A Time Series Analysis of Monthly Oil Prices

Josh Ingram

November 2014

TABLE OF CONTENTS

INTRODUCTION

THE OBJECTIVE OF THE PROJECT

I will analyze the monthly spot prices of West Texas Intermediate (WTI) and my objectives are:

- to model prices from January 1979 to December 2006 using the methods learned in the course;
- select an appropriate model and perform diagnostics; and
- predict potential future oil prices for the first quarter of 2007.

DATA DESCRIPTION

Crude oil prices behave much as any other commodity with wide price swings in times of shortage or oversupply. The crude oil price cycle may extend over several years responding to changes in demand as well as the Organization of the Petroleum Exporting Countries (OPEC) and non-OPEC supply. OPEC's quota decisions are signals to the market about the preferred range of prices; a quota is generated and released to the public by OPEC on a weekly basis.

The 1973 oil crisis started in October 1973, when the members of Organization of Arab Petroleum Exporting Countries (consisting of the Arab members of OPEC, plus Egypt, Syria and Tunisia) proclaimed an oil embargo in response to the U.S. decision to re-supply the Israeli military during the Yom Kippur war; it lasted until March 1974. The effects of the embargo were immediate; OPEC forced the oil companies to increase payments drastically and in turn the price of oil quadrupled by 1974 to nearly $12 per barrel. Until the March 28, 2000 adoption of the $22-$28 price band for the OPEC basket of crude, oil prices only exceeded $24.00 per barrel in response to war or conflict in the Middle East. With limited spare production capacity, OPEC abandoned its price band in 2005 and with that deemed itself powerless to control a surge in oil prices. Oil hit a record high of $145 in July 2008.

Industry analysts have been suggesting a price of $100 per barrel for a few years, especially with much of the world's oil lying under territory prone to unrest. As well, rapidly growing demand from countries like India and China will keep the pressure on prices heading higher.

The training set used in my analysis contains historical data from January 1979 to December 2006, and the testing set contains historical data from January 2007 to April 2007. Data prior to 1979 was not used in the analysis; oil prices flat-lined at around the $3 mark for many years leading up to the aforementioned 1973 crisis. From 2007-2009, oil prices were highly unstable with rapid increases and decreases due to various unpredictable factors including poor political relations between the U.S. and large oil-producing countries such as Nigeria and Iran. I felt that this data did not add value to my analysis due to the uncertainty surrounding the events during those times.

ANALYSIS

LINEAR REGRESSION

My first step in analyzing the data was to observe the time series plot for trends and seasonality.

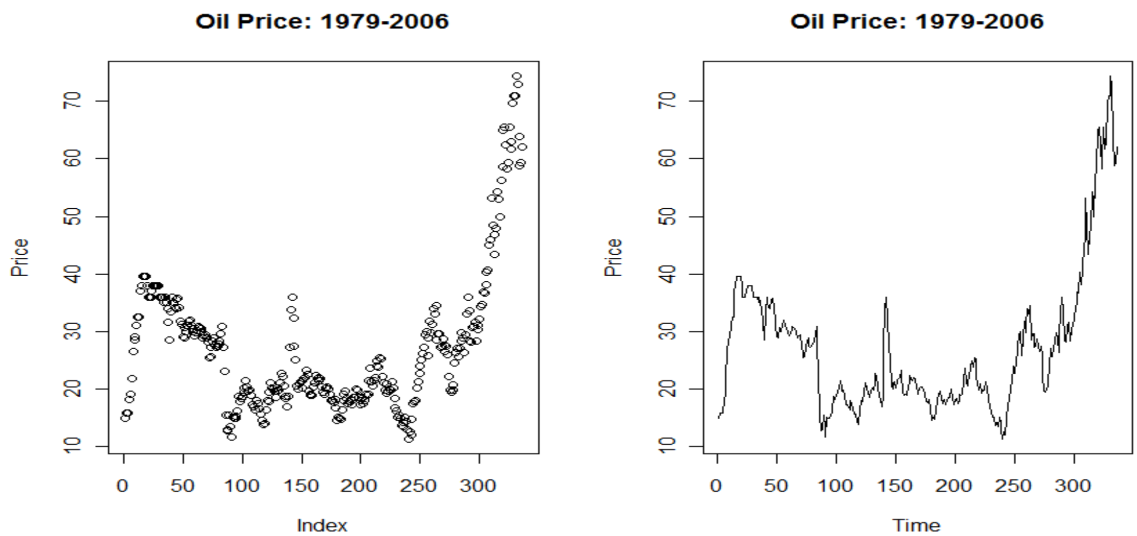

It can be seen that data is not stationary. After analyzing various models as seen in Appendix A, the one with the highest adjusted R^2 of 0.787 was of the form:

Price ≈ Index + Index2 + Index3 + Index4

Using this model, I then tested the OLS assumptions. I first tested whether the residuals were homoscedastic:

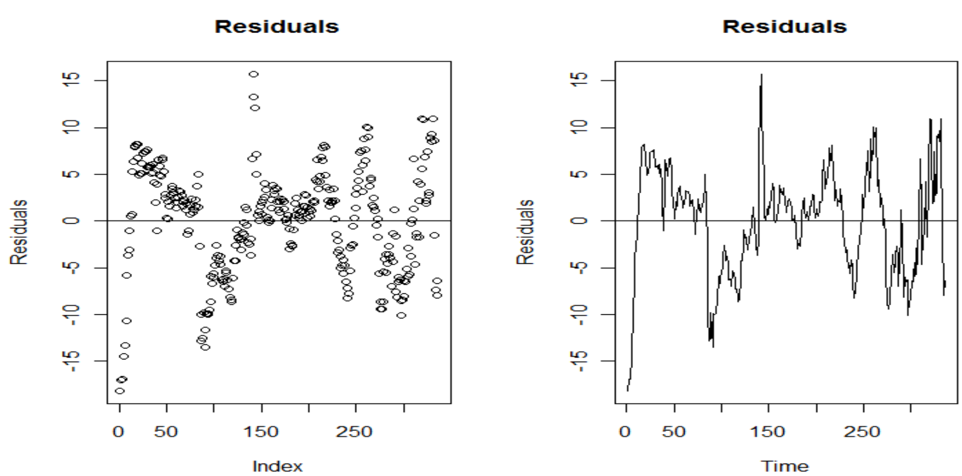

From the plots, the residuals seem to have constant mean, and fluctuate around that mean with more or less constant amplitude. Therefore, the residuals seem to have common variance, indicating that they are homoscedastic.

Then, I wanted to test whether my residuals were uncorrelated.

Test	Statistic	p-value	Conclusion
Dwtest(b)	0.1417	< 2.2e-16	Correlated residuals
Runs.test(res)	-13.9868	< 2.2e-16	Correlated residuals
Bartels.test(res)	-16.7919	< 2.2e-16	Correlated residuals

After examining the DW test, runs test, and Bartels test, I concluded that my residuals were positively correlated at lag 1.

Although I thought of applying an AR(1) model to my residuals to make them uncorrelated, I decided not to go ahead with it because the model is not very apt to begin with: R-squared is not very high and most of the variables are insignificant (see Appendix A). Thus, I decided to model my data with ARIMA, independently of regression.

ARIMA

Let us first observe the ACF and PACF plots of my data.

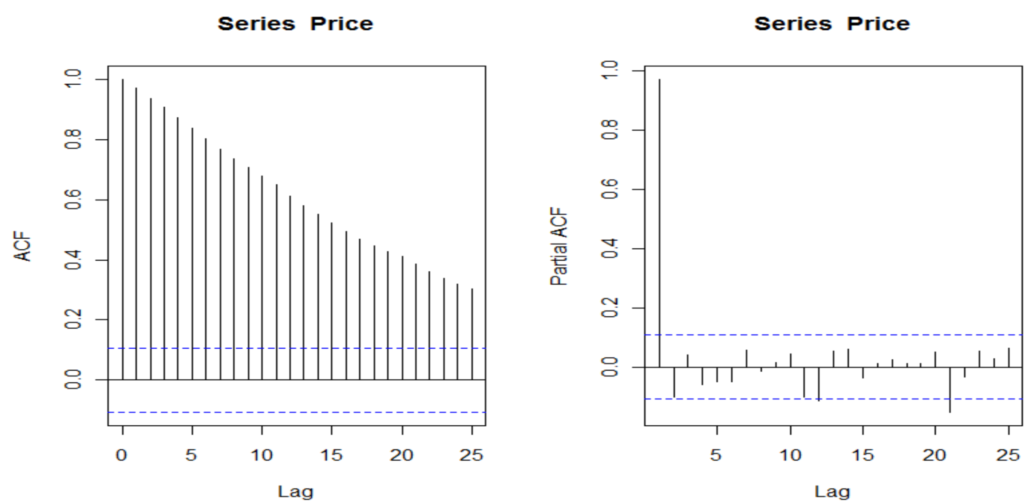

The ACF plot shows a random walk; decay is very slow and the data is not stationary. This indicates that I need to difference my data. After differencing my data, it looked more stationary than my original time series plot (see Appendix B). However, there appeared to be increasing amplitude, so I applied the variance-stabilizing transformation “log” to my data (I differenced the logged data).

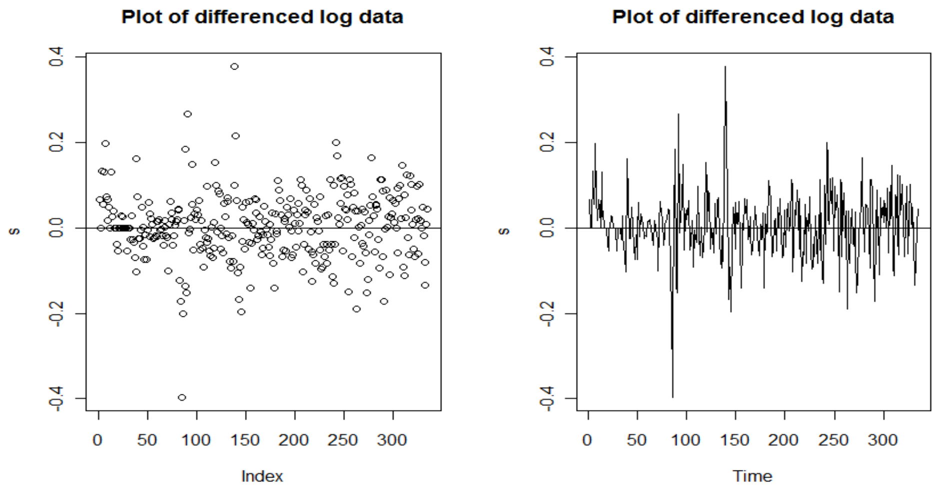

As you can see, the data is clearly stationary with a few outliers that may need to be removed. The Dickey Fuller test also shows a very low p-value, indicating that my data is stationary.

After comparing various ARIMA models as seen in Appendix B, I decided to apply an ARIMA(1,1,1) model to my data.

DIAGNOSTICS OF ASSUMPTIONS FOR ARIMA(1,1,1)

I now need to verify that my OLS assumptions are satisfied.

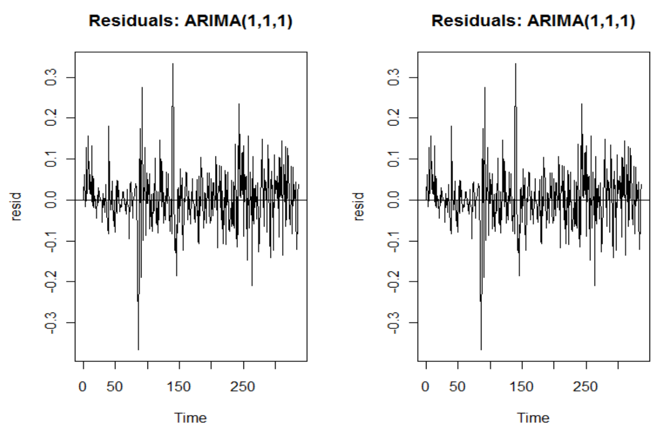

From the plots, I can see that the residuals seem to have constant mean; they do fluctuate around zero with more or less constant amplitude. The residuals also seem to have common variance, indicating that they are homoscedastic. There are, however, a few outliers which may need to be deleted.

I must now check that the residuals are uncorrelated. Please see the output below.

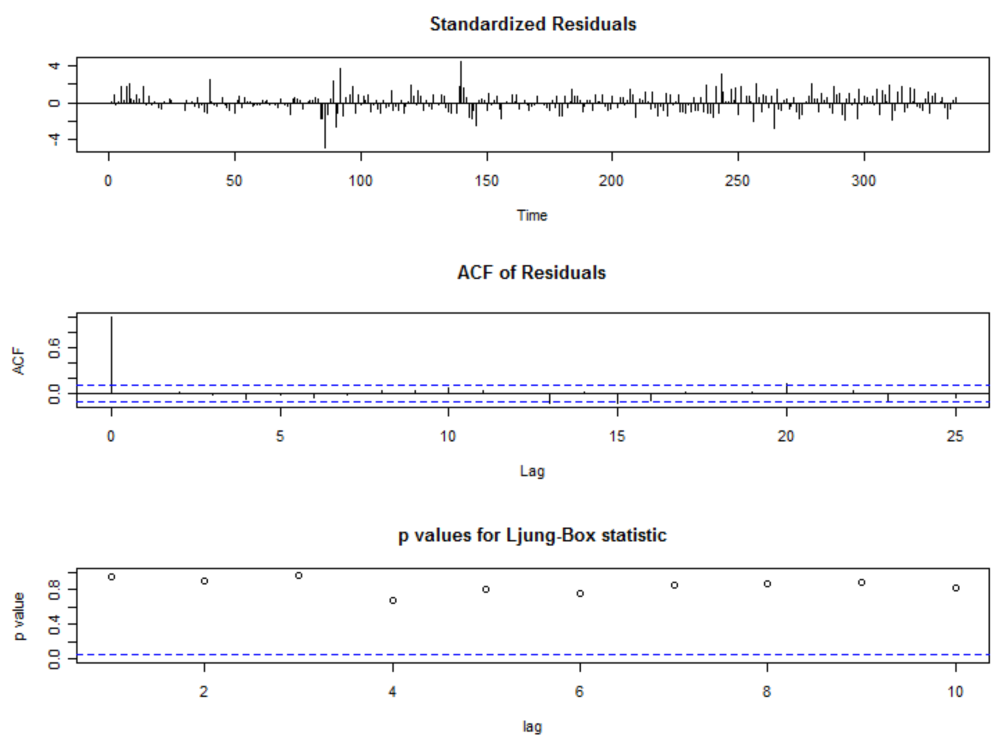

Test	p-value	Conclusion
Runs.test(resid)	0.913	Uncorrelated residuals
Bartels.test(resid)	0.4727	Uncorrelated residuals
Ljung-box	All p-values above threshold	Uncorrelated residuals

The first plot tells us that the residuals are homoscedastic, and the last two plots tell us that the residuals are uncorrelated. The runs test and Bartels test also show that the residuals are uncorrelated.

Now I must check for normality of my residuals.

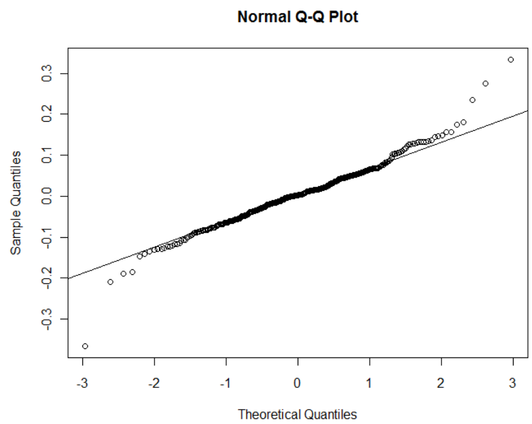

Based on the qqnorm plot, there appears to be some outliers. The Shapiro-Wilk test gives us a p-value of 1.877e-06, indicating that my residuals are not normal. I need to remove some outliers and re-test my normality assumption. After removing three outliers, I obtained the following:

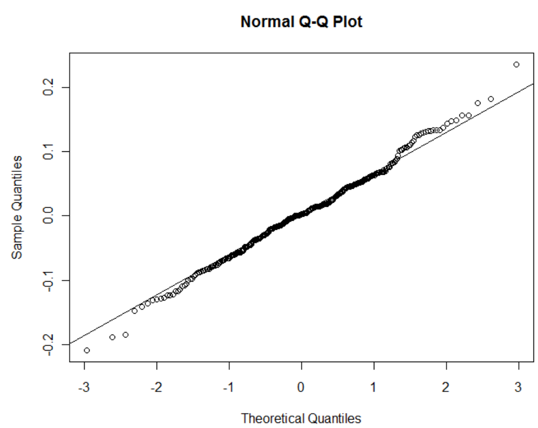

The qqplot shows a fairly straight line, indicating that my residuals are normally distributed, and the Shapiro-Wilk test gives a p-value of 0.3038.

Thus, my OLS assumptions of white noise and normality have been satisfied.

PREDICTIONS

I will now proceed to predict monthly oil prices for the first quarter of 2007, and corresponding 99% prediction intervals.

Steps ahead	Actual	Prediction	Lower bound	Upper bound
1	54.57	62.6654	60.14794	65.1828
2	59.26	62.5275	59.88695	65.16805
3	60.56	62.5572	59.83427	65.28021
4	63.97	62.5508	59.75594	65.34566

It can be seen that three of the four predictions are very close to their actual values, with the exception of January. In addition, the actual value for February is very close to falling within my 99% prediction interval, and March and April actually do fall within it.

CONCLUSIONS

APTNESS OF THE MODEL

The model I have chosen satisfies all of the OLS assumptions of white noise and normality. It also has a very low AIC, high log likelihood, and low sigma^2 comparatively. As a result, an ARIMA(1,1,1) model seems to best model my data.

PRACTICAL IMPLICATIONS

Since three of my four predicted values are fairly close to their actual values, my model has some predictive power and is fairly useful for practical purposes. However, it should be noted that my model cannot predict far into the future, so predictions for more than four months into the future may not be accurate and/or useful for practical purposes such as investing.

DIFFICULTIES AND POSSIBLE IMPROVEMENTS

The difficulties I experienced with my predictions may simply be due to the nature of the data as it is always difficult to accurately model real financial data. A method of improvement could be to use 90% prediction intervals because there would be a greater chance that my true values fit within their respective intervals since the length of the intervals would be larger. The reason I did not do this was because I believed that the standard error would be too large for the purposes of short term investing.

APPENDICES: R-CODE

APPENDIX A: REGRESSION ANALYSIS

```
> t <- read.table("OilData2.txt", header=TRUE)
> attach(t)
> plot(Price, main="Oil Price: 1979-2006", ylab="Price")
> ts.plot(Price, main="Oil Price: 1979-2006", ylab="Price")
```

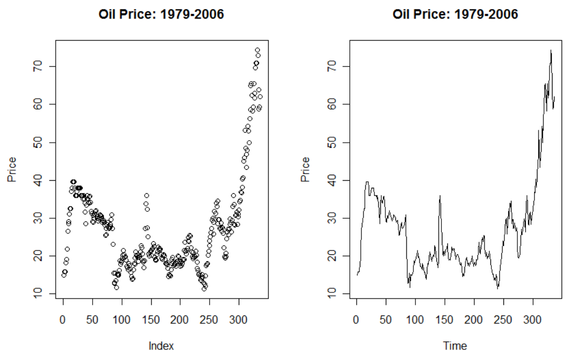

Comparison of various regression models:

Model	Adjusted R-Squared
Price~Index	0.08523
Price~Index2	0.2229
Price~Index3	0.3481
Price~Index4	0.4459
Price~Index+Index2	0.6505
Price~Index+Index2+Index3	0.7799
Price~Index+Index2+Index3+Index4	**0.787**
Price~Index+Index2+Index3+Index4+cosfactor+sinfactor	0.7858

```
> summary(b)

Call:
lm(formula = Price ~ Index + Index2 + Index3 + Index4)

Residuals:
    Min      1Q  Median      3Q     Max
-18.148  -3.769   1.071   3.460  15.647

Coefficients:
              Estimate Std. Error t value Pr(>|t|)
(Intercept)  3.310e+01  1.568e+00  21.115  < 2e-16 ***
Index       -1.057e-01  6.426e-02  -1.646 0.100812
Index2       5.698e-04  7.738e-04   0.736 0.462015
Index3      -5.751e-06  3.447e-06  -1.668 0.096202 .
Index4       1.762e-08  5.075e-09   3.472 0.000585 ***
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 5.645 on 331 degrees of freedom
```

Multiple R-squared: 0.7895,	Adjusted R-squared: 0.787
F-statistic: 310.4 on 4 and 331 DF, p-value: < 2.2e-16

<u>Homoscedasity:</u>

```
> res <- b$residuals
> ts.plot(res, main="Residuals", ylab="Residuals")
> abline(0,0)
```

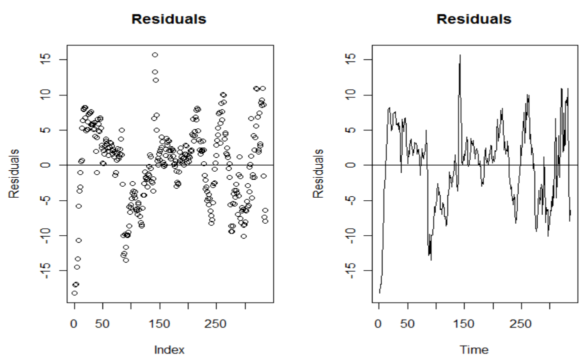

<u>Uncorrelated Residuals:</u>

```
> dwtest(b)

	Durbin-Watson test

data:  b
DW = 0.1417, p-value < 2.2e-16
alternative hypothesis: true autocorrelation is greater than 0

> runs.test(res)

	Runs Test - Two sided

data:  res
Standardized Runs Statistic = -13.9868, p-value < 2.2e-16

> bartels.test(res)

	Bartels Test - Two sided

data:  res
Standardized Bartels Statistic = -16.7919, RVN Ratio = 0.168, p-value <
2.2e-16
```

APPENDIX B: ARIMA ANALYSIS

```
> acf(Price)
> pacf(Price)
```

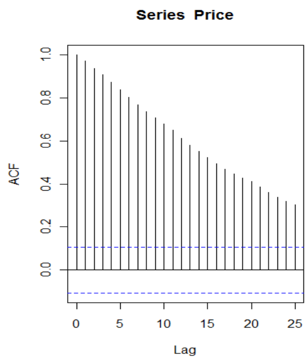

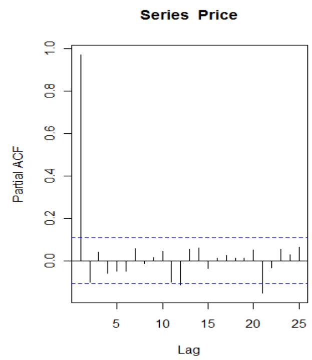


```
> s <- diff(Price)
> plot(s, main="Plot of differenced data")
> abline(0,0)
> ts.plot(s, main="Plot of differenced data")
> abline(0,0)
```

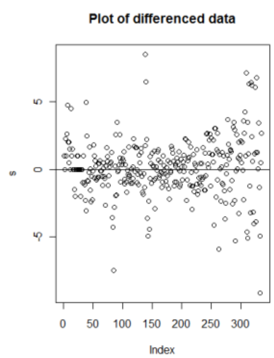

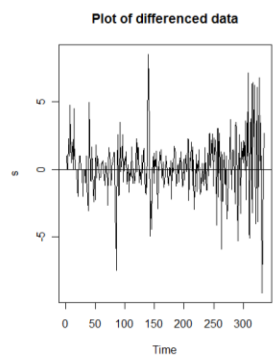


```
> s <- diff(log(Price))
> plot(s, main="Plot of differenced log data")
> abline(0,0)
> ts.plot(s, main="Plot of differenced log data")
> abline(0,0)
```

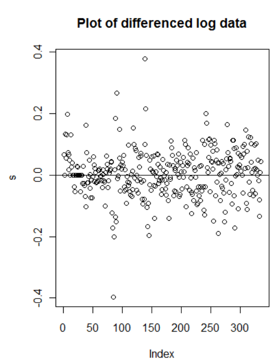

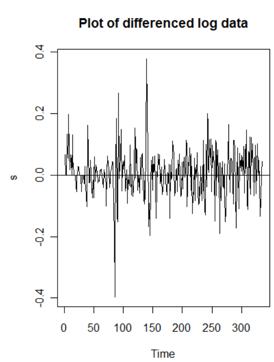

```
> adf.test(s)

        Augmented Dickey-Fuller Test

data:  s
Dickey-Fuller = -7.6361, Lag order = 6, p-value = 0.01
alternative hypothesis: stationary

Warning message:
In adf.test(s) : p-value smaller than printed p-value
```

Comparison of various ARIMA models on the differenced data:

ARIMA Model	Sigma^2	Log-likelihood	AIC
ARIMA(0,0,0)	0.006016	381.14	-758.29
ARIMA(0,0,1)	0.005503	396	-786.01
ARIMA(0,0,2)	0.005489	396.44	-784.87
ARIMA(0,0,3)	0.005483	396.62	-783.24
ARIMA(1,0,0)	0.005581	393.67	-781.33
ARIMA(1,0,1)	**0.005483**	**396.6**	**-785.2**
ARIMA(1,0,2)	0.005473	396.92	-783.85
ARIMA(1,0,3)	0.005471	396.98	-781.96
ARIMA(2,0,0)	0.005513	395.72	-783.44
ARIMA(2,0,1)	0.005472	396.94	-783.89
ARIMA(2,0,2)	0.005472	396.95	-781.9
ARIMA(2,0,3)	0.005415	397.95	-781.9
ARIMA(3,0,0)	0.005509	395.83	-781.67
ARIMA(3,0,1)	0.005471	396.97	-781.94
ARIMA(3,0,2)	0.005454	397.48	-780.97
ARIMA(3,0,3)	0.00532	401.49	-786.97
ARIMA(4,0,0)	0.005455	397.46	-782.91
ARIMA(4,0,1)	0.005448	397.68	-781.36
ARIMA(4,0,2)	0.005449	397.63	-779.26
ARIMA(4,0,3)	0.005411	398.04	-778.08

```
> a <- arima(s, order=c(1,0,1))
> a

Call:
arima(x = s, order = c(1, 0, 1))

Coefficients:
          ar1     ma1  intercept
      -0.2220  0.5251     0.0043
s.e.   0.1873  0.1668     0.0050

sigma^2 estimated as 0.005483:  log likelihood = 396.6,  aic = -785.2
```

***NOTE: ARIMA(1,0,1) with the differenced data should be the same as ARIMA(1,1,1) on the non-differenced data**

```
> slog <- log(Price)
> a <- arima(slog,  order=c(1,1,1))
```

```
> a

Call:
arima(x = log(Price), order = c(1, 1, 1))

Coefficients:
         ar1     ma1
     -0.2161  0.5211
s.e.   0.1886  0.1683

sigma^2 estimated as 0.005495:  log likelihood = 396.23,  aic = -786.47

> resid <- a$residuals
```

Homoscedasity:

```
> plot(resid, main="Residuals: ARIMA(1,1,1)")
> abline(0,0)
> ts.plot(resid, main="Residuals: ARIMA(1,1,1)")
> abline(0,0)
```

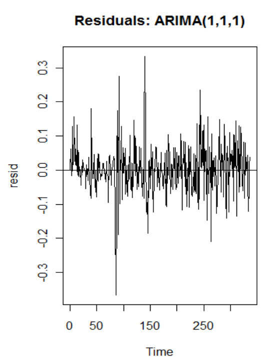

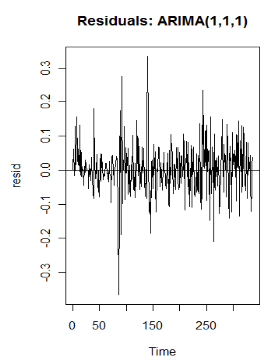

Uncorrelated Residuals:

```
> runs.test(resid)

      Runs Test - Two sided

data:  resid
Standardized Runs Statistic = -0.1093, p-value = 0.913

> bartels.test(resid)

      Bartels Test - Two sided

data:  resid
Standardized Bartels Statistic = 0.7181, RVN Ratio = 2.078, p-value =
0.4727
```

> tsdiag(a)

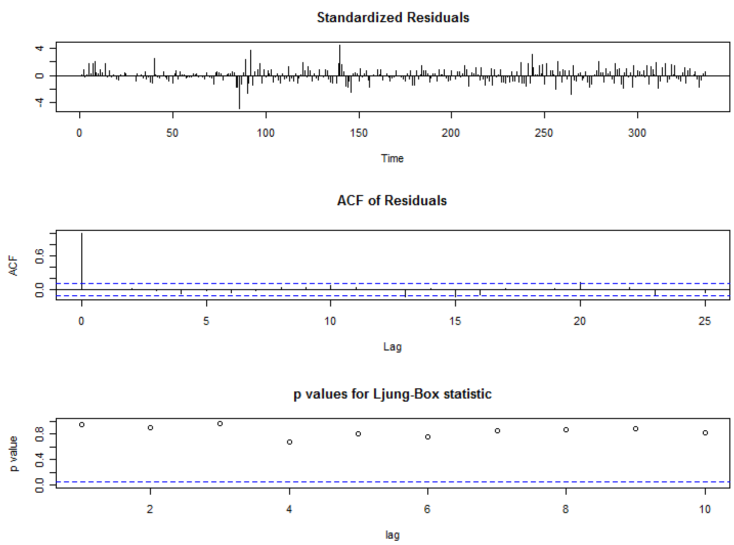

Normality:
> qqnorm(resid)
> qqline resid)

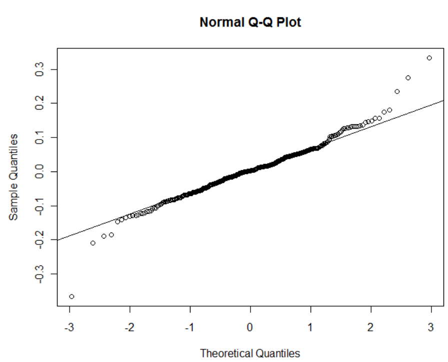

> shapiro.test(resid)

Shapiro-Wilk normality test

data: resid
W = 0.9699, p-value = 1.877e-06

> resid2 <- resid[-86] *removing the -0.36 outlier*
> qqnorm(resid2)
> qqline(resid2)

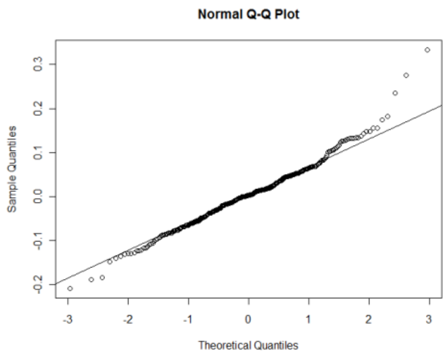

> shapiro.test(resid2)

Shapiro-Wilk normality test

data: resid2
W = 0.9799, p-value = 0.0001221

> resid3 <- resid2[-91] *removes 0.267 outlier*
> qqnorm(resid3)
> qqline(resid3)

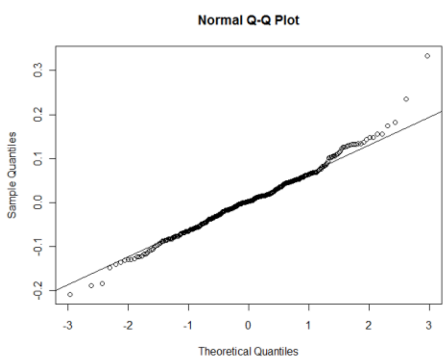

> shapiro.test(resid3)

Shapiro-Wilk normality test

data: resid3
W = 0.9848, p-value = 0.001396

> resid4 <- resid3[-138] *removes the 0.377 outlier*
> qqnorm(resid4)
> qqline(resid4)

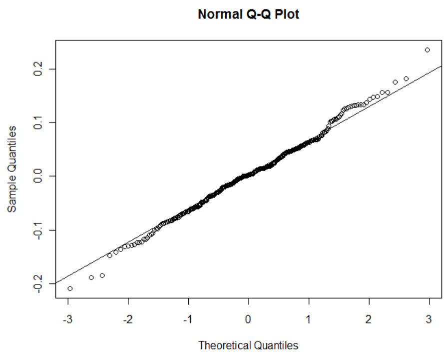


```
> shapiro.test(resid4)

        Shapiro-Wilk normality test

data:  resid4
W = 0.9947, p-value = 0.3038
```

APPENDIX C: R-CODE FOR PREDICTIONS

```
> predict(a, n.ahead=4, se.fit=TRUE, interval="prediction")
$pred
Time Series:
Start = 337
End = 348
Frequency = 1
 [1] 4.137809 4.135606 4.136082 4.135979

$se
Time Series:
Start = 337
End = 348
Frequency = 1
 [1] 0.07413142 0.12187798 0.15261518 0.17867204

> qt(0.99, 334)
[1] 2.337564
```

Summary:

Steps ahead	Prediction	Standard Error	Prediction (Converted)	Standard Error (Converted)
1	4.137809	0.07413142	62.66537	1.076948
2	4.1356064	0.12187798	62.5275	1.129616
3	4.136082	0.15261518	62.55724	1.164877
4	4.135979	0.17867204	62.5508	1.195629

**Note: Since my data is currently logged, I need to convert them back.

Summary:

Steps ahead	Actual	Prediction	Lower bound	Upper bound
1	54.57	62.6654	60.14794	65.1828
2	59.26	62.5275	59.88695	65.16805
3	60.56	62.5572	59.83427	65.28021
4	63.97	62.5508	59.75594	65.34566

REFERENCES

Amadeo, K. (2009, November). *How are Oil Prices Determined?* Retrieved June 2010, from About.com: US Economy: http://useconomy.about.com/od/commoditiesmarketfaq/f/oil_prices.htm

Barsky, R., & Kilian, L. (2004). Oil and the Macroeconomy since the 1970's. *Journal of Economic Perspectives*, 115-134.

BBC News. (2008, April 18). *Oil Reaches $117 for the First Time.* Retrieved June 2010, from BBC News: http://newsvote.bbc.co.uk/2/hi/business/7355121.stm

CBC News. (2007, July 18). *The Price of Oil.* Retrieved June 2010, from CBC: http://www.cbc.ca/news/background/oil/

Dow Jones & Company. (2007, December). *Spot Oil Price: West Texas Intermediate.* Retrieved June 2010, from The Financial Forecast Center: http://www.forecasts.org/data

Oxford Institute for Energy Studies. (2007, January 25). *Analysing Oil Prices: The Usefulness and Limitations of Existing Approaches.* Retrieved June 2010, from Oxford Energy: http://www.oxfordenergy.org/presentations/analysing_oil_prices.pdf