated Asymmetric Power ARCH
FIEGARCH	Fractionally Integrated Exponential GARCH

FIGARCH	Fractionally Integrated GARCH
GARCH	Generalized ARCH
IGARCH	Integrated GARCH
KPSS	Kwiatkowski, Phillips, Schmidt and Shin (test)
LM	Lagrange Multiplier (test)
LR	Likelihood ratio (test)
LSTAR	Logistic STAR
MAPE	Mean absolute percentage error
MS	Markov-switching
MS-AR	Markov-switching autoregressive
OPEC	Organization of the Petroleum Exporting Countries
PACF	Partial autocorrelation function
PP	Phillips-Perron (unit root test)
RMSE	Root mean square error
RSS	Residual sum of squares
SETAR	Self-exciting threshold AR
SIC	Schwartz information criterion
STAR	Smooth transition autoregressive
TAR	Threshold autoregressive
WN	White noise

Each chapter is divided into sections, with consecutive labelling of Equations, Tables and Figures within each chapter.

Chapter 1

Introduction

The Brent commodity, usually known as crude oil or petroleum, has been used since ancient times, and due to its high energy density, easy transportability and relative abundance, it has become the world's most important source of energy since the mid-1950s. Nowadays, it plays an important part in our society, as it supports the maintenance of the industrialised civilization as we know it, being a major issue of concern that has a very high impact on economy, politics, industry and technology.

The Brent[1] blend is a light crude oil that contains approximately 0.37% of sulphur, classifying it as sweet crude, and it is ideal for the production of gasoline and middle distillates. It is typically refined in Northwest Europe, but it can also be refined in East or Gulf coasts of the United States or in the Mediterranean region if the prices are good enough to export.

During the past years, especially in the last two years, the behaviour of the Brent prices has been characterised by high volatility. For example, in July 2008 it reached the 145 USD/bbl and six months latter it fell to 38 USD/bbl. However, according to the Organization of the Petroleum Exporting Countries (OPEC) the high Brent prices in the 2008 year cannot be justified by the fundamentals of physical and demand, which make them even harder to model and predict.

In our study we try to analyse and understand the Brent prices by studying the long memory characteristic or, in other words, the long-range dependence in the data set. The presence of long memory is tested in the conditional mean and modelled by using long memory time series models, namely autoregressive fractionally integrated moving average (ARFIMA) models. The long memory characteristic is also tested in the prices volatility, and in this case we model it with fractionally integrated general-

[1] The name "Brent" comes from the naming policy of Shell UK Exploration and Production, operating on behalf of Exxon and Shell, which originally named all of its fields after birds (in this case the Brent Goose).

ized autoregressive conditional heteroskedastic (FIGARCH) models, or some of their variants, namely fractionally integrated asymmetric power ARCH (FIAPARCH) or fractionally integrated exponential GARCH (FIEGARCH) models.

Additionally, several authors have mentioned a strong relation between the long memory characteristic and stochastic regime-switching models, considering very hard and complex to distinguish between long-range dependence and structural breaks. So, besides the long memory time series models, we also analyse regime-switching models. Regarding this class of models, we focus on the smooth transition autoregressive (STAR), self-exciting threshold autoregressive (SETAR) and Markov-switching autoregressive (MS-AR) models.

We now summarise the content and organization of this dissertation.

In Chapter 2 we present a brief literature review of the modelling techniques we follow in order to better understand, model and predict the Brent prices. We refer to studies that apply the same methods, namely long memory time series models and regime-switching models.

In Chapter 3 we give a brief description of the methods we use for modelling the Brent prices data set. We consider long memory time series models, such as ARFIMA models, for the conditional mean, and FIGARCH and its variants, namely FIAPARCH and FIEGARCH, models for the conditional variance. We also introduce and describe regime-switching models, like STAR, SETAR and MS-AR models. First, we present a brief theoretical introduction of each modelling technique, and then we provide the methodology we follow to analyse and understand the Brent prices behaviour.

In Chapter 4 we present the empirical results. We start by introducing our time series, and two data sets are analysed: the daily and monthly Brent prices. We use data from the beginning of January 2000 until the end of December 2009, and, in order to better understand the behaviour of both data sets, a descriptive analysis is performed and the presence of outliers is tested. Then we model the data sets considering long memory time series models and regime-switching models following closely the methodology described in Chapter 3. After identifying and validating several models we proceed to the model forecasting, and we consider both in-sample and out-of-sample prediction methods. At last, and having the forecasts of the different models, we are able to assess the prediction accuracy of all proposed models in order to choose the best model. We apply the traditional criteria (MAPE, RMSE) but we also use the Diebold-Mariano test to compare the predictive ability of the proposed models.

Finally, in Chapter 5 we draw some conclusions and present the main contributions of this thesis.

Chapter 2

Literature Review

In this chapter we present a brief literature review of the methodologies we apply to model and predict the Brent prices. In our analysis we focus on two main approaches: long memory time series models and regime-switching models. Next, we briefly review some studies concerning these modelling techniques.

The long memory characteristic, also known as long-range dependence, translates a high degree of persistence in the observed data, and the fact of time series exhibiting long memory is becoming more common. In fact, in the recent past, long memory has been found in time series from many fields, such as hydrology, meteorology and financial economics, among others.

The first study concerning this phenomenon was originally noticed by Hurst (1951). The author identified the presence of long-range dependence in hydrology while examining yearly discharges of the Nile River and measuring the long-term storage capacity of the reservoirs. Latter, many other authors, namely Mandelbrot and Wallis (1968) studied this effect, providing several studies related with data persistence in many natural science fields, like hydrology, meteorology, etc.

However, this phenomenon is not exclusive of non-financial time series. Long memory has become another stylised fact of financial time series, especially in terms of volatility. The growing necessity of analysing this kind of behaviour has lead to an increasing number of studies on this topic in order to better understand this phenomenon. Many authors, namely Baillie (1996), Beran (1994) and Robinson (2003), present an introduction on this subject as well as a discussion on several ways on how to estimate long-range dependence. Robinson (2003) edited a book with a collection of articles regarding a variety of topics concerning long memory time series. Baillie (1996), refers to applications in geophysical sciences, macroeconomics, assets pricing, stock returns, and exchange and interest rates, discussing also the ARFIMA models introduced by

Hosking (1981) that are used to describe the long memory in the conditional mean of the processes. Modelling long-range dependence in volatility is also done, by proposing extensions of the conditional heteroskedastic models to fractionally integrated models, namely FIGARCH and FIEGARCH processes.

Concerning other fields of natural sciences, Viswanathan et al. (1997) used long-range correlation measures when analysing the DNA patchiness. A very detailed study has been presented by Braun (2010) who developed a PhD. thesis on long memory applied to psychological behaviour. Besides all these non-financial analysis, the fact is that the long memory phenomenon is becoming a common issue when analysing financial data. Baillie and Chung (1996) present a study where the inflation for ten countries is analysed considering fractionally integrated $ARMA - GARCH$ models. These models are also used by Korkmaz et al. (2009) to analyse Istanbul stock exchange returns, and in this study, the author concludes that although long memory does not exist in the equity return, it exists in the volatility. A variant of these models has been proposed by Tse (1998), who examines the conditional heteroskedasticity of the yen-dollar exchange rate by considering FIAPARCH models.

Although there are a wide range of studies in several fields concerning the long memory phenomenon, one topic underrepresented in the literature is the long-range dependence applied to the Brent prices. However, Chaouachi (2005) developed a very interesting study on this topic, by searching for long memory in the Brent time series. The analysis suggests that long-range dependence is only present on the time series volatility, and therefore the author proposes modelling the conditional mean of the process with ARMA models, and its volatility with models that consider the long memory phenomenon, namely FIAPARCH models.

Another interesting paper was proposed by Medeiros and Scharth (2009) who analyse asymmetric effects and long memory in the volatility of Dow Jones stocks. This analysis focuses our interest, not due to the subject, but because the authors concluded that fractionally integrated models are an incomplete description of the volatility process, and their empirical results showed that multiple regime-switching modelling could be a promising alternative for these kind of applications. Like this author, many others, namely Diebold and Inoue (2001), Granger and Hyung (2004), and Banerjee and Urga (2005), consider a parallel analysis on long memory and structural breaks, and the reasoning is quite interesting and simple. In fact, according to Granger and Hyung (2004) theory and simulation results show that it is very complex to distinguish between the long memory characteristic and structural breaks on the data set, and that neglecting breaks may generate a long memory effect in the autocorrelation function.

The authors analyse the S&P 500 absolute stock returns by comparing the results from the two different methodologies, and they concluded that the performance of both kind of models for the in-sample and out-of sample forecasts are quite similar, showing no statistical difference.

Diebold and Inoue (2001) also argue that regime-switching is easily confused with long memory, and present a theoretical and Monte Carlo study which corroborate the idea that stochastic regime-switching and long memory are strongly related.

Considering this relation, in our study we will not only use long memory time series models but also regime-switching models.

Regarding regime-switching models, different approaches are being presented. Engle and Hamilton (1990), in order to analyse the behaviour of the dollar exchange rate, consider a model which assumes the existence of an unobservable variable that characterises the regime in which the process is at time t, and depending on the value of this variable, the data assumes different characteristics. As the evolution of the unobservable variable is based on a Markov chain, this class of models is a particular case of the regime-switching models, and is referred to as Markov-switching models. A variant of this class of models is proposed by Caraiani (2010). The author, using a two-state Markov-switching autoregressive (MS-AR) approach, analyses and predicts the business cycle in the Romanian economy. The achieved results show that the macroeconomic dynamics in the Romanian economy can be very well explained and forecasted when using regime-switching models. Schindlmayr (2005) also proposes a MS-AR model applied to the daily electricity spot prices from the European Energy Exchange. The author identifies two different regimes, one characterised by the normal behaviour of the data, and the other by the price spikes with a high volatility and strong mean-reversion. Another MS-AR model is presented by Hardy (2001), who analyses the monthly data from the Standard and Poor's 500 and the Toronto Stock exchange 300 indices.

A different type of Markov-switching models has been proposed by Dafas (2004) who tries to characterise the stochastic behaviour of crude oil prices. The author, using market data of international crude oil spot prices for the last seventeen years, calibrates a mean-reverting Markov regime-switching jump-diffusion model and estimates its parameters using the Hamilton filter. A different study, also considering commodities, has been proposed by Chen and Forsyth (2010). The authors propose a one-factor regime-switching model for natural gas spot prices and analyse the implication of the model on the valuation and optimal operations of a natural gas storage facility. Each model follows either a mean-reverting process or a Brownian motion, and the obtained results seemed to reflect the gas prices dynamics in each regime.

A different approach, considering regime-switching models, is proposed by Dijk et al. (2002) and consists on the smooth transition autoregressive (STAR) time series models, as well as some of their variants, namely the logistic STAR (LSTAR) and the self-exciting threshold AR (SETAR) models. The authors, after providing a rich introduction and discussion on these models, present an empirical example concerning the US unemployment. The behaviour of the US unemployment rate series is modelled and analysed, and several extensions of the STAR models are performed in order to choose the best model. Another study concerning the SETAR models is proposed by Boero and Marrocu (2004). The authors, using daily data of the Euro effective exchange rate, evaluate the out-of-sample performance of SETAR models relative to an AR-GARCH model. The results show that the SETAR models have significantly improved the forecasting performance when the forecast origin was conditioned on a specific regime.

A very interesting analysis is presented by Ismail and Isa (2006), who model the monthly returns of exchange rates of three Asian countries (Malaysia, Singapore, and Thailand) against the British pound using two types of regime-switching models: MS-AR and SETAR models. After comparing the results of the two classes of models, the authors concluded that the fitting of the MS-AR model is better than the SETAR model, and also that the regime-switching models outperform linear models in terms of fitting returns. In our study, like in this previous one, we will also consider and compare these two types of regime-switching models.

Chapter 3

Description of Models

In this chapter we provide a brief description of the methodologies we follow when analysing, modelling and forecasting Brent prices.

In the first section, Section 3.1, we introduce time series models that exhibit long memory and in Section 3.2 we describe the regime-switching models. For each section we present a brief theoretical introduction of each approach, as well as the methodology we follow to analyse and understand the Brent prices behaviour. As we only give an overview of the used techniques, along each section we provide references for a more detailed analysis.

Finally, in Section 3.3 we present several procedures in order to evaluate the prediction ability that will help us when selecting the final model.

3.1 Long Memory Time Series Models

As the main purpose of this study is to analyse and characterise the Brent prices behaviour, we will consider long memory time series models by testing long-range dependence not only on the first moment (mean), but also in terms of the variance (that represents volatility).

Therefore in Subsection 3.1.1 we present a general introduction to long memory time series models. We start by introducing the long memory concept, and next we give a brief description on the models we use: the autoregressive fractionally integrated autoregressive moving average (ARFIMA) model for the conditional mean analysis, and some variants of the autoregressive conditional heteroskedastic (ARCH) model for the conditional variance modelling, namely the fractionally integrated generalized ARCH (FIGARCH), asymmetric power ARCH (FIAPARCH), and fractionally integrated exponential generalized ARCH (FIEGARCH) models. In the last section, Subsection 3.1.2, we present the methodology followed in the Brent prices modelling considering long memory time series models.

3.1.1 Introduction

Long memory, also known as long-range dependence, is traditionally defined either in the time domain or in the frequency domain. In terms of time domain it results on a decaying rate of autocorrelation functions (with special interest on the long lags). As described by Beran (1994), let $X = \{X_t; t \in \mathbb{Z}\}$ be a stationary process, then if there exists a constant $c > 0$ and $\alpha \in (0, 1)$ such that

$$\lim_{k \to \infty} \frac{\rho_k}{ck^{-\alpha}} = 1, \tag{3.1}$$

then X is a stationary process with long memory, where ρ_k are the autocorrelation coefficients of order k. A more general definition has been proposed by McLeod and Hipel (1978) and considers that

$$\lim_{n \to \infty} \sum_{j=-n}^{n} |\rho_j| = \infty, \tag{3.2}$$

where n corresponds to the number of observations and ρ_j to the autocorrelation coefficients.

Other definition involving frequency domain considers the spectral density of the process. Therefore, let $X = \{X_t; t \in \mathbb{Z}\}$ be a stationary process and f be the spectral density. If there is a constant $c_f > 0$ and $\beta \in (0, 1)$ such that the condition

$$\lim_{\lambda \to \infty} \frac{f_\lambda}{c_f |\lambda|^{-\beta}} = 1 \tag{3.3}$$

verifies, then the process X exhibits long memory, or long-range dependence.

After defining the long memory concept, we present a class of models that are able to capture long-range dependence in the conditional mean of a process: the ARFIMA models. However, before referring these models, we present a brief introduction on autoregressive moving average (ARMA) models. Following Box and Jenkins (1976) and Wei (1990), a stationary process $X = \{X_t; t \in \mathbb{Z}\}$ is defined as an ARMA(p, q) model if

$$\phi_p(L)(X_t - \mu) = \theta_q(L)\epsilon_t, \tag{3.4}$$

where $p, q \in \mathbb{N}_0$, μ is the process mean, L is such that $L^j X_t = X_{t-j}$ (commonly known as lag operator), and $\{\epsilon_t; t \in \mathbb{Z}\} \sim \mathrm{WN}(0, \sigma_\epsilon^2)$. The terms $\phi_p(L)$ and $\theta_q(L)$ are respectively the autoregressive and moving average polynomials, given by $\phi_p(L) = 1 - \sum_{j=1}^{p} \phi_j L^j$ and $\theta_q(L) = 1 - \sum_{j=1}^{q} \theta_j L^j$.

However in many situations instead of a stationary series, we deal with non-stationary behaviour, and therefore we need to consider the ARFIMA models, discussed in Baillie (1996), and given by

$$\phi_p(L)(1-L)^d(X_t - \mu) = \theta_q(L)\epsilon_t, \tag{3.5}$$

with d being the fractional differencing parameter, and where μ, $\phi_p(L)$, $\theta_q(L)$, and ϵ are as previously defined. Note that the roots of $\phi_p(L)$ and $\theta_q(L)$ lie outside the unit circle. For $-0.5 < d < 0.5$ the process is covariance stationary, for $d < 1$ is mean reverting, and for $0 < d < 0.5$ it is assumed to have long memory. The Equation (3.5) can be reduced to a stable ARMA process when $d = 0$, and to an ARIMA model when $d = 1$. So we can conclude, and as referred by Brunetti (1999), that the ARFIMA models are, in fact, very flexible as they can capture both long memory (by estimating the d parameter) and short term memory (considering the dynamics captured by the autoregressive and moving average polynomials).

The relevant issues previously discussed when dealing with long-range dependence in the first moment of a process, are also important and necessary when focusing on the second moment, its volatility. As the assumption of homoskedastic errors is not always valid, Engle (1982), in order to overcome this situation, proposed the autoregressive conditional heteroskedastic (ARCH) model. Latter a generalisation of this class of models was proposed by Bollerslev (1986), introducing the generalized ARCH (GARCH) model. A GARCH(p,q) process is defined as

$$\begin{aligned}
\epsilon_t &= \upsilon_t \sqrt{h_t} \\
h_t &= \alpha_0 + \sum_{i=1}^{q} \alpha_i \epsilon_{t-1}^2 + \sum_{i=1}^{p} \beta_i h_{t-i} \\
&= \alpha_0 + A(L)\epsilon_t^2 + B(L)h_t
\end{aligned} \tag{3.6}$$

with $\{\upsilon_t\} \sim iid(0,1)$ and

$$\begin{gathered}
p \geq 0, \quad q > 0, \\
\alpha_0 > 0, \quad \alpha_i \geq 0,\ i = 1, \ldots, q \\
\beta_i \geq 0,\ i = 1, \ldots, p.
\end{gathered} \tag{3.7}$$

The GARCH(p,q) model can also be expressed as an ARMA(m,q) process in ϵ_t^2 according to

$$\phi(L)\epsilon_t^2 = \alpha_0 + (1 - \beta(L))\upsilon_t, \tag{3.8}$$

where $\upsilon_t = \epsilon_t^2 - \sigma_t^2$ corresponds to the innovations in the conditional variance process, $m = \max(p, q)$, and $\phi(L) = \{1 - \alpha(L) - \beta(L)\}$.

Like in the analysis of the first moment of a process, long memory and persistence can also be characteristics of the conditional variance of a process. Therefore, and in order to capture the long-range dependence Baillie et al. (1996) proposed the fractional integrated GARCH (FIGARCH) model. A FIGARCH(p, d, q) process can be represented as

$$\phi(L)(1 - L)^d\epsilon_t^2 = \alpha_0 + (1 - \beta(L))\upsilon_t, \tag{3.9}$$

where all roots of $\phi(L)$ and $\{1 - \beta(L)\}$ lie outside the unit circle. When $d = 0$ the Equation (3.9) of the FIGARCH process reduces to a GARCH model, while when $d = 1$, we get an IGARCH process.

As, in financial markets, frequently occurs that positive and negative shocks have a different impact on the volatility, meaning that usually a negative event (bad news) conducts to higher volatility than a positive one. This phenomenon is named as leverage effect, see Black (1976). Therefore, Tse (1998) proposed a model that are able to capture both, the volatility asymmetric behaviour and the long memory presence: the fractional integrated asymmetric power ARCH (FIAPARCH) model. A FIAPARCH(p, d, q) process is, then, given by

$$(1 - L)^d\phi(L)g(\epsilon_t)^\delta = \alpha_0 + (1 - \beta(L))\xi_t \tag{3.10}$$

where $\xi_t = g(\epsilon_t)^\delta - \sigma_t^\delta$, $g(\epsilon_t) = |\epsilon_t| - \gamma\epsilon_t$, with $\alpha > 0$, $\delta \geq 0$, $|\gamma| < 1$, and all the other terms are as previously defined. If the asymmetry parameter, γ, is positive (negative) then it indicates that negative (positive) shocks conducts to higher volatility than positive (negative) ones.

Another model that can capture the non-symmetric behaviour, was proposed by Bollerslev and Mikkelsen (1996) and represents an extension of the FIGARCH model based on the exponential GARCH model previously introduced by Nelson (1991). The developed model is a fractional integrated exponential GARCH (FIEGARCH) model, and a FIEGARCH(p, d, q) process can be written as

$$\log(\sigma_t^2) = \alpha_0 + \phi(L)^{-1}(1 - L)^{-d}[1 - \lambda(L)]g(\xi_{t-1}), \tag{3.11}$$

with

$$g(\xi_t) = \theta\xi_t + \gamma[|\xi_t| - E|\xi_t|]. \tag{3.12}$$

We stress that all roots of $\phi(L)$ and $\lambda(L)$ lie outside the unit circle, and if we consider $d = 0$, the Equation (3.11) simplifies to an EGARCH model.

For a more detailed explanation on these matters we refer to Baillie (1996), Brunetti (1999), Diebold and Inoue (2001), Breidt et al. (1998), Robinson (1995) and Banerjee and Urga (2005).

3.1.2 Modelling Methodology

In this subsection we present our methodology in the Brent prices modelling and forecasting, considering long memory time series models.

Many authors divide the time series modelling methodology in several steps (see, for example, Wei (1990)), and in this thesis we follow the same procedure, but we introduce some variants in order to better achieve our purpose. Therefore the following steps are the ones considered in our analysis:

1. Stationarity and long memory testing;
2. Model identification and validation; and
3. Forecasting and model selection.

Next we briefly describe each of the points referred above.

1. Stationarity and long memory testing

Before testing for the presence of long memory in our time series, it is necessary to check if the data has a stationary behaviour.

Although a process can be non-stationary in the mean, variance, or in other moments, in this study we only analyse and deal with mean and variance non-stationary behaviour. Therefore, and since a process can be stationary in the mean but not necessarily in the variance, it is a common practice to deal firstly with the variance non-stationary behaviour and only afterwards test if it is necessary to apply a mean stabilising transformation.

There are several transformations that can be applied in order to stabilise and homogenise the variance. The most common transformation is the Box-Cox's power transformation proposed by Box and Cox (1964), and that is given by

$$T_\lambda(X_t) = \begin{cases} \frac{X_t^\lambda - 1}{\lambda}, & \lambda \neq 0, \\ \log(X_t), & \lambda = 0 \end{cases} \tag{3.13}$$

where $\{X_t\}$ is the original data set, and $\lambda \in [-1, 1]$ is the transformation parameter. The best transformation corresponds to the value of λ that minimises the sample variance of the transformed series. For more details on this procedure we refer to Pires (2001). We recall that in financial markets, the logarithmic transformation is the most commonly used.

The following step, after stabilising the variance, is to test if the mean non-stationarity remains. This can be accomplished in several ways: by analysing the empirical autocorrelation function (ACF) or partial autocorrelation function (PACF); or by computing some tests (implemented in several softwares, namely *R* and *Eviews*). The tests used in this study are the Kwiatkowski, Phillips, Schmidt and Shin (KPSS) test, the Augmented Dickey-Fuller (ADF) unit root test, and the Phillips-Perron (PP) unit root test. All tests are available in the *R* Software, as well as in Eviews, so for more details on these tests we refer to Trapletti (2005) and Software (2009), respectively.

When the mean non-stationarity is confirmed it is necessary to apply a differencing operator $\nabla^k X_t = X_t - X_{t-k}$, possibly more than once, or of k order, with the purpose of removing, respectively, any trend or cyclic component (of k order).

Now that the series is stationary, it is possible to check for the presence of short or long memory. There are several tests to check for the presence of long memory in a time series, and one of the oldest methods is the rescaled-range, or simply R/S, statistic. This measure was firstly introduced by Hurst (1951) and then developed and refined by Mandelbrot and Wallis (1968). The R/S statistic corresponds to the range of partial sums of deviations of a time series from the mean, rescaled by its standard deviation. Therefore, consider a time series $X = \{X_t,\ t = 1, \ldots, T\}$, where T corresponds to the number of observations. In order to evaluate the R/S statistic, after subdividing our data set into υ partitions we define, for each partition, its range (R) according to

$$R_T = \left\{ \max_{1 \leq i \leq T} \sum_{t=1}^{i} (x_t - \overline{x}) - \min_{1 \leq i \leq T} \sum_{t=1}^{i} (x_t - \overline{x}) \right\}, \tag{3.14}$$

where $\overline{x}$ is the sample mean of the data set. The sample standard deviation is also determined, and, for each partition, we obtain the R/S statistic. As referred in Mandelbrot and Wallis (1968) this statistic follows asymptotically

$$(R/S)_s \propto Cs^H. \tag{3.15}$$

In order to estimate the parameter H, and according to Grau-Carles (2005), it is usual to run a linear regression over the growing temporal horizons, $s = t_1, t_2, \ldots, T$, each referring to its respective partition,

$$\ln(R/S)_s = \ln(C) + H\ln(s). \tag{3.16}$$

If the H estimated value is 0.5 we can conclude that the process has short memory, if the estimate is within the $(0.5, 1)$ interval then it is a stationary process with long memory, and if the H parameter is between $(0, 0.5)$, the process is anti-persistent[1]. Many authors identify a relation between the H parameter and the differencing parameter d. In case of an infinite variance process the relation is given by $H = d + \frac{1}{\alpha}$, but in the case of a finite variance process, the relation is simply $H = d + \frac{1}{2}$. For a more detailed analysis see Chaouachi (2005) and Grau-Carles (2005).

A methodology to estimate the H parameter was developed and implemented by Wuertz (2009) and is available in the R Software. However many authors, namely Lo (1991), point out the incapacity of this statistic to distinguish between long and short memory. Therefore, and to overcome the lack of robustness of the R/S statistic, different statistics are also used.

A different methodology has been proposed by Geweke and Porter-Hudak (1983) and is based on the spectral density function. The idea behind this method is that, for models applied to stationary time series, the spectral density function is bounded at zero frequency, while for long memory models it becomes unbounded at frequency $\lambda = 0$. Therefore Geweke and Porter-Hudak (1983) proposes to estimate d according to

$$\begin{aligned} \ln\{I(\lambda_{j,T})\} &= \ln\left\{\frac{\sigma^2 f_u(0)}{2\pi}\right\} - d\ln\left\{4\sin^2\left(\frac{\lambda_{j,T}}{2}\right)\right\} \\ &+ \ln\left\{\frac{f_u(\lambda_{j,T})}{f_u(0)}\right\} + \ln\left\{\frac{I(\lambda_{j,T})}{f_u(\lambda_{j,T})}\right\}, \end{aligned} \tag{3.17}$$

where $I(\lambda_{j,T})$ and $f_u(\lambda_{j,T})$ are respectively the periodogram and the spectral density function of the time series at frequencies $\lambda_{j,T} = (2\pi j)/T$, with $\lambda = 1, \ldots, T$, and T corresponds to the number of observations of the series. Considering that

[1] A anti-persistent time series reverses more often than a random series would be.

$\ln\{4\sin^2(\lambda_{j,T}/2)\}$ is the explanatory variable, $\ln\{I(\lambda_{j,T})/f_u(\lambda_{j,T})\}$ is the disturbance term, and the remaining two terms correspond to the intercept term, (3.17) can be simply represented, as Grau-Carles (2005) proposed, by

$$\ln\{I(\lambda_{j,T})\} = \beta_0 + \beta_1 \ln\left\{4\sin^2\left(\frac{\lambda_{j,T}}{2}\right)\right\} + u_{\lambda_{j,T}}. \tag{3.18}$$

Therefore the proposed estimator for d, is the slope coefficient obtained in the least squares regression of $\ln\{I(\lambda_{j,T})\}$ on $\ln\{4\sin^2(\lambda_{j,T}/2)\}$.

A very similar statistic, proposed by Robinson (1994), is also computed. It is available in the TSM software and was implemented by Davidson (2010). This statistic gives a semi-parametric estimate of the H parameter (which can be related to the d as previously described) through

$$\hat{H}_{mq} = 1 - \frac{\ln \hat{F}(q\lambda_m)/\hat{F}(\lambda_m)}{2\ln q} \tag{3.19}$$

where F is the average periodogram function, q must be restricted to the interval $(0, 1)$ and it is assumed to be 0.5, and $\lambda_m = 2\pi T^{-0.35}$.

2. Model identification and validation

After estimating the long memory parameter (d) we are able to proceed to the model identification. The initial step is to model the long-range dependence in the first conditional moment of the process, meaning that an adequate ARFIMA model shall be adjusted to the already stationary time series. The orders of the autoregressive and moving average parameters can be identified by analysing the autocorrelation and partial autocorrelation functions.

Then a similar procedure is followed, but now considering the second conditional moment of the process (the time series volatility), and in this case a range of models can be selected to better adjust the data volatility, namely, FIGARCH, FIAPARCH and FIEGARCH models.

To identify the models and estimate its parameters we use the TSM software (Davidson (2010)), as it allows for a very rich analysis on the long memory phenomenon.

After identifying the model and estimating its parameters (through the conditional maximum likelihood estimation method), we need to evaluate its adequacy by testing if its residuals can be considered as realisations of a white noise (WN) process. In other words, we want to test if the residuals are uncorrelated random variables with null mean and constant variance.

We start by analysing the empirical autocorrelation functions of the residuals, and both ACF and PACF must show no statistically significant pattern. The hypothesis of uncorrelated residuals can also be checked by computing the Ljung-Box and the Box-Pierce tests, which are both implemented in *R*, Team (2004). To evaluate whether the variance is constant, we may simply analyse the residuals plot in order to examine if there is any change in the behaviour of the residuals in terms of their volatility. Another way, is to apply the Box-Cox's power transformation (see (3.13)) for different values of λ.

The Lagrange Multiplier test (defined in Eviews, Software (2009)) is also computed in order to check for the presence of conditional heteroskedasticity and for any evidence of ARCH effects, when we are analysing the residuals of the identified model for the conditional mean.

Based on the test results, we may have to adjust other models, so that all assumptions are verified.

3. Forecasting and model selection

After identifying and validating several models in the previous point we are now able to proceed with the forecasting by taking into account two types of prediction methods. The first one is an in-sample method which consists of one-step-ahead prediction and its main purpose is to test the model stability. The second approach, from our point of view much more interesting, consists on the prediction of k out-of-sample values.

Having the forecasts of the identified models, the next step is to select the one that better adjusts reality, and there are several criteria, described in Section 3.3, that can help us to choose which is, in fact, the best model.

3.2 Regime-switching Models

In this section we present the alternative approach used to model and forecast the Brent prices: regime-switching models. The idea is to analyse if the behaviour of the data set can be described by different states, or regimes. Therefore in Subsection 3.2.1 we present a brief introduction to regime-switching models by considering the two types of switching mechanisms used in this study, the smooth transition autoregressive (STAR) models and the Markov-switching (MS) models.

In Subsection 3.2.2 we describe our methodology to understand and model the Brent prices behaviour considering these state-dependent models.

3.2.1 Introduction

The regime-switching models can be defined as non-linear models that allow for the possibility that the dynamic behaviour of the time series depends on regimes that occur at any given point in time. Therefore this class of models lets the series switch stochastically between alternative cases of the conditional mean and variance equations. In other words, it allows for certain properties of the time series, such as mean, variance and/or autocorrelation to be different across regimes.

The first class of these models to be introduced is the STAR models. Initially, Tong (1983) proposed the threshold autoregressive (TAR) model, which assumes that the regime in a given point in time, t, can be determined by an observable variable. Therefore the regime is defined considering the value of the observable variable (also named as threshold variable) relative to a define threshold value, which we refer to as c.

A general class of state-dependent models is the STAR model, which assumes that the transition between regimes is gradual. A two-regime STAR(p) model is given by

$$\begin{aligned} y_t &= (\phi_{1,0} + \phi_{1,1}\, y_{t-1} + \cdots + \phi_{1,p}\, y_{t-p})\ (1 - G(s_t; \gamma, c)) \\ &+ (\phi_{2,0} + \phi_{2,1}\, y_{t-1} + \cdots + \phi_{2,p}\, y_{t-p})\ G(s_t; \gamma, c) + \epsilon_t, \end{aligned} \tag{3.20}$$

with $t = 1, \ldots, T$. The s_t variable, as referred in Dijk et al. (2002), can assume several forms. As discussed by Teräsvirta (1994), s_t is considered to be a lagged endogenous variable, $s_t = y_{t-\delta}$ (with δ defined as the delay parameter), while Frances and Dijk (2000) assume s_t to be an exogenous variable ($s_t = z_t$). The other term is the transition function, $G(s_t; \gamma, c)$, whose choice determines different regime-switching behaviours.

A popular choice for this function is the logistic function,

$$G(s_t; \gamma, c) = \frac{1}{1 + \exp(-\gamma(s_t - c))}, \tag{3.21}$$

where γ is the smoothness parameter, c is the threshold value, and the STAR model turns into a logistic STAR (LSTAR) model. When γ tends to zero the logistic function approaches a constant value (equal to 0.5), and when $\gamma = 0$ the LSTAR model reduces to a linear model. Another particular model occurs when γ becomes very large and $s_t = y_{t-\delta}$, transforming the transition function into an indicator function $I[s_t > c]$ and leading the transition between different regimes to be instantaneous. This class of models is called a self-exciting TAR (SETAR) model, and a two-regime SETAR(p) is represented as

$$\begin{aligned} y_t &= (\phi_{0,1} + \phi_{1,1}\ y_{t-1} + \cdots + \phi_{1,p}\ y_{t-p})\ (1 - I[s_t > c]) \\ &+ (\phi_{2,0} + \phi_{2,1}\ y_{t-1} + \cdots + \phi_{2,p}\ y_{t-p})\ I[s_t > c] + \epsilon_t, \end{aligned} \quad (3.22)$$

where $I[A] = 1$ if the event A occurs, and $I[A] = 0$ otherwise. For a more detailed description on these models we refer to Tong (1983), and Frances and Dijk (2000).

In some cases is interesting to allow for more than two regimes. Therefore a m-regime SETAR model can be obtained by assuming $m+1$ threshold values, $c_0, c_1, \ldots, c_{m-1}, c_m$, such that $-\infty = c_0 < c_1 < \cdots < c_{m-1} < c_m = \infty$, and the extension of (3.22) to a m-regime of order p is given by

$$y_t = \phi_{0,j} + \phi_{1,j} y_{t-1} + \cdots + \phi_{p,j} y_{t-p} + \epsilon_t \qquad \text{if } c_{j-1} < y_{t-1} \leq c_j, \quad (3.23)$$

for $j = 1, \ldots, m$ states.

A similar approach is done for the STAR model. Consider the same subset of threshold values defined above, $c_1, \ldots, c_{m-1}$, and an additional set of $m-1$ smoothness parameters, $\gamma_1, \ldots, \gamma_{m-1}$, then a m-regime STAR model can be defined as

$$y_t = \phi_1' x_t + (\phi_2 - \phi_1)' x_t G_1(s_t) + \cdots + (\phi_m - \phi_{m-1})' x_t G_{m-1}(s_t) + \epsilon_t, \quad (3.24)$$

where $x_t = (1, \tilde{x}_t')'$ with $\tilde{x}_t = (y_{t-1}, \ldots, y_{t-p})'$, $\phi_i = (\phi_{i,0}, \ldots, \phi_{i,p})'$ with $i = 1, \ldots, m$, and $G_j(s_t) \equiv G_j(s_t; \gamma_j, c_j)$ are the logistic functions defined in (3.21) for $j = 1, \ldots, m-1$.

The other class of regime-switching models, contrary to the one previously described, assumes that the regime that occurs at time t cannot be determined by an observable variable. In fact, the switching between regimes depends on an unobservable variable, which we denote as s_t, and that characterises the regime or state in which the process is at a given time t.

Many authors, namely Hamilton (1989) and Teräsvirta (1994) among others, have made many relevant contributions among non-linear models, especially regime-switching models. In this analysis we focus on Markov-switching autoregressive (MS-AR) models, considering a finite number of states. Additionally, MS-AR models can be seen as an extension of autoregressive models to the non-linear case. A common representation (see, for example, Frances and Dijk (2000)) of a MS-AR(p) model is given by

$$A(n) = \begin{cases} \phi_{0,1} + \sum_{i=1}^{p} \phi_{i,1}\ y_{t-i} + \epsilon_t, & \text{if } s_t = 1 \\ \phi_{0,2} + \sum_{i=1}^{p} \phi_{i,2}\ y_{t-i} + \epsilon_t, & \text{if } s_t = 2 \\ \vdots & \\ \phi_{0,m} + \sum_{i=1}^{p} \phi_{i,m}\ y_{t-i} + \epsilon_t, & \text{if } s_t = m, \end{cases} \tag{3.25}$$

or, in a shorthand version by

$$y_t = \phi_{0,s_t} + \sum_{i=1}^{p} \phi_{i,s_t}\ y_{t-i} + \epsilon_t, \qquad s_t = j, \tag{3.26}$$

where the parameters ϕ_0, and ϕ_i, with $i = 1, \ldots, p$ are estimated for each state j, with $j = 1, \ldots, m$ regimes.

However, as the switching is under the control of a Markov-chain updating mechanism with fixed transition probabilities, the model is only completely characterised after defining these probabilities, which are given by

$$p_{ij} = P(s_t = j | s_{t-1} = i), \quad i, j = 1, \ldots, m, \tag{3.27}$$

with $p_{ij} \geq 0$ and $\sum_{j=1}^{m} p_{ij} = 1$ for all $i, j = 1, \ldots m$. Once again we refer to Engle and Hamilton (1990), Frances and Dijk (2000) and Hamilton (1994) for a more detailed analysis.

3.2.2 Modelling Methodology

In this section we present a brief description of the methodology we use in the modelling and forecasting of the Brent prices when considering regime-switching models.

Like in the methodology regarding long memory time series models, previously discussed in Subsection 3.1.2, we also consider some key steps, much in line with the ones proposed by Frances and Dijk (2000). Therefore we assume the following steps

1. Model identification and regime-switching testing;
2. Model validation; and
3. Forecasting and model selection.

Next we present a brief discussion on each of the above topics.

1. Model identification and regime-switching testing

We start by discussing the model identification for the STAR models as well as for a restriction of this class, the SETAR models. The first step consists on identifying an appropriate order p for the autoregressive model, and the technique used for this procedure has already been explained in Subsection 3.1.2.

We also need to choose the number of regimes, m, and consequently the number of the threshold values c_i, for $i = 1, \ldots, m-1$ states, as well as their estimates. The choice of the threshold values is generally dependent on the number of observations in each regime, so that each state contains a minimum number of observations to obtain reliable estimates of the autoregressive parameters. Therefore the threshold values are usually chosen so that each regime contains at least a specified fraction of observations. There is no clear choice for this fraction but the R software assumes that 15% is reasonable.

Another important aspect is to select the threshold variable, and in this case for the STAR model, we also consider the data itself, lagged by some positive integer δ. The estimate for the delay parameter, δ, is the value that minimises the residual variance.

Regarding the STAR model, it is also necessary to estimate the smoothness parameter, γ. According to Frances and Dijk (2000) the estimate of γ is usually imprecise and, when its significance is analysed, it is usually considered as statistically insignificant. For more details on the identification of SETAR and STAR models we refer to Antonio (2008).

One important and relevant issue regarding regime-switching models, and in particular SETAR models, is to test if the time series really needs to be modelled by non-linear processes. We follow closely the tests described by Hansen (1999) that, generally, considers an appropriate least squares test of the null hypothesis of SETAR(i) against SETAR(k), with $i < k$. The test statistic is given by

$$F_{ik} = n\left(\frac{S_i - S_k}{S_k}\right), \tag{3.28}$$

and is rejected for large values of F_{ik}. The parameter S_i is the residual sum of squares (RSS) of the SETAR(i) model, and S_k is the RSS of the SETAR(k) model. This ratio is equivalent to the likelihood ratio test when the errors are independent and follow a $N(0, \sigma^2)$ process, and is also equivalent to the Lagrange multiplier test. So this test can be used for testing linearity (SETAR(1)) against the SETAR(2) alternative (a non-linear model with two regimes). A large value of F_{12}, as previously mentioned, leads to the rejection of the null hypothesis, but in order to implement this test we must know the distribution of F_{12} under the null hypothesis. In most testing contexts, test statistics like F_{ik}, or in particular F_{12}, are expected to have an asymptotic χ^2

distribution. However, this is not the case, and an asymptotic distribution theory must be developed. In effect, if F_{12} is not statistically significant when compared to the χ^2, it will be certainly not significant when compared to the correct asymptotic distribution. For a more detailed description on this procedure we refer to Hansen (1999). This test is implemented in R software and allows testing for a linear autoregressive model against a SETAR(2) and against a SETAR(3) models, and it also tests the null hypothesis of SETAR(2) against SETAR(3). For further discussion we refer to Antonio (2008) and Hansen (1999).

For the Markov-switching models, we also start by specifying an appropriate order p for the AR(p) model, and this procedure is described in Subsection 3.1.2. As for this class of models, the regimes are determined by an unobservable variable (s_t), we need to choose the number of states, or regimes, for the data under investigation. For each state the autoregressive parameters are estimated using the maximum likelihood techniques and the transition probability from one regime to another is also determined. Additionally, the probability that a process is in a particular state s_t at time t, given all information available at that time is also determined

$$p(s_t|y_1, \ldots, y_t) \tag{3.29}$$

and is defined as the filter inference of the process to be at a certain regime in time t. We also perform an inference about the historical state of the process at some date t, but now based on all given observations,

$$p(s_t|y_1, \ldots, y_T) \tag{3.30}$$

and is referred to as the smoothness inference about the regime at time t.

The Markov-switching model relevance is tested by computing a Likelihood Ratio (LR) statistic, which assumes for H_0 the hypothesis of linearity against the alternative of a relevant MS model. In this case, the test statistic is

$$LR_{MS} = L_{MS} - L_{AR} \tag{3.31}$$

where L_{MS} and L_{AR} correspond, respectively, to the values of the likelihood functions of Markov-switching and linear autoregressive models. As described by Frances and Dijk (2000) this test statistic cannot be characterised analytically, so critical values must be computed in terms of computational simulation.

2. Model validation

After identifying and estimating the model we proceed with the model validation by performing several diagnostic tests. Although the residuals can be subjected to the diagnostic tests proposed and described by Box and Jenkins (1976) (seeSubsection 3.1.2), some test statistics are not suitable in the context of non-linear models.

Regarding the SETAR and STAR models we follow closely the methodology proposed by Frances and Dijk (2000) and test mainly two aspects: serial correlation and remaining non-linearity.

For testing the serial correlation, we need to assume a general non-linear AR(p) model,

$$y_t = F(x_t; \theta) + \epsilon_t, \tag{3.32}$$

where $x_t = (1, y_{t-1}, \dots, y_{t-p})$, and $F(x_t; \theta)$ corresponds to a general non-linear function with parameters θ and twice continuously differentiable. The Lagrange Multiplier (LM) approach is used and its test for the qth order serial dependence in ϵ_t is checked by computing

$$F = nR^2 \sim \chi^2(q) \tag{3.33}$$

where R^2 is the coefficient of determination of $\hat{\epsilon}_t$ on $\partial F(x_t; \theta)/\partial\theta$ and q lagged residuals, $\hat{\epsilon}_{t-1}, \dots, \hat{\epsilon}_{t-q}$.

When testing for remaining non-linearity in models such as STAR and SETAR models, a very simple analysis can be done. It is possible to test if the time series under investigation requires another regime, and the techniques for this evaluation have been previously described when performing the regime-switching testing. For example, if a two-regime model has been identified, we may test the null hypothesis of a two-regime model against the alternative that a third regime is required, and so on.

Considering Markov-switching models, Hamilton (1996) proposes a wide diagnostic checking analysis by testing for residuals autocorrelation, heteroskedasticity, misspecification of the Markov-process s_t, and omitted explanatory variables. The used tests are LM-type tests and make heavy use of the score $h_t(\theta)$ which is defined as

$$h_t(\theta) \equiv \frac{\partial \ln f(y_t|\Omega_{t-1}; \theta)}{\partial\theta} \tag{3.34}$$

where $f(y_t|\Omega_{t-1}; \theta)$ corresponds to the likelihood, or to the conditional density function. For a detailed description on these tests we refer to Frances and Dijk (2000) and

Hamilton (1996).

3. Forecasting and model selection

At last, after identifying and testing several models, we are able to proceed with the forecasting. Analogous to what has been done for the long memory time series models, discussed in Subsection 3.1.2, we consider two different approaches: the in-sample one-step-ahead and the out-of-sample forecasting.

Then the achieved results, obtained for each model, are evaluated considering the methods introduced and discussed in Section 3.3, and the best model is, finally, chosen.

3.3 Model Selection

Analogous to the previous sections, we also consider two subsections. In the first, Subsection 3.3.1, we introduce some concepts and techniques used for the model selection. We refer to simple methods that allow for the evaluation of in-sample and out-of-sample forecasts but we also discuss some formal tests that evaluate the prediction accuracy of two competing forecasts.

In Subsection 3.3.2 we present the key steps followed in the model selection.

3.3.1 Introduction

There are two different approaches that help us to select the best model. The first is based on an in-sample fit evaluation, and the second considers an out-of-sample performing analysis. Regarding the in-sample approach the most appropriate method is based on the information criteria, namely the Akaike's information criterion (AIC), the Bayesian information criterion (BIC), or the Schwartz information criterion (SIC). These model selection criteria are goodness-of-fit measures which compare the in-sample fit, by measuring the residual variance, against the number of estimated parameters.

The other approach is to evaluate the prediction accuracy of the proposed models by considering out-of-sample forecasting. Once again, there are several procedures to check on the quality of the forecasts. The most simple one can be obtained by simply calculating the percentage of the m observations which are within the 95 percent confidence interval (CI). Others commonly used are the root mean square error (RMSE), and the mean absolute percentage error (MAPE) that gives a relative error measure.

These last two statistics compare prediction errors from different models. However a very important issue is being neglected, i.e., testing if there is any qualitative difference between the forecasts of the evaluated models. The original tests of forecasting accuracy imposed certain conditions to the loss function and to the forecast errors that are very difficult to achieve. The loss function has to be quadratic and the forecast errors have to be Gaussian, with null mean, and serially and contemporaneously uncorrelated. To overcome these restrictive limitations, Diebold and Mariano (1995) proposed several tests which test the null hypothesis of no difference in the accuracy of two competing forecasts, but allow a certain flexibility in the loss function, as well as in the forecast errors.

That is, assume two forecasts $\{\hat{y}_{it}\}_{t=1}^{T}$ and $\{\hat{y}_{jt}\}_{t=1}^{T}$ of the time series $\{y_t\}_{t=1}^{T}$ and their respective forecast errors, $\{e_{it}\}_{t=1}^{T}$ and $\{e_{jt}\}_{t=1}^{T}$ with T being the number of predicted values. The idea is to assess the expected loss function (defined as $g(.)$) associated to each of the forecasts. As in many applications the loss function is assumed to be a direct function of the forecast errors, what we want to test for the null hypothesis is if $E[g(e_{it})] = E[g(e_{jt})]$, or simply $E[d_t] = 0$ with $d_t = g(e_{it}) - g(e_{jt})$. In other words, we want to test the null hypothesis that the population mean of the loss differential series is zero.

In order to test the null hypothesis of a zero mean of the loss differential, Diebold and Mariano (1995) assume $\{d_t\}_{t=1}^{T}$ as the sample path of a loss differential series. The asymptotic distribution of the sample mean loss differential can be defined as

$$\sqrt{T}(\bar{d} - \mu) \rightarrow N(0, 2\pi f_d(0)) \tag{3.35}$$

where $\bar{d} = \frac{1}{T}\sum_{t=1}^{T} d_t$, and $f_d(0)$ is the spectral density of the loss differential at frequency zero. Thus, considering large samples, $\bar{d}$ is approximately distributed as $N(\mu, 2\pi f_d(0))$, and for testing the null hypothesis of equal forecast accuracy the test statistic is given by

$$DM = \frac{\bar{d}}{\sqrt{\frac{2\pi \hat{f}_d(0)}{T}}} \sim N(0,1) \tag{3.36}$$

where $\hat{f}_d(0)$ is a consistent estimator of $f_d(0)$. For more details on this test we refer to Diebold and Mariano (1995).

3.3.2 Model Selection Methodology

In this subsection we present the methodology we use for selecting the model that better describes the behaviour of the Brent prices.

As previously mentioned, there are two different approaches to verify the adequacy of the model in terms of forecasting. For the in-sample approach, we compute the information criteria, such as AIC, BIC, and SIC.

For the out-of-sample approach, we consider not only forecast accuracy measures, but we also compare the forecast accuracy of two different models. In order to assess the prediction accuracy we compute the usual accuracy measures, namely the root mean square error (RMSE),

$$\text{RMSE} = \sqrt{\frac{1}{T}\sum_{t=1}^{T}(p_t - \hat{p_t})^2}, \tag{3.37}$$

and the mean absolute percentage error (MAPE),

$$\text{MAPE} = \frac{1}{T}\sum_{t=1}^{T}\frac{|p_t - \hat{p_t}|}{p_t} \times 100\%, \tag{3.38}$$

where T is the number of the predicted observations, and p_t and $\hat{p_t}$ are the actual and the forecasted price, respectively.

Finally, the Diebold-Mariano (DM) test is also computed to compare the forecast accuracy of the different identified models.

Chapter 4

Empirical Results

In this chapter we present all empirical results. In Section 4.1 we start by introducing the Brent prices data set. We analyse its empirical distribution by performing a graphical analysis, and by computing some descriptive statistics, in order to evaluate some location and dispersion measures, as well as its symmetry and kurtosis. The presence of outliers is also tested.

Next we proceed with the data modelling. In Section 4.2 we analyse and forecast the Brent prices using long memory time series models. During the model identification process, we evaluate the long memory phenomena not only in the conditional mean but also in the conditional variance, and we model the data with ARFIMA-FIGARCH models, and some of their variants. In Section 4.3 we characterise and predict the Brent prices behaviour considering regime-switching models, and in this methodology we consider two different approaches. The first considers that the changes in regime occur due to an observable and known variable, and the data is modelled with STAR and SETAR models. The other, assuming that the switching between regimes results from an unobservable variable, considers the application of MS-AR models.

Finally, in Section 4.4 we analyse the forecasting values of the previous identified models, and by performing a prediction accuracy assessment we select the model that better describes and adjusts our Brent prices behaviour. We start with a in-sample fit evaluation by comparing the information criteria of the different identified models, as well as some forecast accuracy measures, namely the RMSE and the MAPE. Then we proceed with an out-of-sample prediction accuracy evaluation, and besides the RMSE and MAPE measures, the Diebold-Mariano test is also computed in order to compare the forecast accuracy of the different models.

We refer that in all sections the study is divided according to the data granularity, meaning that there is a subsection for the daily analysis and another for the monthly Brent prices study.

4.1 Data Analysis

In this section we introduce our time series: the Brent prices data set. First we evaluate its behaviour by presenting its chronogram and by performing a descriptive analysis. The data consists of the *Bloomberg European Dated Forties Oseberg Erofisk Price* and was obtained through a Bloomberg[1] terminal.

As our main purpose is to analyse and understand the Brent prices behaviour, we have used both daily and monthly prices data sets. For the daily data set we consider historical data (on a daily basis and only trading days) from the fourth of January 2000 until the 31st of December 2009, for a total of $n = 2560$ observations. The monthly data, considering the same historical period, is achieved by computing the monthly average of the daily prices, for a total of $n = 120$ observations. However, in the modelling and for estimation purposes, we consider historical data only until September 2009 (for a total of 117 observations) in order to keep an out-of-sample of a year period. Next we present an introduction to each data set. In Subsection 4.1.1 we present the daily Brent prices while in Subsection 4.1.2 we give an overview of the monthly data behaviour.

4.1.1 Daily Brent Prices

First, we start by introducing the daily Brent prices (in USD/Barrel, or \$/bbl) by showing its behaviour in Figure 4.1. From its graphical analysis it is possible to identify

Figure 4.1: Daily Brent prices chronogram (in USD/barrel), from January 4, 2000, until December 31, 2009 (2560 observations).

an increasing trend from January 2000 to July 2006, that leads the Brent prices to reach the value of 76.29\$/bbl. Then a slightly decrease can be observed until January 2007, preceding the most expressive raise that lasts for one year and an half, and that guides

[1]Service that provides financial news and data to some companies and organisations.

the Brent prices to touch its maximum value of 145\$/bbl on the third of July 2008. Afterwards, its price falls again to 34.04\$/bbl at the 24th December 2008, and latter it inverts its trend, starting an increase that remains until the end of December 2009.

To complement the graphical analysis, some descriptive statistics are also provided in Table 4.1. Based on the ratio between the standard deviation and the mean, we evaluated the coefficient of variation which is, approximately, 52%, meaning that the price variation represents about 52% of the mean. Regarding the measures of location, we can verify that the series has a wide range of variation from a minimum of 16.62\$/bbl to a maximum of 145.66\$/bbl, and that 5% of the observations are higher than 101.61\$/bbl. The asymmetry measure (skewness) of 1.10 suggests that the mass of the distribution is concentrated on the left side of the curve, but there are a set of high values that lead to a long right tail. Finally, we have analysed the kurtosis measure which represents the "peakedness" of the distribution. It is a leptokurtic function, as its value is 3.94, indicating that the probability of occurring extreme variations is higher than in a normal distribution (kurtosis is 3).

	Brent prices
Mean	49.83
Std. Deviation	26.06
Skewness	1.10
Kurtosis	3.94
Minimum	16.62
Percentile 0.05	22.55
Median	43.90
Percentile 0.95	101.61
Maximum	145.66

Table 4.1: Descriptive statistics of the daily Brent prices in USD/Barrel, from January 4, 2000, until December 31, 2009.

We follow with a brief analysis on the presence of outliers. Regarding this point, our main goal is simply to identify if there are any severe outliers. The methodology used is based on the interquartile range and assume that an observation is considered outlier if it lays outside the interval $[Q_1 - k(Q_3 - Q_1), Q_3 + k(Q_3 - Q_1)]$, for some constant k. Assuming $k = 3$, and since our series has $Q_1 = 28$ and $Q_3 = 66.23$, we checked all observations and concluded that there is no severe outliers. The presence of outliers is also checked by computing the Grubbs' test described in Komsta (2005), and implemented in the R software, Team (2004). Its null hypothesis assumes that there is no outlier in the data set against an alternative hypothesis that there is at least one observation statistically different from all the others. This test consists of calculating the test statistic G and comparing it to an appropriate critical value, where

$G = \frac{\max_{i=1,\dots,N}|Y_i - \bar{Y}|}{s}$, with $\bar{Y}$ and s denoting the mean and the standard deviation, respectively, and N being the number of observations. Although this test is based on the assumption of normality, and the performed normality tests (see Table A.1) rejected this hypothesis we decided to compute it anyway. The observation in evaluation is the maximum value of the series, 145.66\$/bbl, and as the p-value of the Grubbs' test is 0.30, we do no reject the null hypothesis, and, consequently, the highest value is not considered an outlier, corroborating with the previous result.

4.1.2 Monthly Brent Prices

The behaviour of the monthly Brent prices data set is presented in Figure 4.2. The

Figure 4.2: Monthly Brent prices chronogram (in USD/barrel), from January 2000 until December 2009 (120 observations).

conclusions from the graphical analysis are, as expected, very similar to the ones obtained for the daily data. We stress that the monthly prices are obtained by computing the monthly average of the daily prices. Therefore, an increasing trend is also exhibited from January 2000 to July 2006 followed by a decline for six months. The Brent prices touch its maximum of 134 \$/bbl in July 2008, then decreases, and after December 2008 it increases again till the end of the 2009 year.

Analogous to the analysis implemented for the daily data, we also computed some descriptive statistics. The results, presented in Table 4.2, are practically the same as the ones from the daily data set. The coefficient of variation is still 52%, the leptokurtic characteristic remains, and the mass of the distribution is concentrated on the left side of the curve exhibiting a long right tail.

Regarding the presence of outliers, we have applied the methodology based on the interquartile range, and assuming $k = 3$ no severe outliers were identified. The Grubb's test was also computed, though the hypothesis of normality had been rejected (see

	Brent prices
Mean	49.86
Std. Deviation	25.97
Skewness	1.07
Kurtosis	3.85
Minimum	18.61
Percentile 0.05	23.17
Median	43.38
Percentile 0.95	100.94
Maximum	133.90

Table 4.2: Descriptive statistics of the monthly Brent prices in USD/Barrel, from January 2000 until December 2009.

Table A.2), and as its p-value is 0.06, for a significance level of 5%, the maximum value is not considered as an outlier, supporting the result achieved from the interquartile range methodology.

Next we model the Brent prices data sets considering long memory time series models and regime-switching models.

4.2 Long Memory Time Series Models

In this section we model the Brent prices data, on a daily and monthly basis, using long memory time series models.

Analogous to what has been done for the data analysis, we also consider two distinct subsections. We start by analysing the daily data in Subsection 4.2.1 and in Subsection 4.2.2 we characterise the monthly data set.

4.2.1 Daily Brent Prices

Following closely the methodology described in Subsection 3.1.2 we start with the analysis of the daily Brent prices data set.

1. Stationarity and long memory testing

As previously referred, before testing the long memory presence in our data set, we need to check if the time series has a stationary behaviour. We started to test the variance stationarity by computing the Box-Cox's power transformation. Instead of choosing a grid of values for the λ, we decided only to test the logarithmic hypothesis (as it is the most commonly used in financial markets) against the data set with no transformation. So, for the transformation parameter $\lambda \in \{0, 1\}$, the minimum of the sample variance of the transformed series is obtained for $\lambda = 0$ (the results are shown in Table A.3).

After applying the logarithm to our data set in order to stabilise the variance, we tested if the series has a mean non-stationary behaviour. We computed some tests, such as KPSS, ADF and PP tests, and all results suggest that, for a significance level of $\alpha = 5\%$, the logarithmic transformation is not enough to stabilise the mean of the series. Additionally, we need to apply to it the first order differences. The results, presented in Table A.4, show that after differentiating the data set the time series is finally stationary.

We stress that the presence of any seasonal or cyclic component was also tested by examining the periodogram of the data set and by computing the Fisher test, see Fisher (1929), and no significant component has been identified.

At this moment, we are able to check for the presence of short or long memory on the mean and volatility of our data set. We remark that we have assumed the Brent prices squared returns as a proxy for the Brent prices volatility. An alternative approach might have been the use of the realised volatility, however we leave this analysis for further investigation. In Table 4.3 we present the results obtained when estimating the R/S, GPH and Robinson's d statistics. Although each gives a different estimate value for the long memory parameter, the estimates of the Robinson's d and the GPH (considering a bandwidth of 0.7) tests are very similar for the Brent prices squared returns. However, a very important step is to evaluate if the estimates are statistically significant. Since the GPH test provides not only the estimate but also its standard deviation (contrary to the Robinson's d), we computed the t-test and concluded that, for a significance level of $\alpha = 5\%$, the d estimate for the mean is not statistically significant, while for the volatility it is, and for the further analysis we have considered the value of 0.3152 for the long memory parameter in the volatility of the data set.

	Brent prices returns	Brent prices squared returns
R/S statistic	0.0376	0.1186
Bw=0.6	0.0281	0.4922
GPH Bw=0.7	0.0053	0.3152
Bw=0.8	0.0343	0.1741
Robinson's d	0.0920	0.3159

Table 4.3: Estimates of the long memory parameter, d, for the daily Brent prices.

2. Model identification and validation

After estimating the long memory parameters for the mean and volatility of our data set, we began the model identification. First, and as there is no evidence of long memory in the conditional mean, we computed the empirical autocorrelation and partial autocorrelation functions of our stationary time series, and represented them in Figure 4.3.

(a) (b)

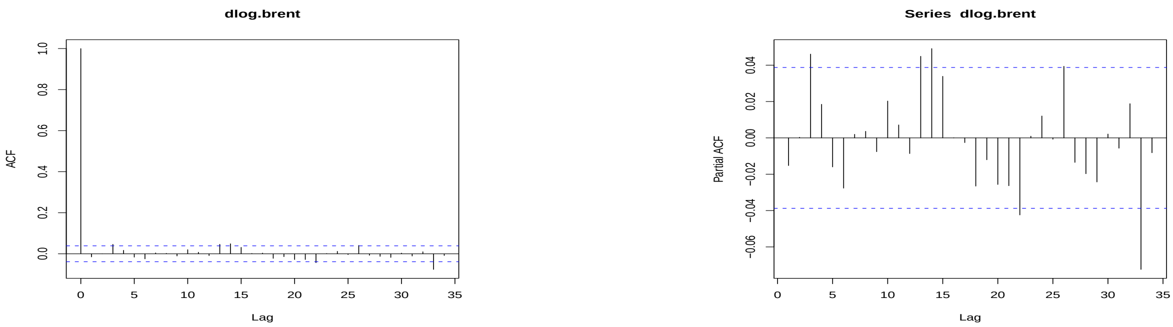

Figure 4.3: Empirical ACF (a) and empirical PACF (b) of daily Brent prices returns.

From the analysis of both functions we conclude that there is no ARMA structure on the mean of the series, so we modelled it with as a WN process with null mean. However, when evaluating the adequacy of the residuals, we concluded that although the residuals are uncorrelated as the Box-Pierce and Ljung-Box tests have confirmed, there is structure in the squared residuals (see Figure A.1). The hypothesis of homokedastic errors was also rejected when we computed the ARCH test. The results of the residuals analysis are shown in Table A.5.

Therefore, as the residuals of the model have a non constant conditional variance and exhibit long memory, we have considered the FIGARCH and some of its variants, namely FIAPARCH and FIEGARCH models, due to the asymmetry of the volatility. We have also modelled the data set without the long memory characteristic for comparison purposes. We remark that in the model estimation we have considered the conditional maximum likelihood estimation method, and as the series of returns rejects the hypothesis of normality we have also assumed a Student's t distribution (see Table A.1). We refer that the estimation of these models, especially the ones considering the long memory characteristic is complex, and sometimes the algorithm failed to converge.

Several models were tested and we present the results of the best model, in terms of the residual analysis, within each class of models. The information regarding the proposed models is presented in Table A.6, and the diagnostic checking for each model is summarised and shown in Table 4.4.

The BP and LB tests correspond to the Box-Pierce and Ljung-Box tests, respectively. The column named as "Var. homog.", refers to the variance homogeneity checking that is done by applying the Box-Cox's power transformation, and in case no trans-

Model	Param. Signif.	Model validation: Residuals BP test	LB test	Var. homog.	Squared residuals BP test	LB test
GARCH(1,1)	✓	0.861	0.860	✓	0.718	0.717
APARCH(2,2)	✓	0.712	0.711	✓	0.912	0.912
EGARCH(1,1)	✓	0.725	0.725	✓	0.575	0.575
FIGARCH(1,1)	✓	0.794	0.794	✓	0.717	0.717
FIAPARCH(3,3)	✓	0.841	0.841	✓	0.889	0.889
FIEGARCH(2,2)	✓	0.340	0.340	✓	0.048	0.048

Table 4.4: Summary results of the proposed models for the daily Brent prices.

formation is needed (and consequently the variance is stabilised), we simply represent it using a check mark.

From the model validation analysis, we verify that modelling the daily Brent prices with time series models without the long memory characteristic seems to be enough to describe the returns behaviour, at least in terms of validation. In effect, all models passed the diagnostic checking, except the FIEGARCH(2,2) model that was not able to remove the autocorrelation among the squared residuals for the lower lags. An additional examination regarding the models validation is shown from Figure A.2 to Figure A.7. Next we followed to the forecasting and model selection, and we considered all identified models except the FIEGARCH(2,2) model.

3. Forecasting and model selection

As we have several identified and validated models, we followed with the models forecasting. Although we have computed the predicted values, we decided only to present the results in Subsection 4.4.1, where a prediction accuracy assessment is performed by comparing the results of all identified daily models.

4.2.2 Monthly Brent Prices

Analogous to the procedure followed for the daily data set, in the modelling of the monthly Brent prices we also consider the modelling methodology described in Subsection 3.1.2. We start by evaluating the stationary behaviour and the long memory in our data set.

1. Stationarity and long memory testing

The first step to follow is to verify the variance homogeneity, and we considered the same two values for the transformation parameter, $\lambda \in \{0, 1\}$. The results, presented in Table A.7, are very similar to the ones obtained for the daily data set, as they also suggest a logarithmic transformation.

Analogous to the daily results, and since no periodic or seasonal component were identified, we have computed the KPSS, PP and ADF tests to check for the series stationarity and concluded that we needed to apply to the series the first order differences, see Table A.8.

As our data set is now stationary, we are able to proceed with the short and long memory testing. The R/S, GPH and Robinson's d statistics were computed and the estimates of the d parameter are given in Table 4.5. In this case, the results are quite

	Brent prices returns	Brent prices squared returns
R/S statistic	0.0736	-0.2528
Bw=0.6	-0.0610	0.1208
GPH Bw=0.7	0.1037	0.2044
Bw=0.8	0.1221	0.4249
Robinson's d	0.1605	0.0689

Table 4.5: Estimates of the long memory parameter, d, for the monthly Brent prices.

different from the ones obtained for the daily data, as we are not able to identify similar values for two different statistics, which makes it harder to choose the right estimate. Additionally, the analysis of the different d estimates for the mean and volatility show simultaneously negative and positive values depending on the performed statistic, which may lead to contradictory conclusions.

The estimate significance was also tested, and for the conditional mean none of the estimates was statistically significant. Regarding the conditional variance, both results for the 0.7 and 0.8 bandwidths (in the GPH test) revealed to be statistically significant, but as they are quite different and there is no obvious reason to choose one against the other, the choice of the d parameter will be done during the model identification.

2. Model identification and validation

We are now able to proceed to the model identification. As there is no evidence of long memory in the conditional mean, we computed the empirical ACF and PACF of the stationary series in Figure 4.4, in order to identify the order of the autoregressive and moving average parameters. From the analysis of both empirical autocorrelation functions, we verify that both exhibit a higher value for the sixth lag, so we adjusted an AR(6) to our stationary time series. We stress that, in the monthly modelling we have considered historical data only until the end of September 2009, for a total of 117 observations, in order to keep an out-of-sample of twelve points, meaning that a year will be forecasted.

The identified model, an AR(6), passed all diagnostic tests. The ACF and PACF of its residuals were computed and shown in Figure A.8, and they illustrate the behaviour

(a) (b)

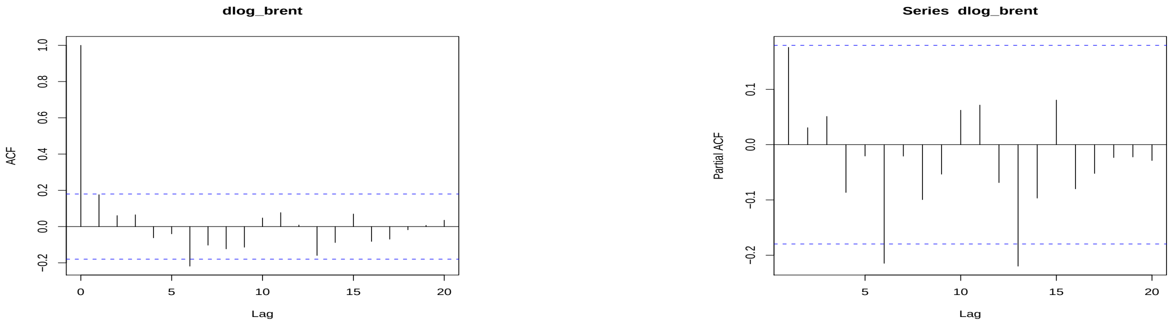

Figure 4.4: Empirical ACF (a) and empirical PACF (b) of the monthly Brent prices returns.

of a WN process. This behaviour is also confirmed by the results of some diagnostic tests, namely the Box-Pierce and Ljung-Box tests and the ARCH LM test, that suggest no correlation among the residuals and no conditional heteroskedasticity, respectively (see Table A.9). Although the results suggest no ARCH effect, we tried to force a conditional heteroskedastic model to the volatility in order to include the long memory characteristic (using the previous estimates, as well as allowing the model to estimate itself the d parameter). However, and as expected, none of the tried models was statistically significant, meaning that the monthly Brent prices data set has no long-range dependence, neither in the conditional mean, nor in the variance. So, the only proposed model for the monthly Brent prices returns is the AR(6) model (described in Table A.10), and its results are summarised in Table 4.6

		Model validation				
	Param.	Residuals			Squared residuals	
Model	Signif.	BP test	LB test	Var. homog.	BP test	LB test
AR(6)	✓	0.952	0.951	✓	0.176	0.170

Table 4.6: Summary results of the proposed model for the monthly Brent prices returns.

3. Forecasting and model selection

Now, given the identified model we proceeded with the models forecasting. Although we have computed the predicted values, we decided, like in the daily case, only to present the results in Subsection 4.4.2, where a prediction accuracy assessment of all identified monthly models is done.

4.3 Regime-switching Models

In this section we present the Brent prices modelling by considering the regime-switching models. Similar to the previous sections, we have also divided the analysis and modelling of the Brent prices in two different subsections, according to the data granularity. Thus, in Subsection 4.3.1 we perform the daily analysis, and in Subsection 4.3.2 we characterise the Brent prices behaviour considering monthly data.

4.3.1 Daily Brent Prices

Considering the steps previously described in Subsection 3.2.2, we started with the model identification.

1. Model identification and regime-switching testing

As previously referred, first we need to identify an appropriate order for the autoregressive parameter. So we started by running several autoregressive models, and the choice of the p parameter corresponds to the model for which the information criterion AIC is minimum, and in this case we identified an AR(1) model. Before we passed on to the other parameters estimation, we have confirmed that modelling the data with an AR(1) model was enough to remove all data autocorrelation.

As the threshold variable is assumed to be the data itself, as previously mentioned, we followed with the estimation of the delay parameter, δ. We considered a grid of values for this parameter, $\delta \in \{1, 2, \dots, 9\}$, and the first models we tried to adjust to the daily Brent prices were the LSTAR models, but as the estimated values for the smoothness parameter were too high (see Table A.11) we decided to discard these models. We stress that for high values of the smoothness parameter (γ), the transition between regimes is practically instantaneous, and as our threshold variable is the data itself, these models are in effect SETAR models. Therefore, we continued this analysis considering only this last class of models.

Once again, we assumed the same grid of values for the delay parameter and estimated the corresponding SETAR(1) models. The results are shown in Table 4.7, and the minimum residual variance is achieved for $\delta = 7$ and $\delta = 8$. We decided to consider $\delta = 7$.

We proceeded with the regime-switching testing, and the results illustrated in Table A.12 show that the hypothesis of linearity is not rejected, meaning that there is no need to model the daily data with regime-switching models. Even though the achieved conclusion for the SETAR models, we have continued the analysis and tried to adjust a MS-AR(1) model to the daily Brent prices. The model was estimated in the TSM

δ	Residual variance	Threshold value	% Low regime	% High (observations)
1	1.678	35.02	43%	57%
2	1.677	35.97	44%	56%
3	1.677	34.83	43%	57%
4	1.678	33.30	41%	59%
5	1.677	36.16	44%	56%
6	1.676	36.16	44%	56%
7	1.675	38.19	45%	55%
8	1.675	37.95	45%	55%
9	1.676	39.40	46%	54%

Table 4.7: Summary results of the two-regime SETAR models for the daily Brent prices.

software, considering the conditional maximum likelihood estimation method and using the Student's t distribution (as the normality hypothesis is rejected for the daily data, see Table A.1). The results were very interesting, though not surprising, as they corroborate with the SETAR model conclusions: the identified model has also rejected the need for modelling the daily data with regime-switching models. Instead of computing the LR test, we checked if the estimated values for the parameters of the second regime were statistically different from the ones of the regime one. However, since the results showed that they were not statistically different, we could not reject the hypothesis of linearity.

Given the previous results, we conclude that there is no improvement in modelling the daily Brent prices with regime-switching models, as both regimes are statistically identical, and consequently the data should be modelled with linear models.

4.3.2 Monthly Brent Prices

Analogous to the procedure followed for the daily data, we started with the model identification and the regime-switching testing.

1. Model identification and regime-switching testing

We started by computing several models to identify the order of the AR(p) model, and the model that presented the lowest AIC was an AR(2) model. After confirming that the correlation among the residuals has been removed, we continued with the other parameters estimation.

Since we are considering as threshold variable the data itself, we only need to estimate the delay parameter δ. First, we tried to apply LSTAR models, but as the achieved values for the smoothness parameter γ were too high (see Table A.13), and consequently the regime transition was practically instantaneous, we decided to use only SETAR models. The results for the SETAR models are shown in Table 4.8. For each model, we have assumed an AR(2) model and for the following grid of delay

	δ	Residual variance	Threshold value(s)		% Low	% Medium	% High
					regime (observations)		
2 Regimes	1	20.79	71.32		84%	-	16%
	2	18.75	71.32		85%	-	15%
	3	18.02	70.19		82%	-	18%
	4	18.44	70.80		84%	-	16%
3 Regimes	1	22.63	43.03	58.48	49%	19%	32%
	2	17.49	33.36	71.32	44%	41%	15%
	3	16.69	48.67	70.19	58%	25%	17%
	4	16.92	43.03	70.80	51%	33%	16%

Table 4.8: Summary results of the SETAR models for the monthly Brent prices.

values, $\delta \in \{1, 2, 3, 4\}$, we computed the residual variance, the threshold value or values (depending if we are considering a two or a three-regime model), and the percentage of observations in each regime. As the residual variance is minimum for $\delta = 3$, we have considered a three month delay in the SETAR modelling.

Next, in Table A.14 we present all tests regarding the number of regimes in the SETAR models. We started by testing the hypothesis of linearity against the alternatives SETAR(2) or SETAR(3), and the null hypothesis is clearly rejected meaning that the monthly data should be modelled with regime-switching models. The hypothesis of a two-regime model (SETAR(2)) against a three-regime model (SETAR(3)) was also tested, and in this case we do not reject the null hypothesis, and consequently, a two-regime model must be considered.

We embarked on the modelling of the MS-AR model. We started to consider an AR(2) model for each regime, and then using the TSM software, we adjusted a MS-AR(2) model to the monthly Brent prices. In the model estimation, we have considered the conditional maximum likelihood estimation method and a Student's t distribution, as the monthly data rejects the normality hypothesis (see Table A.2). Thus, for each state the TSM program estimated the parameters and the transition probability was also calculated.

Regarding the Markov-switching testing, and in order to test the hypothesis of linearity against a two-regime model, we considered the same procedure used for the daily data. Instead of computing the LR statistic, we computed the test that allowed to verify if the parameters of higher order regimes were different from the linear (one-regime) model, and as they revealed to be statistically different, we concluded that the data should be modelled with a two-regime model. The hypothesis of another state was also tried, however, due to estimation problems as the algorithm failed to converge, we were not able to estimate a three-regime model. Therefore, we assume that our final model is a two-regime MS-AR(2) model, in line with the two-regime SETAR(2) model previously identified.

So, to summarise we identified two different models: one assuming that the monthly Brent prices data set depends on itself with a three month lag, and the other considering that it depends on an external and unknown variable. Both models, described in Table A.16 follow an AR(2) process and are a two-regime models.

2. Model validation

We proceeded with the model validation by testing the serial correlation, and if there is any remaining non-linearity.

First we analysed the serial correlation of the two-regime SETAR(2) model. However, as pointed by Frances and Dijk (2000), the SETAR models do not satisfy the requirement of the linear function $F(x_t;\theta)$ being twice continuously differentiable, and the LM statistic cannot be applied to the residuals from SETAR models. As a solution, we considered testing the correlation among the residuals in each regime by computing the Box-Pierce and Ljung-Box tests, and the results suggest no correlation. The test for remaining non-linearity has already been done. In the previous point, when we were estimating the number of regimes, we tested the null hypothesis of a two-regime model against the alternative that a third regime is required. In effect, we were testing if with a two regime model there was any remaining non-linearity, and as the null hypothesis is not rejected we can conclude that there is no remaining non-linearity.

The model validation for the MS-AR(2) model is similar to the one followed by the SETAR models. The Box-Pierce and Ljung-Box tests were computed and their results indicate that there is no serial correlation when testing the MS-AR(2) model residuals. The analysis of remaining non-linearity, as previously mentioned, could not be done due to estimation problems. We refer to the difficulty in estimating non-linear models, as their estimating process involves complex mathematical calculations and requires advanced mathematical techniques. Nonetheless we assumed the MS-AR(2) model as our final model.

We stress that all results regarding the model validation for both models are shown in Table A.15.

At last, and after having the two proposed models validated we present their characteristics in Table A.16. Additionally, we have computed some statistics in order to compare the results of both modelling approaches. The comparison, illustrated in Table 4.9, is quite interesting as it indicates a great similarity between the results of the two models. Both models show a clear difference between the two identified regimes. While the first regime, represented by the majority of observations, presents the lowest Brent prices and the highest variation, the other regime has the opposite characteristics, i.e., it exhibits the highest Brent prices and the lowest variation. Regarding each state duration, regime one lasts for more than one year (almost two years in the SETAR(2) model case), and the second state has a duration of only a few months.

	SETAR(2) model		MS-AR(2) model	
	Regime 1	Regime 2	Regime 1	Regime 2
Mean [\$/bbl]	41.55	87.56	47.26	60.61
Std. Deviation [\$/bbl]	17.21	27.70	26.30	23.33
Coef. variation [%]	41%	32%	55%	38%
% Observations	82%	18%	88%	13%
Average duration [months]	23.5	6.7	12.3	2.0

Table 4.9: Statistics for the monthly regime-switching models.

The analysis and comparison of the regime transition is also done by comparing the smoothed probabilities of the MS-AR(2) model with the regime of the SETAR(2) model in which the time series is at a certain point in time, see Figure 4.5. The

(a) (b)

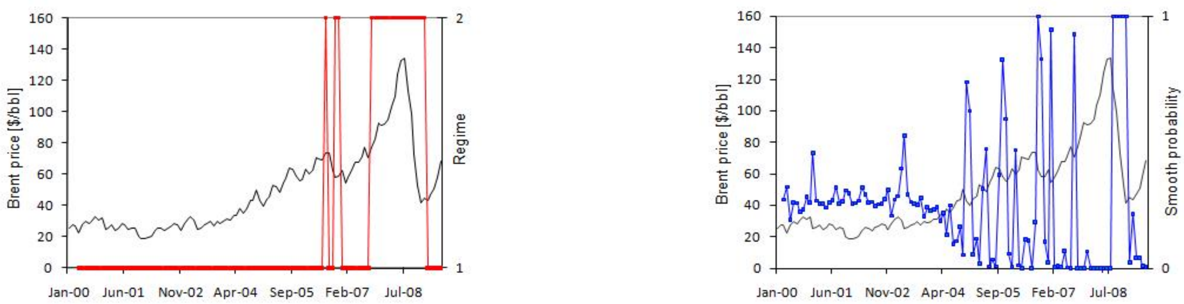

Figure 4.5: Regime of the SETAR(2) model (a) and smoothed transition probability of the MS-AR(2) model (b) for the monthly Brent prices.

SETAR(2) model seems to capture quite well the highest peaks of the monthly Brent prices, identifying them as part of the second regime and maintaining the remaining observations in regime one. The smoothed probabilities estimated by the MS-AR(2) model are also able to identify the Brent prices peaks, however with a lag difference. For example, when analysing the maximum of the Brent prices we verify that the smoothed probability reaches the value one only after the maximum occurs. The results of these two models in terms of fitting are compared in the next section.

3. Forecasting and model selection

Analogous to the procedure followed for all previous models, the assessment of the prediction accuracy of the proposed regime-switching models is presented in Subsection 4.4.2.

4.4 Model Selection

In this section we select the final models for the daily and monthly Brent prices. We start by performing an in-sample analysis, comparing the information criteria of the different models, and some forecast accuracy measures, namely the RMSE and the MAPE. Then we follow with an out-of-sample prediction accuracy evaluation by computing the RMSE and MAPE measures, and by performing the Diebold-Mariano test in order to compare the forecast accuracy of the different models. In terms of organisation, the daily results are presented in Subsection 4.4.1, and in Subsection 4.4.2 we select the final model for the monthly Brent prices.

4.4.1 Daily Brent Prices

Considering the methodology previously described, we start by presenting some results regarding the in-sample analysis. We recall that we have only analysed the time series models, as we were not able to identify any regime-switching model for the daily Brent prices. Therefore, in Table 4.10, we compare the results of the proposed models for the daily Brent prices, presenting the information criteria, and some measures of prediction accuracy. From the information criteria analysis, we verify that the intro-

	Inf. criteria		For. measures	
Model	AIC	SIC	RMSE	MAPE
GARCH(1,1)	6043	6028	1.002	27.32%
APARCH(2,2)	6039	6013	1.001	25.01%
EGARCH(1,1)	6037	6019	1.003	29.07%
FIGARCH(1,1)	6037	6023	0.997	30.02%
FIAPARCH(3,3)	6032	5999	0.996	30.39%

Table 4.10: In-sample evaluation for the daily models.

duction of the long memory characteristic has a positive impact on both GARCH(1,1) and APARCH(2,2), as the AIC and SIC measures decreases. This idea is corroborated with the RMSE values, though the conclusions when checking the MAPE are contradictory.

So, we continued the forecasting and model selection analysis, but now we evaluated the out-of-sample forecasts. The results, illustrated in Table 4.11 show that the RSME and MAPE values for the proposed models are very similar.

	GARCH(1,1)	APARCH(2,2)	EGARCH(1,1)	FIGARCH(1,1)	FIAPARCH(3,3)
RMSE	5.699	5.432	5.494	5.795	5.773
MAPE	6.79%	6.46%	6.54%	6.90%	6.88%

Table 4.11: Out-of-sample forecast accuracy evaluation for the daily models.

Moreover, we computed the Diebold-Mariano (DM) test, in order to compare the forecast accuracy of the different prediction samples. This test, implemented in the R software by Hyndman (2010), was computed for the proposed different models, and the results are shown in Table A.17. We can conclude that although the models have very similar prediction error values, their forecasting samples are statistically different. When we compare the values of the DM test statistic for the proposed models we verify that the models that incorporate the long memory characteristic show a better predictive ability. Therefore, and to conclude, we decided to select the FIGARCH(1,1) model as our final daily model, since it is the one that presents the highest predictive accuracy, according to the results of the Diebold-Mariano test.

4.4.2 Monthly Brent Prices

We now present the results for the monthly Brent prices. Analogous to the procedure for the daily data, we started with the in-sample analysis and in Table 4.12 we illustrate the obtained results. The information criteria measures, unlike the daily case, are not

	Inf. criteria		For. measures	
Model	AIC	SIC	RMSE	MAPE
AR(6)	95	83	5.298	7.92%
SETAR(2)	369	-	4.301	7.01%
MS-AR(2)	-344	-358	3.104	5.71%

Table 4.12: In-sample evaluation for the monthly models.

very useful in this context, as we are comparing models of different classes. Regarding the forecast accuracy measures and though the results are very similar, we verify that the MS-AR(2) model is the one that presents the best in-sample fitting.

Next, we followed with the out-of-sample analysis, whose results are shown in Table 4.13. We have compared the predicted values from the three models with the real monthly prices of Brent for October 2009 until October 2010, for a one year and a month period. Once again, the results are very similar. Finally, and in order to check

	AR(6)	SETAR(2)	MS-AR(2)
RMSE	13.390	19.578	13.172
MAPE	16.49%	19.78%	16.36%

Table 4.13: Out-of-sample forecast accuracy evaluation for the monthly models.

if the differences from the three forecasting samples are statistically significant, we computed the Diebold-Mariano (DM) test. We applied the DM test implemented in R (see Hyndman (2010)) and the results, presented in Table A.18, indicate that the null hypothesis is always rejected, meaning that there is a difference statistically significant among the three forecasting samples.

Therefore, we chose the SETAR(2) model as our monthly final model since it is the one that, according to the results from the Diebold-Mariano test, has the highest predictive accuracy. We also remark that this model is very interesting as it considers two different regimes depending on the Brent prices characteristics.

Chapter 5

Conclusion

During the past years, the Brent commodity has become the most important source of energy in the world, and due to its impact on society it has turned into an issue of major concern. Therefore, the aim of this work is to analyse, understand and forecast the Brent prices behaviour. To model and predict the Brent prices we considered data on a daily and monthly basis, and we used long memory time series models and regime-switching models.

For the daily Brent prices analysis we considered historical data from the fourth of January 2000 until the 31st of December 2009, for a total of 2560 observations. The results show that the daily data set should be modelled with long memory time series models, as there is evidence of long-range dependence on the conditional variance of the Brent prices. The selected model is the FIGARCH(1,1) applied to the Brent prices returns since it is the one, according to the Diebold-Mariano test, that presents the highest predictive accuracy. To the daily data we also applied regime-switching models, but the results suggest that the data should be modelled by linear models, as the hypothesis of non-linearity is rejected.

The monthly Brent prices analysis was done considering the monthly average of the daily prices from January 2000 until September 2009, for a total of 117 observations. The reason why we did not use the data until the end of the year, as we have done for the daily case, was to guarantee an out-of-sample period of one year. The conclusions for the monthly Brent prices analysis are quite different from the ones obtained for the daily data. Regarding time series models, we identified an AR(6) model for the monthly returns, and neither conditional heterokedasticity nor long-range dependence were identified. The regime-switching models were also applied and we chose the two-regime SETAR(2) model as, according to the DM test, it is the one with the highest predictive accuracy. This model is very interesting as it shows clear differences between the two identified regimes. While the first regime, represented by the majority of observations, exhibits the lowest Brent prices and the highest variation, the other regime has the opposite characteristics, i.e., the highest Brent prices and the lowest variation. Besides, regime one lasts for almost two years while the second state has a duration of only a few months.

In this analysis we were not able to identify a strong relation between stochastic regime-switching and long memory. Nonetheless, we think that for a future investigation a very interesting topic to analyse is the application of regime-switching models to the conditional variance of the daily Brent prices.

Finally, we present the main contributions of this thesis for the econometric and financial literature concerning the Brent prices modelling and forecasting.

First, to model and predict the Brent prices we consider long memory time series models and regime-switching models and compare their predictive accuracy. Concerning time series models, several models are used to capture the long memory in the Brent prices returns and some of them, namely FIAPARCH and FIEGARCH models are not so commonly used in the econometric literature. We note that the long memory characteristic is tested simultaneously in the conditional mean and variance. For regime-switching models we analyse and compare if the Brent prices are better described by its past behaviour or by an unobservable and unknown variable.

Another interesting issue is the parallel analysis made for both daily and monthly Brent prices that leads to very different and interesting conclusions. At last, and in order to capture the models prediction accuracy we do not use just the traditional criteria, such as MAPE and RMSE measures, but we also apply the Diebold-Mariano test to check if the differences in the forecast accuracy are statistically significant.

Bibliography

Antonio, F. (2008, March). Nonlinear autoregressive time series models in R using tsDyn. Version 0.7. http://cran.r-project.org/web/packages/tsDyn.

Baillie, R. (1996). Long memory processes and fractional integration in econometrics. *Journal of Econometrics 73*, 5–59.

Baillie, R., T. Bollerslev, and H. Mikkelsen (1996, September). Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics 74*, 3–30.

Baillie, R. and C. Chung (1996). Analysing inflation by the fractionally integrated ARFIMA-GARCH model. *Journal of Applied Econometrics 11*, 23–40.

Banerjee, A. and G. Urga (2005). Modelling structural breaks, long memory and stock market volatility: an overview. *Journal of Econometrics 129*, 1–34.

Beran, J. (1994). *Statistics for Long-Memory Processes.* United States of America: Chapman and Hall/CRC.

Black, F. (1976). Studies in stock price volatility changes. In Business and E. S. Section (Eds.), *American Statistical Association*, pp. 177–181.

Boero, G. and E. Marrocu (2004, April-June). The performance of SETAR models: a regime conditional evaluation of point, interval and density forecasts. *International Journal of Forecasting 20*(2), 305–320.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics 31*, 307–327.

Bollerslev, T. and H. Mikkelsen (1996, September). Modeling and pricing long memory in stock market volatility. *Journal of Econometrics 73*, 151–184.

Box, G. and D. Cox (1964). An analysis of transformations. *Journal of the Royal Statistical Society 2*(26), 211–252.

Box, G. and G. Jenkins (1976). *Time Series Analysis: forecasting and control* (Revised Edition ed.). Oakland, California, USA: Holden-Day.

Braun, S. (2010, June). *Memory diagnostic in time series analysis.* Ph. D. thesis, Faculty of Verhaltens- und Empirische Kulturwissenschaften of the Ruprecht-Karls-Universität Heidelberg.

Breidt, F., N. Crato, and P. Lima (1998). The detection and estimation of long memory in stochastic volatility. *Journal of Econometrics 73*, 325–348.

Brunetti, C. (1999). Long memory, the "Taylor effect" and intraday volatility in commodity futures markets. Technical report, Department of Business Studies. The University of Edinburgh.

Caraiani, P. (2010). Modelling business cycles in the Romanian economy using Markov Switching approach. *Romanian Journal of Economic Forecasting 1*, 130–136.

Chaouachi, M. (2005, September). Detection de la memoire longue dans la chronique de petrole. Technical report, Laboratoire: L'E.R.M.E. Equipe de Recherche sur les Marchés, lÉmploi et la Simulation affilié au CNRS; Université Panthéon-Assas Paris-II.

Chen, Z. and P. Forsyth (2010). Implications of a regime-switching model on natural gas storage valuation and optimal operation. *Quantitative Finance 10*, 159–176.

Dafas, P. (2004, August). Estimating parameters of a mean-reverting Markov-switching jump-diffusion model for crude oil spot prices. Working paper. School of economics and statistics.

Davidson, J. (2010, January). *TSM - Time Series Modelling Version 4.31.* University of Exeter.

Diebold, F. and A. Inoue (2001). Long memory and regime switching. *Journal of Econometrics 105*, 131–159.

Diebold, F. and R. Mariano (1995, July). Comparing predictive accuracy. *Journal of Business & Economic Statistics 13*(3), 134–144.

Dijk, D., T. Teräsvirta, and P. Frances (2002). Smooth transition autoregressive models - a survey of recent developments. *Econometric Reviews 21*(1), 1–47.

Engle, C. and J. Hamilton (1990, September). Long swings in the dollar: are they in the data and do markets know it? *The American Economic Review 80*(4), 689–713.

Engle, R. (1982, July). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica 50*(4), 987–1008.

Fisher, R. (1929). Tests of significance in harmonic analysis. *Proc. Roy. Soc. A125*, 54–59.

Frances, P. and D. Dijk (2000). *Non-linear time series models in empirical finance.* Cambridge University Press.

Geweke, J. and S. Porter-Hudak (1983). The estimation and application of long-memory time series models. *Journal of Time Series Analysis 4*, 221–238.

Granger, C. and N. Hyung (2004). Occasional structural breaks and long memory with an application to the S&P 500 absolute stock returns. *Journal of Empirical Finance 11*, 399–421.

Grau-Carles, P. (2005). Tests of long memory: a bootstrap approach. *Computational Economics 25*, 103–113.

Hamilton, J. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica 57*, 39–70.

Hamilton, J. (1994). *Time Series Analysis.* Princeton, New Jersey: Princeton University Press.

Hamilton, J. (1996). Specification testing in Markov-switching time series models. *Journal of Econometrics 70*, 127–157.

Hansen, B. (1999). Testing for linearity. *Journal of Economic Surveys 13*(5), 551–576.

Hardy, M. (2001, April). A regime-switching model of long-term stock returns. *North American Actuarial Journal 5*(2), 41–53.

Hosking, J. (1981, April). Fractional differencing. *Biometrica 68*(1), 165–176.

Hurst, H. (1951). Long term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineer 116*, 770–799.

Hyndman, R. (2010, October). Package forecast. Version 2.09. http://cran.r-project.org/web/packages/forecast.

Ismail, M. and Z. Isa (2006). Modelling exchange rates using regime-switching models. *Sains Malaysiana 35*(2), 55–62.

Komsta, L. (2005, August). The outliers package: Tests for outliers. Version 0.12. http://cran.r-project.org/web/packages/outliers.

Korkmaz, T., E. Çevik, and N. Özataç (2009). Testing for long memory in ISE using ARFIMA-FIGARCH model and structural break test. *International Research Journal of Finance and Economics 26*, 186–191.

Lo, A. (1991). Long-term memory in stock market prices. *Econometrica 59*, 1279–1313.

Mandelbrot, B. B. and J. R. Wallis (1968, October). Noah, Joseph, and operational hydrology. *Water Resources Research 4*(5), 909–918.

McLeod, A. and K. Hipel (1978, June). Preservation of the rescaled adjusted range, 1. A reassessment of the Hurst phenomenon. *Water Resources Research 14*(3), 159–178.

Medeiros, M. and M. Scharth (2009). Asymmetric effects and long memory in the volatility of Dow Jones stocks. *International Journal of Forecasting 25*, 304–327.

Nelson, D. (1991, March). Conditional heteroskedasticity in asset returns: a new approach. *Econometrica 59*(2), 347–370.

Pires, A. (2001, March). Notas de séries temporais 2000/01. https://dspace.ist.utl.pt/bitstream/2295/53453/1/Notas.pdf.

Robinson (2003). *Time Series with Long Memory.* Oxford University Press.

Robinson, P. (1994, March). Semiparametric analysis of long-memory time series. *The Annals of Statistics 22*(1), 515–539.

Robinson, P. (1995, June). Log-periodogram regression of time series with long range dependence. *The Annals of Statistics 23*(3), 1048–1072.

Schindlmayr, G. (2005, May). A regime-switching model for electricity spot prices. Working paper. EnBw Trading GmbH.

Software, Q. M. (1994-2009). *EViews 7 User's Guide I.* United States of America: Quantitative Micro Software, LCC.

Team, T. R. D. C. (2004, November). R: A language and environment for statistical computing. Version 2.0.1. http://cran.r-project.org/doc/manuals.

Teräsvirta, T. (1994). Specification, estimation, and evaluation of smooth transition autoregressive models. *Journal of the American Statistical Association 89*, 208–218.

Tong, H. (1983). *Threshold Models in Non-linear Time Series Analysis.* Springer-Verlag.

Trapletti, A. (2005, October). The tseries package: Time series analysis and computational finance. Version 0.10-0. http://cran.r-project.org/web/packages/tseries.

Tse, Y. (1998, Jan-Feb). The conditional heteroscedasticity of the yen-dollar exchange rate. *Journal of Applied Econometrics 13*(1), 49–55.

Viswanathan, G., S. Buldyrev, S. HavIin, and H. Stanley (1997, February). Quantification of DNA patchiness using long-range correlation measures. *Biophysical Journal 72*, 866–87.

Wei, W. (1990). *Time Series Analysis* (First ed.). California, Massachusetts, New York, Amsterdam, Ontario, Sydney, Madrid, Bonn, Singapore, Tokyo, San Juan, United Kingdom: Addison-Wesley Publishing Company, Inc.

Wuertz, D. (2009, September). ARMA time series modelling. 2100.76. http://cran.r-project.org/web/packages/fArma.

Appendix A

Empirical Results

A.1 Data Analysis

Test	p-value	
	Brent	Brent returns
Jarque-Bera	0.00	0.00
Anderson-Darling	0.00	0.00
Lilliefors (Kolmogorov-Smirnov)	0.00	0.00
Pearson chi-square	0.00	0.00
Shapiro-Francia	0.00	0.00

Table A.1: Normality tests for the daily data.

Test	p-value	
	Brent	Brent returns
Jarque-Bera	0.00	0.00
Anderson-Darling	0.00	0.00
Lilliefors (Kolmogorov-Smirnov)	0.00	0.00
Pearson chi-square	0.00	0.00
Shapiro-Francia	0.00	0.00

Table A.2: Normality tests for the monthly data.

A.2 Long Memory Time Series Models

λ	Sample variance
0	634
1	1737980

Table A.3: Sample variance in the Box-Cox's power transformation for the daily data.

Test	H_0		p-value	
		Brent	ln(Brent)	Returns
KPSS	Stat.	0.01	0.01	0.10
ADF	Non-Stat.	0.62	0.29	0.01
PP	Non-Stat.	0.40	0.42	0.01

Table A.4: Stationarity testing of the daily Brent prices.

(a) (b)

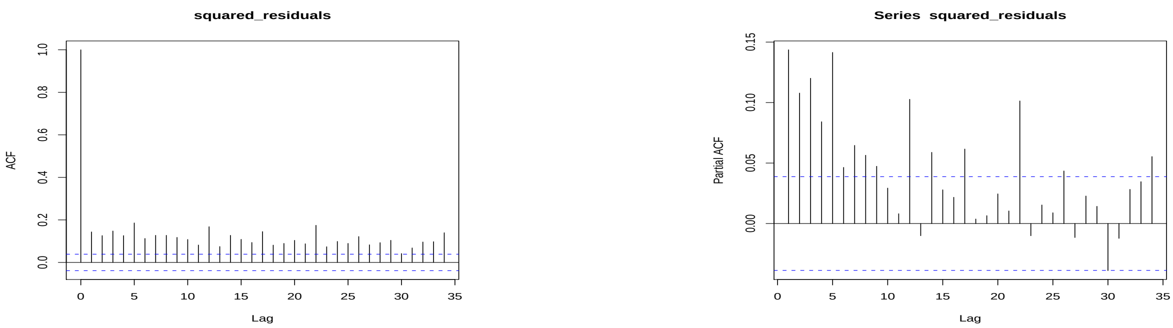

Figure A.1: Empirical ACF (a) and empirical PACF (b) of the squared residuals of the WN model applied to the daily Brent prices returns.

	Test	p-value
Residuals	Box-Pierce	0.440
	Ljung-Box	0.441
	ARCH	0.000
Squared residuals	Box-Pierce	0.000
	Ljung-Box	0.000

Table A.5: Dignostic testing of the residuals of the WN model applied to the daily Brent prices returns.

		Estimate	Std. Err.	t Ratio	p-value
GARCH(1,1)	const	0.0012	0.0004	2.804	0.005
	α_0	0.0122	0.0012		
	α_1	0.0495	0.0114	4.349	0
	β_1	0.9329	0.0158	58.969	0
APARCH(2,2)	const	0.0010	0.0004	2.351	0.019
	α_0	0.0175	0.0138		
	Asymmetry	1.0065	0.5899	1.706	0.088
	Power	1.7970	0.4233		
	α_1	0.0337	0.0163	2.063	0.039
	α_2	0.0273	0.0128	2.142	0.032
	β_2	0.8695	0.0677	12.842	0
EGARCH(1,1)	const	0.0011	0.0005	2.298	0.022
	α_0	10.6974	1.9017		
	Asymmetry	-0.3786	0.1770	-2.139	0.033
	α_1	0.1297	0.0451	2.874	0.004
	β_1	0.9683	0.0114	84.645	0
FIGARCH(1,1)	const	0.0013	0.0004	2.966	0.003
	α_0	0.0092	0.0008		
	d	0.3152	Fixed		
	α_1	-0.2589	0.0238	-10.867	0
	β_1	0.5442	0.0495	10.996	0
FIAPARCH(3,3)	const	0.0013	0.0004	2.93	0.003
	α_0	0.0141	0.0059		
	d	0.3152	Fixed		
	Asymmetry	0.2875	0.1710	1.681	0.093
	Power	1.7219	0.1908		
	α_1	-0.2567	0.0222	-11.555	0
	α_3	-0.1853	0.0293	-6.331	0
	β_1	0.3737	0.0703	5.32	0
	β_2	-0.5077	0.0484	-10.499	0
	β_3	0.2874	0.0417	6.901	0
FIEGARCH(2,2)	const	0.0008	0.0004	1.851	0.064
	α_0	8.1815	0.0585		
	d	0.3152	Fixed		
	Asymmetry	-1.6233	0.3502	-4.635	0
	α_1	-0.2308	0.0209	-11.071	0
	α_2	-0.2276	0.0211	-10.784	0
	β_1	-0.4153	0.0689	-6.028	0
	β_2	0.5761	0.0686	8.396	0

Table A.6: Identified models for the daily Brent prices returns.

(a)

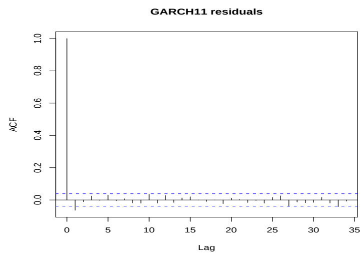

(b)

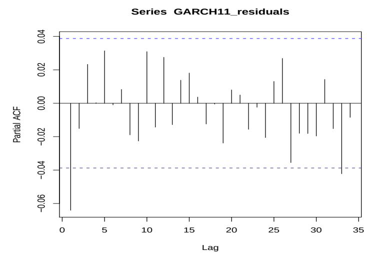

(c)

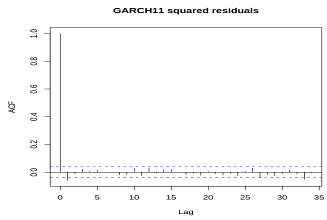

(d)

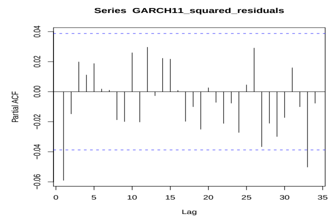

Figure A.2: For the GARCH(1,1) identified model: Empirical ACF (a) and empirical PACF (b) of the residuals; Empirical ACF (c) and empirical PACF (d) of the squared residuals.

(a)

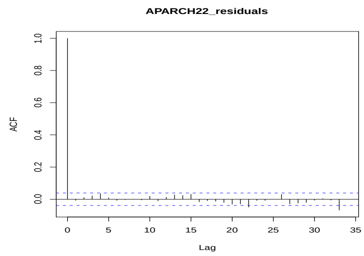

(b)

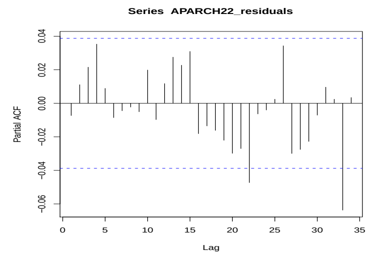

(c)

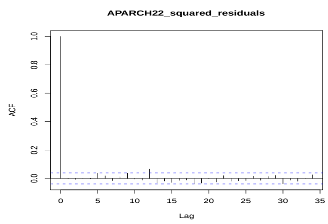

(d)

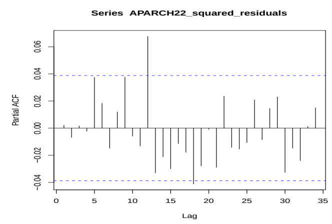

Figure A.3: For the APARCH(2,2) identified model: Empirical ACF (a) and empirical PACF (b) of the residuals; Empirical ACF (c) and empirical PACF (d) of the squared residuals.

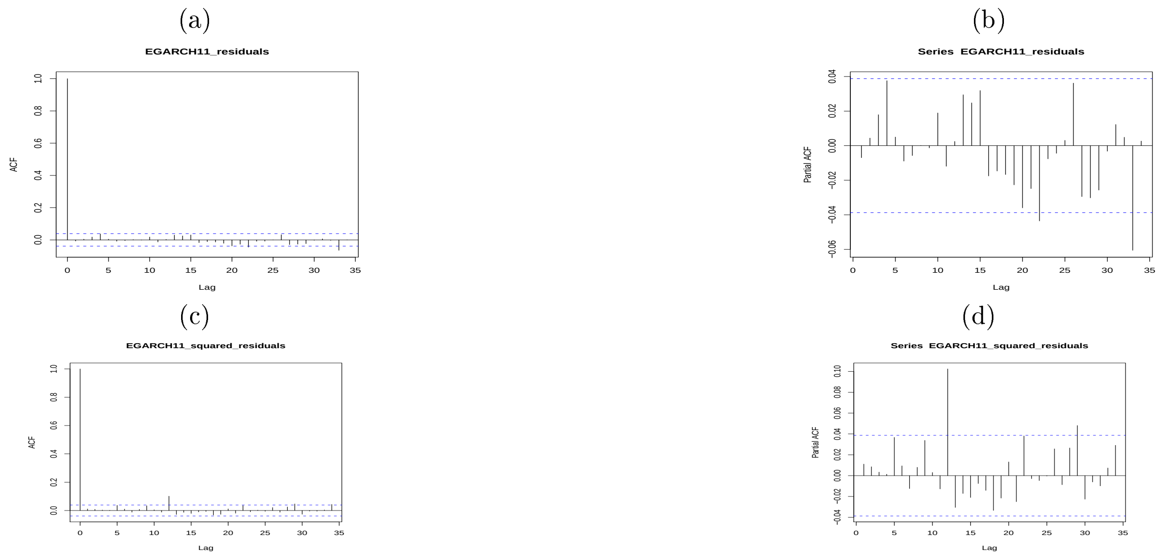

Figure A.4: For the EGARCH(1,1) identified model: Empirical ACF (a) and empirical PACF (b) of the residuals; Empirical ACF (c) and empirical PACF (d) of the squared residuals.

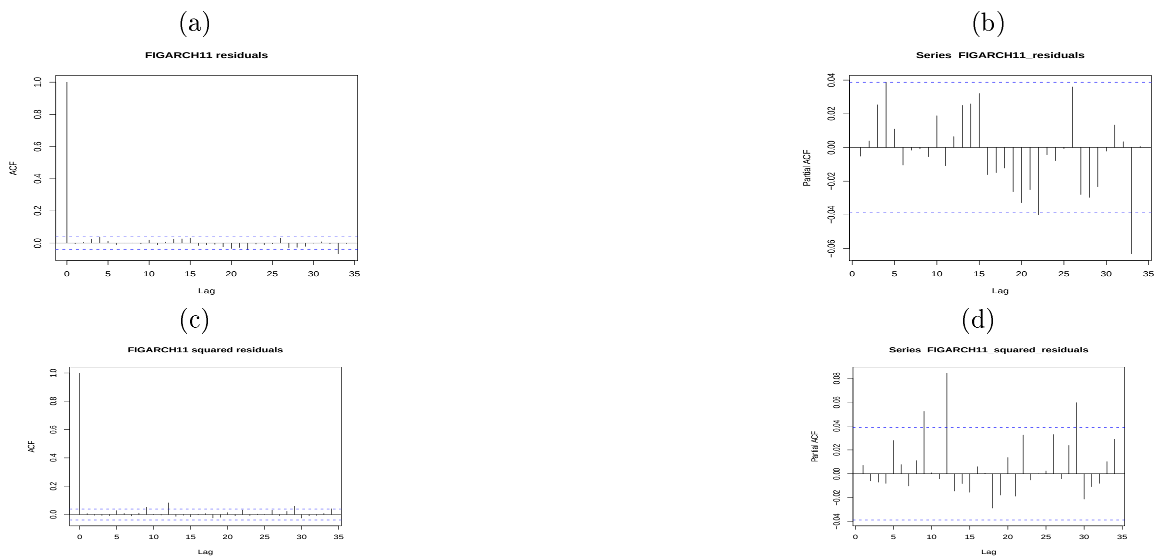

Figure A.5: For the FIGARCH(1,1) identified model: Empirical ACF (a) and empirical PACF (b) of the residuals; Empirical ACF (c) and empirical PACF (d) of the squared residuals.

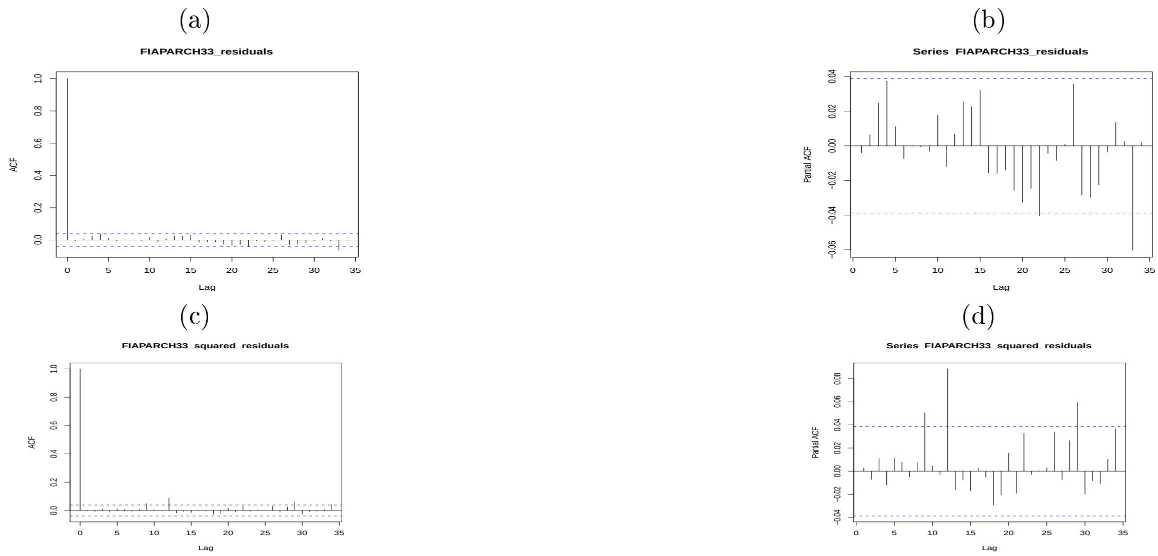

Figure A.6: For the FIAPARCH(3,3) identified model: Empirical ACF (a) and empirical PACF (b) of the residuals; Empirical ACF (c) and empirical PACF (d) of the squared residuals.

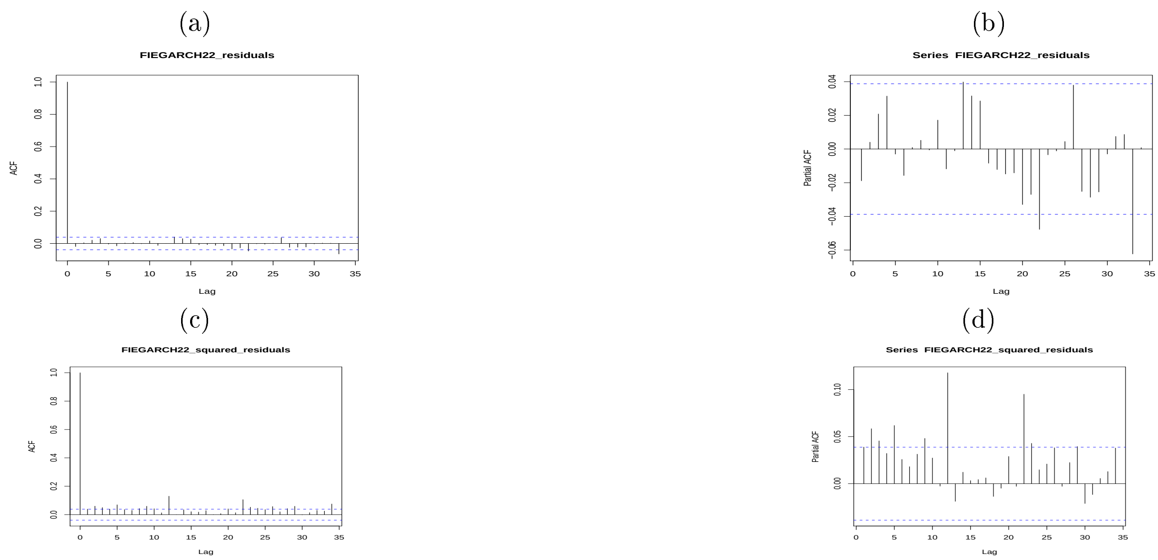

Figure A.7: For the FIEGARCH(2,2) identified model: Empirical ACF (a) and empirical PACF (b) of the residuals; Empirical ACF (c) and empirical PACF (d) of the squared residuals.

λ	Sample variance
0	29
1	80516

Table A.7: Sample variance in the Box-Cox's power transformation for the monthly data.

			p-value	
Test	H_0	Brent	ln(Brent)	returns
KPSS	Stat.	0.01	0.01	0.10
ADF	Non-Stat.	0.05	0.18	0.01
PP	Non-Stat.	0.17	0.32	0.01

Table A.8: Stationarity testing of the monthly Brent prices.

	Test	p-value
Residuals	Box-Pierce	0.952
	Ljung-Box	0.951
	ARCH	0.092
Squared residuals	Box-Pierce	0.176
	Ljung-Box	0.170

Table A.9: Diagnostic testing of the residuals of the AR(6) model applied to the monthly Brent prices returns.

		Estimate	Std. Err.	t Ratio	p-value
AR(6)	ϕ_6	-0.226	0.082	-2.767	0.007

Table A.10: Identified model for the monthly Brent prices returns using time series models.

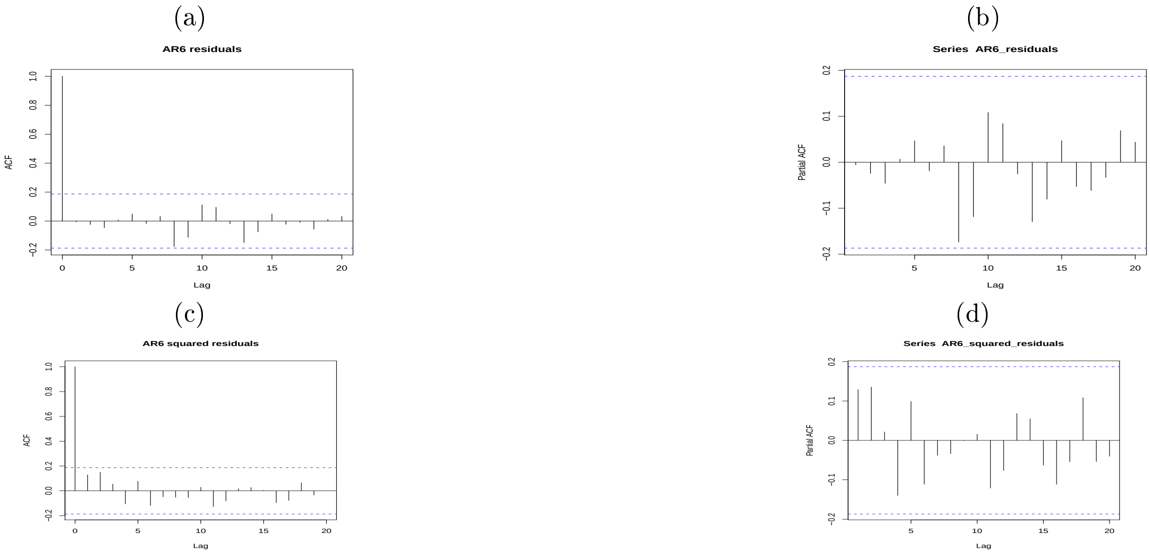

Figure A.8: For the AR(6) identified model: Empirical ACF (a) and empirical PACF (b) of the residuals; Empirical ACF (c) and empirical PACF (d) of the squared residuals.

A.3 Regime-switching Models

Delay parameter (δ)	Smoothness parameter (γ)	Threshold
1	40.00	80.29
2	40.00	36.01
3	40.00	34.87
4	40.00	33.36
5	40.00	35.28
6	40.00	36.27
7	26.00	38.19
8	17.00	38.19
9	40.00	39.46

Table A.11: LTAR models for the daily data, considering a two-regime model.

Test	p-value
SETAR(1) against SETAR(2)	0.4
SETAR(1) against SETAR(2)	0.3

Table A.12: SETAR tests for the daily data.

Delay parameter (δ)	Smoothness parameter (γ)	Threshold
1	13.39	72.76
2	40.00	72.17
3	63.13	77.21
4	40.00	77.16

Table A.13: LTAR models for the monthly data, considering a two-regime model.

Test	p-value
SETAR(1) against SETAR(2)	0
SETAR(1) against SETAR(2)	0
SETAR(2) against SETAR(3)	0.8

Table A.14: SETAR tests for the monthly data.

		p-values	
	Test	Regime 1	Regime 2
SETAR(2)	Box-Pierce	0.831	0.121
Model	Ljung-Box	0.828	0.095
MS-AR(2)	Box-Pierce	0.777	0.324
Model	Ljung-Box	0.774	0.277

Table A.15: Correlation testing for the regime-switching models applied to the monthly data.

		Regime 1			Regime 2		
		Estimate	Std. Err.	p-value	Estimate	Std. Err.	p-value
SETAR(2)	θ_0	0.554	1.227	0.653	21.581	4.135	0.000
Model	ϕ_1	0.886	0.128	0.000	1.721	0.088	0.000
	ϕ_2	0.121	0.136	0.375	-0.951	0.101	0.000
Threshold value: 70.19							
Lagged value: 3 (months)							
MS-AR(2)	θ_0	27.623	4.437	0.000	16.400	9.282	0.080
Model	ϕ_1	1.143	0.084	0.000	0.981	0.148	0.000
	ϕ_2	-0.320	0.095	0.001	0.091	0.154	0.557
Transition probability:			P(.|1)	P(.|2)			
		P(1.|)	0.463	0.192			
		P(2.|)	0.537	0.808			

Table A.16: Identified models for the monthly Brent prices using regime-switching models.

A.4 Model Selection

	DM	p-value
GARCH(1,1) vs APARCH(2,2)	290.09	0.00
GARCH(1,1) vs EGARCH(1,1)	220.97	0.00
GARCH(1,1) vs FIGARCH(1,1)	-103.91	0.00
GARCH(1,1) vs FIAPARCH(3,3)	-80.58	0.00
APARCH(2,2) vs EGARCH(1,1)	-69.12	0.00
APARCH(2,2) vs FIGARCH(1,1)	-394.00	0.00
APARCH(2,2) vs FIAPARCH(3,3)	-370.67	0.00
EGARCH(1,1) vs FIGARCH(1,1)	-324.88	0.00
EGARCH(1,1) vs FIAPARCH(1,1)	-301.55	0.00
FIGARCH(1,1) vs FIAPARCH(3,3)	23.33	0.00

Table A.17: Diebold-Mariano test applied to the daily forecasted samples.

	DM	p-value
AR(6) vs SETAR(2)	-10.93	0.00
AR(6) vs MS-AR(2)	-5.35	0.00
SETAR(2) vs MS-AR(2)	8.99	0.00

Table A.18: Diebold-Mariano test applied to the monthly forecasted samples.