Scientific
Research
Publishing

Financial Time Series Modelling of Trends and Patterns in the Energy Markets

Jane Aduda[1], Patrick Weke[2], Philip Ngare[2], Joseph Mwaniki[2]

[1]Department of Statistics and Actuarial Sciences, Jomo Kenyatta University of Agriculture and Technology, Nairobi, Kenya
[2]School of Mathematics, University of Nairobi, Nairobi, Kenya
Email: jaduda@jkuat.ac.ke, pweke@uonbi.ac.ke, pngare@uonbi.ac.ke, jimwaniki@uonbi.ac.ke

Abstract

Precise recognition of a time series path is important to policy makers, statisticians, economists, traders, hedgers and speculators alike. The correct time series path is also a key ingredient in pricing models. This study uses daily futures prices of crude oil and other distillate fuels. This paper considers the statistical properties of energy futures and spot prices and investigates the trends that underlie the price dynamics in order to gain further insights into possible nuances of price discovery and energy market dynamics. The family of ARMA-GARCH models was explored. The trends depict time varying variability and persistence of oil price shocks. The return series conform to a constant mean model with GARCH variance.

Keywords

1. Introduction

Energy use is behind virtually everything a person comes into contact with. The energy industry has rapidly expanded and become increasingly interdependent. In developed economies, the increase in energy consumption indicates a reliance on energy and its related products for continued and sustainable economic growth and development. Developing economies also rely on the development of energy resources to drive their growth. Energy was once viewed just as a utility, and an enabler with limited consumer interest, but now, it is key in the struggle for sustainable future economic growth [1] [2].

Energy prices, which are largely linked to oil prices, are a major concern for most economies. The recent financial crises and their ripple effects and after shocks have been largely unprecedented in terms of timing,

speed and magnitude of impact on the world economies. Forecasting of crude oil prices is important for better investment and risk management and policy development, and econometric models are the most commonly used. Various authors have employed time series [3] [4], financial [5] and structural models [6] in forecasting crude oil prices. Financial studies involve asset returns [7]-[9], instead of prices, since, firstly, returns give a complete and scale-free summary of the investment opportunities and secondly, return series have nice statistical properties known as stylised facts. These properties have important implications in model selection. Some of these interesting phenomena include volatility clusters which relate to the observation that the magnitudes of volatilities of financial returns tend to cluster together [9] [10], the leverage effect which relates to the asymmetry of news impact [11], leptokurticity, calendar effects and non-linear dependence.

In this paper, we regard an observed price series, $p_1, p_2, \cdots, p_T$, as a particular realisation of a stochastic process $\{p_t\}$. Standard time series analysis rests on important concepts such as stationarity, ergodicity, autocorrelation, white noise, innovation, and on a central family of models, the autoregressive moving average (ARMA) models [12] [13]. A stationary time series can be usefully described by it's mean, variance and autocorrelation or spectral density function [12] [14], and the joint probability distribution of any two observations, say, p_t, and p_{t+k} is the same for any two time periods t and $t+k$ separated by the same lag k [15]. For an ergodic process, we need not observe separate independent replications of the entire process in order to estimate its mean value or other moments [16].

In addition to the ARMA models, for the analysis of financial time series, we introduce the concept of volatility, which is pivotal in finance. Along side volatility we also consider the main stylised facts concerning financial time series [8] [13] [17]. In order to capture these stylised facts, we fit the autoregressive conditional heteroscedastic (ARCH) model first developed by Engle [18] and later generalized by Bollerslev [19] to generalised autoregressive conditional heteroscedastic (GARCH) model. A GARCH model is a parametric error distribution, which does not only capture volatility clustering but also accommodates some of the thick tails commonly found in financial time series.

This paper details the empirical analysis of financial time series data. We consider the official daily prices from the trading floor of the New York Mercantile Exchange (NYMEX) for a specific delivery month for Cushing Oklahoma West Texas Intermediate (OK WTI), Reformulated Blendstock for Oxygenate Blending (RBOB), and the number 1 heating oil futures contracts at close of business at 2:30 p.m. We begin the analysis from model identification, selection, estimation, diagnostic checking and finally forecasting. The analysis is done using MATLAB [20]. The rest of the paper is organized as follows: Section 2 discuses the methods used in empirical analysis, Section 3 gives a brief discussion of the data used, and the empirical analysis and the results obtained and finally Section 4 gives a summary of the work and the findings.

2. Methodology

2.1. Models

The full model under consideration falls under the general class of the combined $ARMA(p,q) - GARCH(P,Q)$ models which include an ARMA specification in the conditional mean and a GARCH specification in the conditional variance. The application of $ARMA(p,q) - GARCH(P,Q)$ model is a common approach in time series analysis that considers autocorrelation, volatility clustering, and heteroscedasticity. This class caters for models with constant variance under the $ARMA(p,q)$, and heteroscedastic models under the $GARCH(P,Q)$.

If $\{p_t\}$ denotes the price series at time t, let r_t be the log-return, and R_t, the relative or simple return so that

$$r_t = \log p_t - \log p_{t-1} = \ln\left(\frac{p_t}{p_{t-1}}\right) = \ln\left(1 + \frac{p_t - p_{t-1}}{p_{t-1}}\right) = \ln(1 + R_t), \tag{1}$$

an $ARMA(p,q) - GARCH(P,Q)$ model for the log-return series r_t, is specified as follows:

$$r_t = \sum_{i=1}^{p} \phi_i r_{t-i} + \sum_{i=1}^{q} \psi_i a_{t-i} + a_t \quad \text{and } a_t = \sigma_t \varepsilon_t \quad \text{where } \varepsilon_t \sim N(0,1)$$

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{P} \alpha_i a_{t-i}^2 + \sum_{j=1}^{Q} \beta_j \sigma_{t-j}^2. \tag{2}$$

When $P = Q = 0$, the model reduces to an $ARMA(p, q)$. Other sub-models can be deduced depending on the respective values of the parameters p, q, P, Q.

2.2. Model Identification, Parameter Estimation and Model Selection

We start with an exploratory analysis of the price data to establish if the set of properties, common across many stocks, in many markets over time periods, that have been observed by many studies [7] [8] [17] are present. We examine spot and futures prices for these stylised facts and other statistical properties such as stationarity. We analyse the descriptive statistics and test for normality using the Jarque-Bera (JB) test [21].

To test for stationarity, the Dickey-Fuller (DF) test [22] is employed. If the data is non-stationary, differencing is performed until stationarity is achieved. Once stationarity is established, we study the autocorrelation function (ACF) and partial autocorrelation function (PACF) in order to determine the form of the $ARMA(p, q)$ model for each series. A set of candidate models, around this initial model are then identified for further comparison.

Once the models for the conditional mean and variance have been identified from the differenced series, we employ the maximum likelihood estimation (MLE) method to fit parameters for the specified model of the differenced series.

The final model for each series is then selected based on Akaike's Information Criterion (AIC) which measures of goodness of fit or uncertainty for the range of values of the data. It measures the difference between a given model and the "true" underlying model. AIC is a function of the squared sum of errors (SSE), number of observations *n*, and the number of independent variables $k \leq p + 1$ where k includes the intercept, and is calculated as $AIC = n \cdot \ln\left(\dfrac{SSE}{n}\right) + 2k$. The best model is the one with the smallest AIC.

2.3. Forecasting

Model selection does not depend only on the goodness-of-fit of a model to the data, but also on the objective of the analysis. A model that is best in the in-sample fitting will not necessarily provide more accurate forecasts [8]. For this reason, it is common to use the performance of out-of-sample forecasts to aid in the selection of an adequate statistical model.

The forecasting ability of the model can only be determined by considering how well a model performs on data not used in estimating the model. It is common practice to partition the data into two sets, use the larger portion for estimating the model and the smaller one for testing the model. The test data can be used to measure model accuracy on new data. The size of the test data set should typically be about 20% of the total sample, although this depends on the sample size and the forecast horizon. In an ideal sense, the size of the test sample should be at least as large as the maximum forecast horizon required [23].

The sample standard deviation, σ of return is used as a simple forecast of volatility of returns, r_t, over the future period $[t, t + k]$. The k-days period historical variance is calculated as $\hat{\sigma}_t^2(k) = \dfrac{1}{k-1}\sum_{i=0}^{k-1}(r_{t-i} - \bar{r})$, where the sample mean return $\bar{r} = \dfrac{1}{k}\sum_{i=0}^{k-1}r_{t-i}$, is the estimate of the mean return μ. In most cases, the mean is usually set to zero in order to get a better forecast. Multiplying the variance by *N*, the number of trading days in a year and taking the square root gives the annualised volatility $\hat{\sigma} = \sqrt{N \cdot \widehat{\sigma^2}}$. The number of trading days in a year is taken to be $N = 250$. $\hat{\sigma}$ is the best estimator for the volatility from the available price data and the volatility of any period of length, *k*, can be estimated from it.

The k-day volatility forecast can also be found by using the GARCH model. Under the condition that returns are uncorrelated across days, the k-day variance as of $t - 1$ is given by

$$E\left[r_{t,t+k}^2 \mid \mathcal{F}_{t-1}\right] = E\left[r_t^2 \mid \mathcal{F}_{t-1}\right] + E\left[r_{t+1}^2 \mid \mathcal{F}_{t-1}\right] + \cdots + E\left[r_T^2 \mid \mathcal{F}_{t-1}\right] \tag{3}$$

The volatility forecast over the future period from $t + 1$ to $t + k$, denoted by σ_{Ft}, is an average of the

expected volatility on each day from t to $t+k$ *i.e.*

$$\sigma_{Ft} = \sqrt{\frac{1}{k} \sum_{i=1}^{k} E\left(\sigma_{t+k}^2 \mid \mathcal{F}_t\right)} \qquad (4)$$

where

$$E\left(\sigma_{t+k}^2 \mid \mathcal{F}_t\right) = \alpha_0 \frac{1-\left(\alpha_1+\beta_1\right)^k}{1-\left(\alpha_1+\beta_1\right)} + \left(\alpha_1+\beta_1\right)^k \sigma_t^2. \qquad (5)$$

Note that if $\left(\alpha_1+\beta_1\right) < 1$,

$$\lim_{k \to 0} E\left(\sigma_{t+k}^2 \mid \mathcal{F}_t\right) = \frac{\alpha_0}{1-\alpha_1-\beta_1}, \qquad (6)$$

which is the unconditional variance.

In-sample and out-of-sample forecasting ability of models can be measured using the mean absolute error (MAE) and root mean squared forecasting errors (RMSFE) which measures the out-of-sample losses. The RMSFE assigns greater weight to large forecast errors. This fact is handled using the MAE which on the contrary assigns equal weights to both over and under predictions of volatility.

The realised volatility at each forecast date is calculated from the expression $\sigma_{R,t} = \sqrt{\frac{1}{N} \sum_{i=1}^{N} r_{t+i}^2}$, so that, if the forecast volatility is σ_{Ft}, the RMSFE for a model is given by $RMSFE = \sqrt{\frac{1}{n} \sum_{t \in s} \left(\sigma_{Ft} - \sigma_{R,t}\right)^2}$ and the MAE is given by $MAE = \frac{1}{n} \sum_{t=1}^{n} \left|\sigma_{Ft} - \sigma_{R,t}\right|$ where n and s denote the number of forecasts and the set of times at which ex-ante forecasts are produced respectively in the above expression.

Forecasting using GARCH models is obtained recursively [24]. If we left t be the forecast origin, then, the k-step ahead forecast for σ_{t+k}^2 is

$$\sigma_t^2(k) = \hat{\alpha}_0 + \left(\hat{\alpha}_1 + \hat{\beta}_1\right)\sigma_t^2(k-1) \quad \text{for } k > 1. \qquad (7)$$

An alternative way of writing Equation (7) is as shown in Equation (8)

$$\sigma_t^2(k) = \hat{\sigma}^2 + \left(\hat{\alpha}_1 + \hat{\beta}_1\right)^{k-1}\left(\sigma_t^2 - \hat{\sigma}^2\right) \quad \text{for } k > 1. \qquad (8)$$

where

$$\hat{\sigma}^2 = \frac{\hat{\alpha}_0}{1 - \hat{\alpha}_1 - \hat{\beta}_1}.$$

From Equation (8), we can see that $\sigma_t^2 \to \hat{\sigma}^2$ as $k \to \infty$

3. Data and Results

This paper investigates the trends and patterns of return series for spot and futures prices of Cushing OK WTI, RBOB and the number 1 heating oil traded in the NYMEX, for the period running from 2nd January 2006 to 22nd May 2015. The data was obtained from the US Energy Information Administration (EIA) website http://www.eia.gov/petroleum/data.cfm, accessed on June 25, 2015. The US EIA is the statistical and analytical agency within the US Department of Energy. It is the principal agency of the US Federal Statistical System responsible for collecting, analysing, and disseminating energy information to promote sound policy-making, efficient markets, and public understanding of energy and its interaction with the economy and the environment. EIA is the premier source of energy information in the US and, by law, its data, analyses, and forecasts are independent of approval by any other officer or employee of the US government. [25].

Figure 1 shows various series plots for daily spot and futures prices for Cushing OK WTI, RBOB gasoline
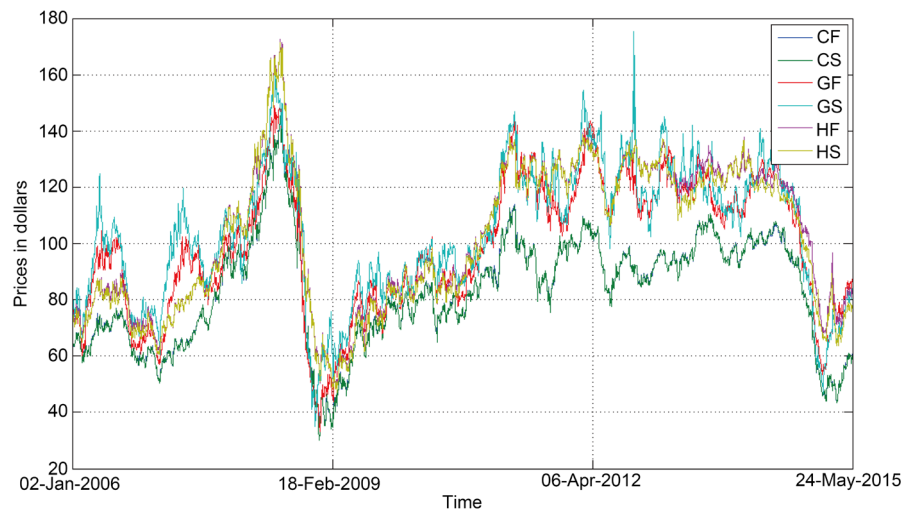
Figure 1. Time series plots for spot, futures prices and crack spreads in energy markets. CF = crude futures, CS = crude spot, GF = gasoline futures, GS = gasoline spot, HF = heating futures and HS = heating spot.

and No. 1 heating oil traded at the NYMEX. The data appears non-stationary, with occasional jumps and spikes indicating heteroscedasticity. Testing the price series for stationarity using DF test shows that they all contain unit roots.

Table 1 gives a summary of the descriptive statistics and results for the JB test for the differenced price series. From the measure of skewness, all the series except the crude spot are skewed to the left. These series exhibit positive excess kurtosis, and these are some of the stylised facts observed in financial time series data. Based on the *p*-values of the JB test, we reject the null hypothesis of normality for all the differenced series.

Figure 2 shows plots of the differenced crude futures and spot prices. The series appear mean stationary but are not variance stationary. Volatility clustering, jumps and spikes are also observed.

We employ the DF test for stationarity on the differenced series. This is a one-sided left tail test and **Table 2** gives the summary report of the DF test on the six series. Stationarity is confirmed for all these series, and hence they can be subjected to further time series analysis

Figure 3 and **Figure 4**, show the ACF and PACF plots of crude futures and spot price series. From these plots, it is not easy to tell the form of $ARMA(p,q)$, significant lags are noticed up to lag 39 of the 40 plotted. The same behaviour is exhibited by the other four series. The ACF and PACF fail to identify the form of the model. We proceed to fit several $ARMA(p,q)$ models for $0 \leq p \leq 40$ and $0 \leq q \leq 40$, and the one with the smallest AIC is picked. **Figure 5** shows a plot of AICs for various models fit for ΔCF, from which we see that the best model is an $ARMA(6,11)$ with an AIC of -4684.1. This model has the parameters estimated via the maximum likelihood, summarized in **Table 3**. For the crude spot data, the best model is an $ARMA(7,11)$. A similar analysis was also carried out on the differenced gasoline futures and prices and the results were similar with high orders of autoregressive and moving averages.

From **Table 3**, we can use the values of the *t*-statistic to check the significance of the parameters. The greater the absolute value of *t*, the greater the evidence against the null hypothesis that there is no significance difference between the parameter estimate and 0. We can use the general rule of thumb that we reject the null hypothesis whenever $|t| \geq 2$.

After fitting the $ARMA(6,11)$ model, we standardized the inferred residuals and checked for normality. **Figure 6** shows two plots. The first one is the residual series which shows a lot of clustering, suggesting non-normality and non-constant variance. The second plot of **Figure 6** is a histogram of the residuals also showing evidence of non-normality. **Figure 7** shows a quantile-quantile plot and a box plot of the residuals and from these, we can also deduce non-normality of the residuals largely because of the extreme values.

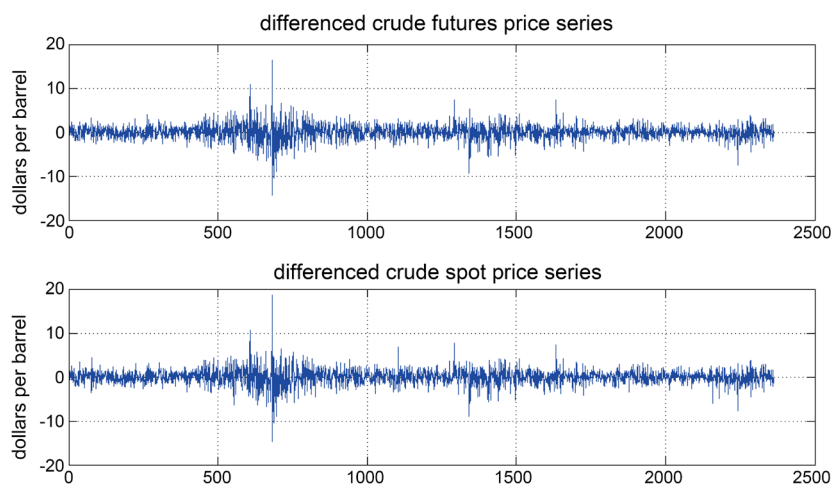The Kolmogorov-Smirnov (KS) test was employed to check normality of the residuals and the *p*-value was

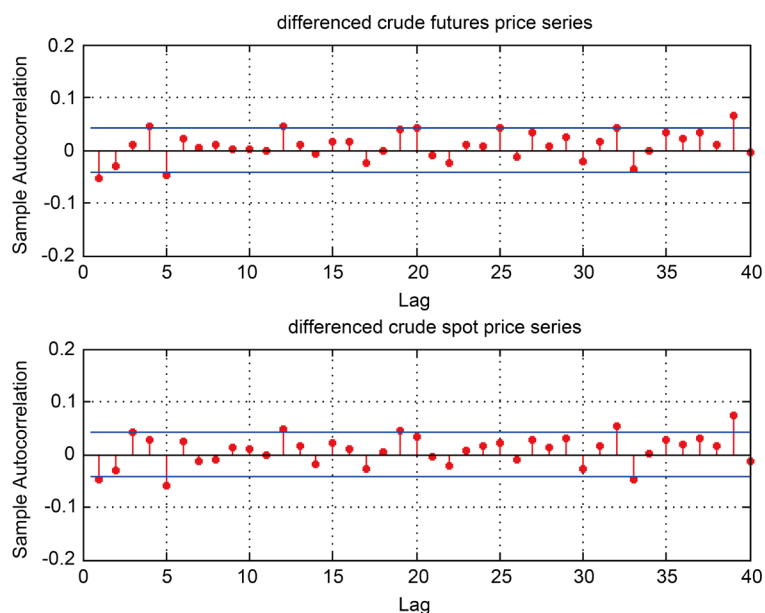**Figure 2.** Differenced series for crude futures and spot prices.

Figure 3. ACF plots for the differenced crude futures and spot price series.

Table 1. Descriptive statistics for energy spot and futures prices.

	ΔCF	ΔCS	ΔGF	ΔGS	ΔHF	ΔHS
Mean	−0.001	−0.002	0.005	0.003	0.003	0.001
Median	0.035	0.040	0.042	0.042	0.000	0.000
Maximum	16.370	18.560	13.020	21.336	12.306	11.508
Minimum	−14.310	−14.760	−17.724	−21.756	−17.304	−10.668
Skewness	−0.051	0.042	−0.456	−0.173	−0.467	−0.090
Kurtosis	10.252	12.046	8.008	8.312	8.699	5.837
JB p-value	0.001	0.001	0.001	0.001	0.001	0.001

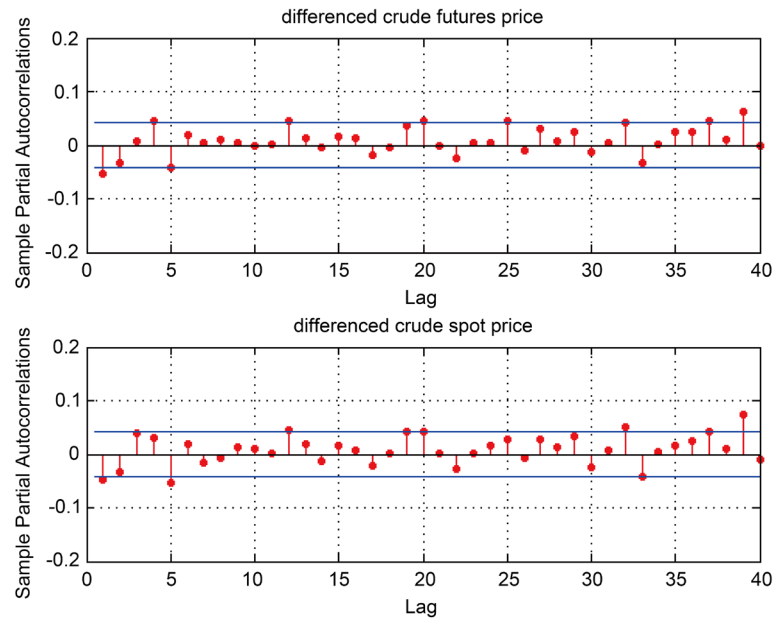Δ denotes differenced series.

Figure 4. PACF plots for the differenced crude futures and spot price series.
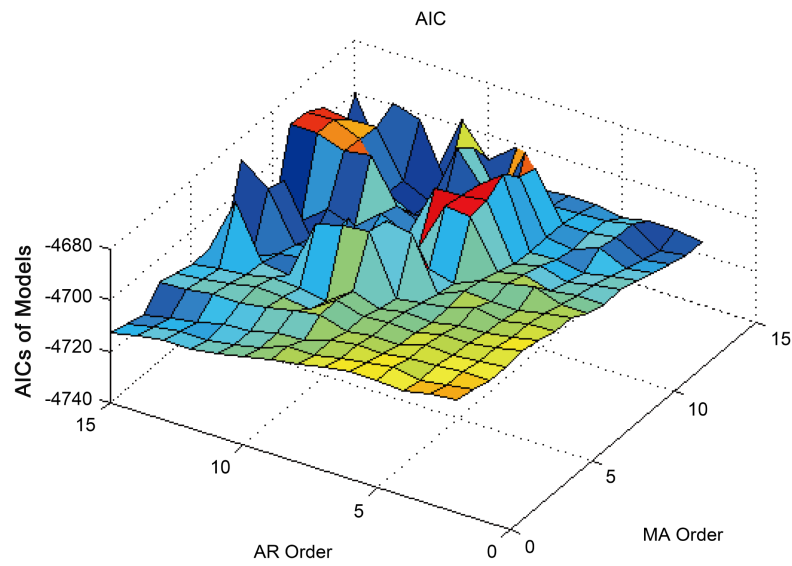

Figure 5. Plot for the AICs for the various ARMA models fitted for the differenced crude futures price data.

Table 2. Results from the Dickey-fuller test on the six series.

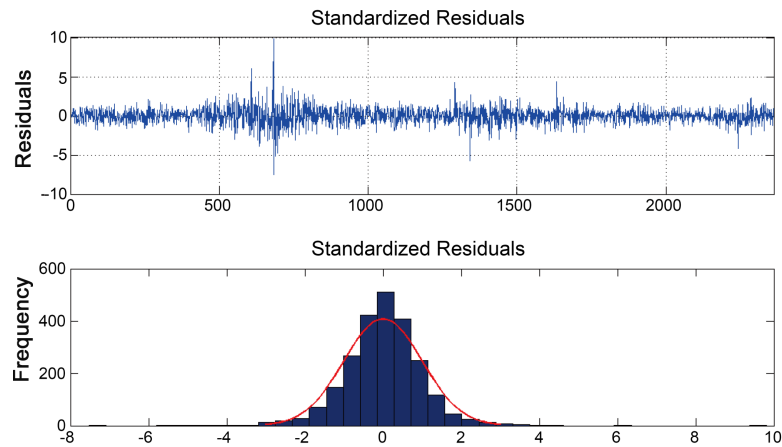| Series	*p*-value	Test statistic	Critical value
ΔCF	0.001	−51.2762	−1.9416
ΔCS	0.001	−50.9918	−1.9416
ΔGF	0.001	−49.5626	−1.9416
ΔGS	0.001	−50.9792	−1.9416
ΔHF	0.001	−50.0207	−1.9416
ΔHS	0.001	−50.6481	−1.9416

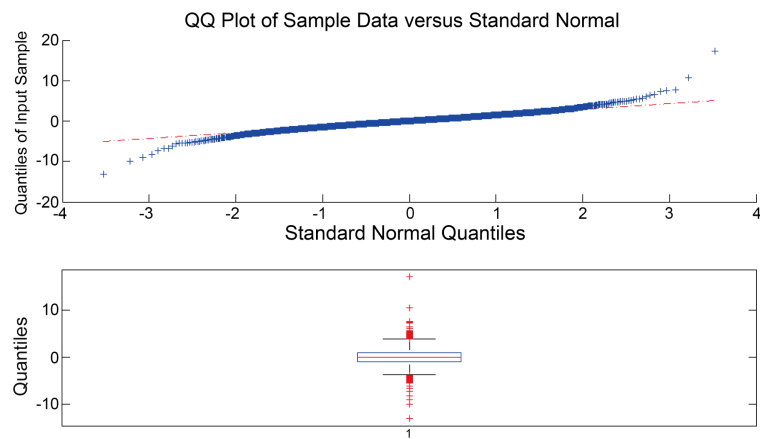**Figure 6.** Normality plots for residual checks.

Figure 7. QQ-plot and Box plot for residuals from the ARMA(6, 11) model.

Table 3. Results from the $ARMA(6,11)$ Model fitted to the differenced crude futures price series.

Parameter	Estimate	Std Error	*t*-Statistic
Constant	−0.0003	0.0055	−0.0494
AR(1)	0.1064	0.2303	0.4622
AR(2)	0.5423	0.2234	2.4273
AR(3)	0.0277	0.0757	0.366
AR(4)	−0.7031	0.0729	−9.652
AR(5)	0.0806	0.2265	0.3558
AR(6)	0.8426	0.2163	3.896
MA(1)	−0.1625	0.2304	−0.7054
MA(2)	−0.5671	0.2374	−2.3886
MA(3)	0.0127	0.0852	0.1488
MA(4)	0.7652	0.0712	10.7484
MA(5)	−0.1797	0.2365	−0.7597
MA(6)	−0.8514	0.2529	−3.3666
MA(7)	0.0874	0.0380	2.2975
MA(8)	0.0383	0.0189	2.0253
MA(9)	−0.0402	0.0185	−2.1675
MA(10)	<0.001	0.0257	−0.001
MA(11)	0.0432	0.0242	1.7835
Variance	3.069	0.0503	61.0522

found to be <0.001 implying the residuals are not normally distributed. In this case the $ARMA(6,11)$ does not conform to the assumption normality in the innovations.

The Ljung-Box (LB) test on the residual series resulted in a *p*-value of 0.9781 indicating the residuals are not autocorrelated. Based on these results, we test for ARCH effects in the squared residuals using the LB test and obtained a *p*- value <0.001, indicating significant ARCH effects in the residuals.

Figure 8 also shows evidence of serial autocorrelation on the squared residuals. The PACF plot shows autoregression of order one so $ARCH(1)$ should capture the ARCH effects in the residuals.

We then specify a combined ARMA-GARCH model to capture the ARCH effects. Fitting a ARMA(6,11)-GARCH(1,1) had an AIC of 8810.9. However, a search for the best combined model revealed that an $ARMA(0,0) - GARCH(1,1)$ model gave the best fit with an AIC of −4.4037. From this analysis, we see that a more appropriate model to use would be one with a constant conditional mean and a conditional variance.

The analysis shows the differenced series exhibit "stylised facts" typically seen in high frequency financial data. volatility clustering manifests itself as autocorrelation in squared and absolute returns, or in the residuals of the estimated conditional mean equation. Examining the ACF and PACF plots for the squared differenced series shown in **Figure 9** and **Figure 10**, reveal serial correlation of the squared differenced series. Gasoline and heating oil series exhibit the same behaviour.

The significance of these autocorrelations at various lags was tested using the LB test and the Lagrange multiplier (LM) test at lags 1, 5 and 10, and the results are summarized in **Table 4**. As is evident, there is enough evidence to reject the null hypothesis of no ARCH effects as per the LB test since the test indicates that there are significant ARCH effects in the all squared series. However, based on Engle's ARCH test, which is the LM test, we do not have sufficient evidence to reject the hypothesis of no ARCH effects for the gasoline futures and the heating oil futures series. For all the other series we reject that hypothesis of no ARCH effects.

Serial correlation in squared returns, or conditional heteroscedasticity can be modelled using GARCH models. GARCH models allow for the volatility to evolve with time. A GARCH model can be expressed as an ARMA model of squared residuals and hence many of it's properties follow easily from those of the corresponding ARMA process. If the fourth order moment of a GARCH (1,1) exists, the kurtosis implied by a GARCH (1,1) process with normal errors is greater than that of the normal distribution which is 3 [19]. For this reason, a GARCH model with normal errors can replicate some of the fat tailed behaviour observed in financial time series, though, most often, a GARCH model with a non-normal error distribution is required in order to fully capture the observed fat tail behaviour in financial time series.
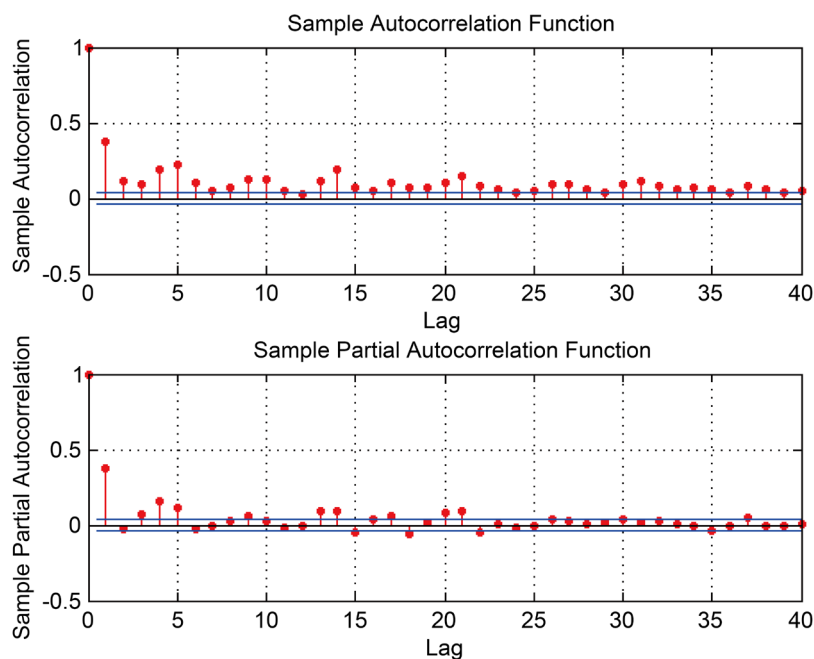
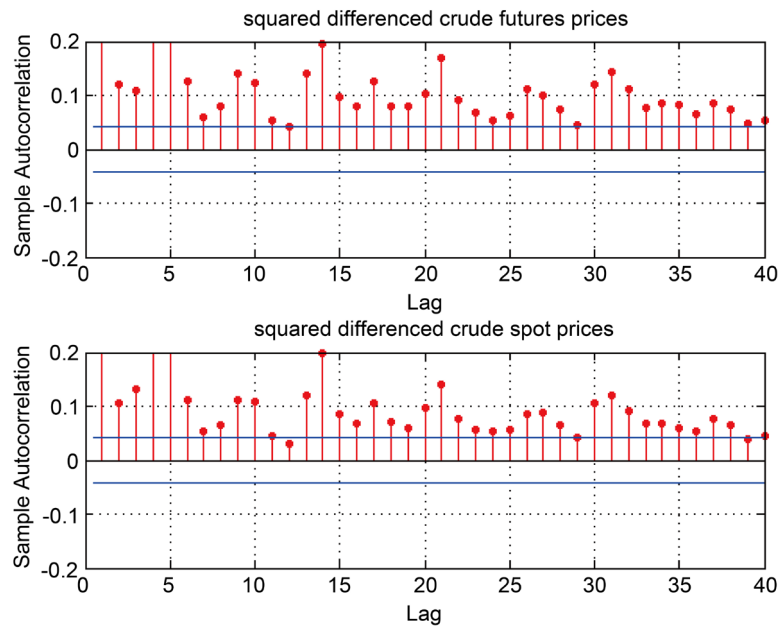**Figure 8.** ACF and PACF for squared residuals from the ARMA(6,11) model.

Figure 9. ACF plots for the squared differenced crude futures and spot price series.
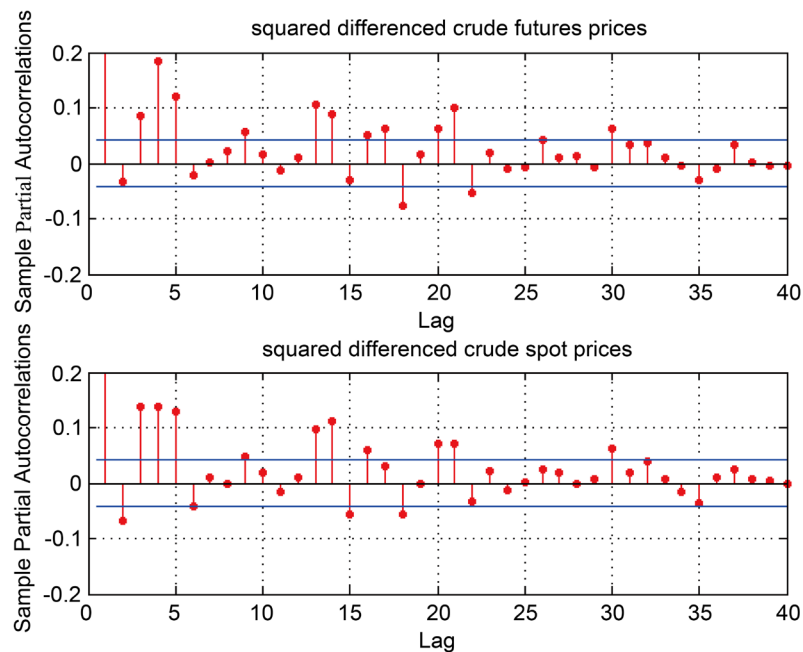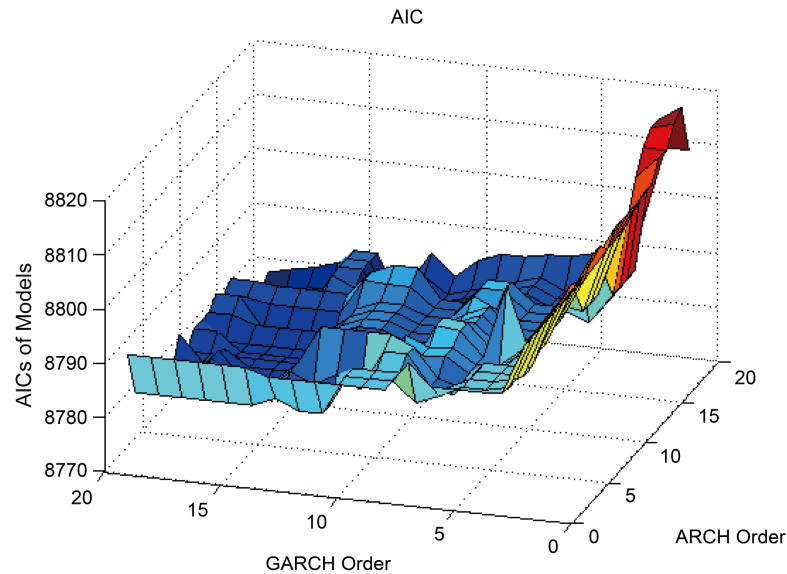

Figure 10. PACF plots for the squared differenced crude futures and spot price series.

We fit several $GARCH(p,q)$ model to the differenced crude futures series for $p = 20, q = 20$. We estimate the parameters α_0, α_i and β_j of the $GARCH(p,q)$ model and the in-sample estimates of the volatility. The Matlab command *garch*, models the return series as $GARCH(p,q)$ and estimates the parameters α_0, α_i and β_j via maximum likelihood.

For the differenced crude futures series, the model with the smallest AIC was GARCH (19, 16) with an AIC of 8776.4 as shown in **Figure 11** which shows the AICs of all the models fitted. The best model according to these fits have the parameters summarized in **Table 5**.

Figure 11. A plot for the AICs for various GARCH models fitted.

Table 4. Results from the Ljung-Box test and LM test for ARCH effects on the squared differenced series.

Lag	$Q(p)$			LM -test		
	1	5	10	1	5	10
$(\Delta CF)^2$	348.786	677.980	821.867	431.691	564.068	608.241
	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)
$(\Delta CS)^2$	383.810	696.811	800.000	287.646	346.920	358.553
	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)
$(\Delta GF)^2$	39.142	101.436	157.990	3.614	4.324	5.329
	(<0.001)	(<0.001)	(<0.001)	(0.0573)	(0.5038)	(0.8681)
$(\Delta GS)^2$	317.298	858.416	962.298	311.786	792.509	866.792
	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)
$(\Delta HF)^2$	45.923	180.863	271.397	1.344	5.341	5.720
	(<0.001)	(<0.001)	(<0.001)	(0.2463)	(0.3758)	(0.8382)
$(\Delta HS)^2$	79.112	251.739	379.028	23.281	38.898	46.177
	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)

p-values are indicated in the parenthesis.

From these results, it can be seen that most of the in-between lags are not significant, the most important lags at 16 and 19. An analysis on the residuals of this model show that this model gives a fairly good fit. As per the AIC values, most these models do not seem to be significantly different and so the lags may be reduced to gain parsimony. An examination of the plots of the standardised residuals after fitting the GARCH(1,1) model for the return series indicates that the residuals are, on the overall, stable with some clustering. The standardised residuals exhibit no residual autocorrelation. It therefore makes more sense just to fit a $GARCH(1,1)$ for this data.

The other series also behave in much the same way as the crude futures. Fitting a $GARCH(1,1)$ model on all the six series gives the results in **Table 6**. From the measure of persistence, we see the highest persistence of 0.9938 in the ΔCF series, although, all other return series are quite persistent. This indicates that the volatility

Table 5. Results from the $GARCH(19,16)$ model fitted to the differenced crude futures price series.

Parameter	Estimate	Std Error	*t*-Statistic
Constant	0.175	0.103	1.701
GARCH (16)	0.090	0.169	0.530
GARCH (19)	0.348	0.200	1.743
ARCH (1)	0.121	0.022	5.475
ARCH (2)	0.071	0.053	1.340
ARCH (5)	0.038	0.052	0.729
ARCH (6)	0.037	0.046	0.811
ARCH (8)	0.026	0.033	0.780
ARCH (9)	0.007	0.037	0.188
ARCH (10)	0.050	0.037	1.346
ARCH (13)	0.022	0.037	0.585
ARCH (14)	0.036	0.038	0.955
ARCH (15)	0.013	0.036	0.348
ARCH (16)	0.083	0.036	2.288

Table 6. Results from the $GARCH(1,1)$ model fitted to all the differenced price series.

| Parameter	ΔCF	ΔCS	ΔGF	ΔGS	ΔHF	ΔHS
α_0	0.0191	0.0265	0.3354	0.5388	0.0743	0.0428
α_1	0.9464	0.9388	0.8236	0.8318	0.915	0.9429
β_1	0.0475	0.058	0.1072	0.1036	0.0642	0.0451
$(\alpha_1 + \beta_1)$	0.9938	0.9916	0.9308	0.9354	0.9792	0.9878

$(\alpha_1 + \beta_1)$ measures the persistence.

processes return to their means after some time. The implication of this high persistence is that these shocks have a significant effect on the prices in these markets.

Forecasting

The total length of the data under consideration is 2363 data points. This data set was divided into two sets, the first sub-sample, used for training the model contained 2063 points and the second one, used for testing, contained 300 points.

In order to regain a forecast of the price series, the differencing on the forecast series obtained from the forecasts of the inferred conditional variances is undone. However because the inferred variances are conditioned on a standard normal error, there are several possibilities and, having many different forecasts and averaging would give a stable forecast for the crude futures prices, which can then be used to predict future crude oil futures prices. In **Figure 12**, we see one of the possibilities of the forecast prices. The RMSFE for this model is 0.2072, an indication that the forecasts are generally good.
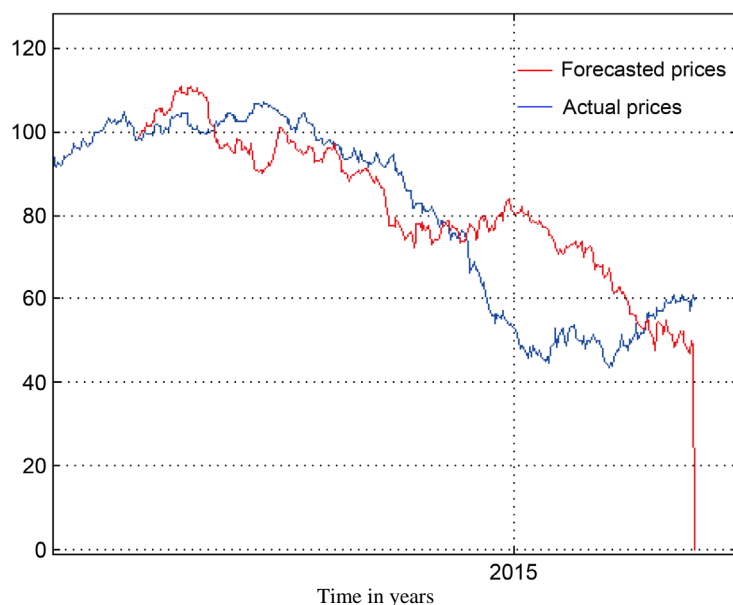
Figure 12. Crude futures price series showing foretasted prices.

4. Discussion

This paper employs statistical and econometric techniques to investigate and model financial time series trends in energy markets. To do this, daily closing prices for a period of about 10 years for Cushing OK WTI, RBOB and number 1 heating oil spot and futures contracts traded in the NYMEX are considered. The paper also investigates the existence of stylised facts in these series, in order to fit an appropriate model that adequately describes the market dynamics.

Price data are tested for stationarity using the DF test and results show the existence of unit roots. Normality is tested for using the JB test and non-normality is established. Return series are then generated from the price series through differencing, and then tested for stationarity using the DF test. The results reveal that return series are indeed mean stationary, but are definitely not variance stationary. Several ARMA models are fit to the return series and the standardised residuals analysed. For the crude futures return series, the best model turns out to be an $ARMA(6,11)$ an for the crude spot it is an $ARMA(7,11)$ among others, based on their AICs. These resultant models however contravene the Gaussian innovation assumption.

We then propose a combined $ARMA(p,q)-GARCH(P,Q)$ model to capture the ARCH effects in the variance, then find that the best model under these circumstances is the $ARMA(0,0)-GARCH(1,1)$, implying a constant mean conditional variance equation. We finally find that for example, for the crude futures return series, the $GARCH(16,19)$ is the best model based on the AIC, although a $GARCH(1,1)$ does quite well for this and all the other six series with the residual analysis conforming to the assumption of Gaussian Innovations. GARCH models can therefore adequately model the trends and patterns in the energy markets. The trends also depict time varying variability and high persistence of oil price shocks. These shocks therefore have a significant impact on the prices of these energy prices.

References

[1] Dorsman, A., Simpson, J.L. and Westerman, W. (2013) Energy Economics and Financial Markets. Springer, Berlin. http://dx.doi.org/10.1007/978-3-642-30601-3

[2] Ramos, S.B. and Veiga, H. (2014) The Interrelationship between Financial and Energy Markets. Springer, Berlin. http://dx.doi.org/10.1007/978-3-642-55382-0

[3] Liu, H.P., Erdem, E. and Shi, J. (2011) Comprehensive Evaluation of ARMA-GARCH (-M) Approaches for Modeling the Mean and Volatility of Wind Speed. *Applied Energy*, **88**, 724-732. http://dx.doi.org/10.1016/j.apenergy.2010.09.028

[4] Beck, A., Shin, Y., Kim, A., Rachev, S., Feindt, M. and Fabozzi, F. (2013) Empirical Analysis of ARMA-GARCH Models in Market Risk Estimation on High-Frequency US Data. *Studies in Nonlinear Dynamics and Econometrics*, **17**, 167-177. http://dx.doi.org/10.1515/snde-2012-0033

[5] Chinn, M.D., LeBlanc, M. and Coibion, O. (2005) The Predictive Content of Energy Futures: An Update on Petroleum, Natural Gas, Heating Oil and Gasoline. National Bureau of Economic Research, Cambridge.

[6] Alvarez-Ramirez, J., Rodriguez, E., Martina, E. and Ibarra-Valdez, C. (2012) Cyclical Behavior of Crude Oil Markets and Economic Recessions in the Period 1986-2010. *Technological Forecasting and Social Change*, **79**, 47-58. http://dx.doi.org/10.1016/j.techfore.2011.04.010

[7] Cont, R. (2001) Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues. *Quantitative Finance*, **1**, 223-236. http://dx.doi.org/10.1080/713665670

[8] Tsay, R.S. (2010) Analysis of Financial Time Series. 3rd Edition, John Wiley & Sons, Inc., Hoboken.

[9] Mandelbrot, B. (1963) The Variation of Certain Speculative Prices. *The Journal of Business*, **36**, 394-419. http://dx.doi.org/10.1086/294632

[10] Engle, R.F. and Patton, A.J. (2001) What Good Is a Volatility Model. *Quantitative Finance*, **1**, 237-245. http://dx.doi.org/10.1088/1469-7688/1/2/305

[11] Black, F. (1976) The Pricing of Commodity Contracts. *Journal of Financial Economics*, **3**, 167-179. http://dx.doi.org/10.1016/0304-405x(76)90024-6

[12] Box, G.E.P. and Jenkins, G.M. (1976) Time Series Analysis: Forecasting and Control, Revised Edition, Holden-Day.

[13] Francq, C. and Zakoian, J.-M. (2011) GARCH Models: Structure, Statistical Inference and Financial Applications. John Wiley & Sons.

[14] Hamilton, J.D. (1994) Time Series Analysis. Princeton University Press, Princeton.

[15] Montgomery, D.C., Jennings, C.L. and Kulahci, M. (2011) Introduction to time Series Analysis and Forecasting. John Wiley & Sons, Hoboken.

[16] Chan, N.H. (2011) Time Series: Applications to Finance with R and S-Plus. John Wiley & Sons, Hoboken.

[17] Zivot, E. and Wang, J.H. (2007) Modeling Financial Time Series with S-PLUS. Springer Science & Business Media, Berlin.

[18] Engle, R.F. (1982) Autoregressive Conditional Heteroscedasticity with Estimates of the variance of United Kingdom Inflation. *Econometrica*, **50**, 987-1007. http://dx.doi.org/10.2307/1912773

[19] Bollerslev, T. (1986) Generalized Autoregressive Conditional Heteroscedasticity. *Journal of Economics*, **31**, 307-327. http://dx.doi.org/10.1016/0304-4076(86)90063-1

[20] The MathWorks (2013) MATLAB Version 7.14.0.739 R2013a. The MathWorks Inc., Torrance.

[21] Jarque, C.M. and Bera, A.K. (1987) A Test for Normality of Observations and Regression Residuals. *International Statistical Review/Revue Internationale de Statistique*, **55**, 163-172. http://dx.doi.org/10.2307/1403192

[22] Dickey, D.A. and Fuller, W.A. (1979) Distribution of the Estimators for Autoregressive Time Series with a Unit Root. *Journal of the American Statistical Association*, **74**, 427-431.

[23] Hyndman, R.J. and Koehler, A.B. (2006) Another Look at Measures of Forecast Accuracy. *International Journal of Forecasting*, **22**, 679-688. http://dx.doi.org/10.1016/j.ijforecast.2006.03.001

[24] Jondeau, E., Poon, S.-H. and Rockinger, M. (2006) Financial Modeling under Non-Gaussian Distributions. Springer, Berlin.

[25] EIA (2015) US Energy Information Administration. http://www.eia.gov/about/