NATIONAL TECHNICAL UNIVERSITY OF ATHENS

ΣΧΟΛΗ ΝΑΥΠΗΓΩΝ ΜΗΧΑΝΟΛΟΓΩΝ ΜΗΧΑΝΙΚΩΝ
ΠΕΡΙΟΧΗ ΘΑΛΑΣΣΙΩΝ ΜΕΤΑΦΟΡΩΝ

«FORECASTING TIME CHARTER RATES USING ADVANCED ECONOMETRIC TECHNIQUES»

THESIS

Πολυμενόπουλος Δημήτριος

JUNE 2006

ATHENS

Πολυμενόπουλος Δημήτριος

FORECASTING TIME CHARTER RATES USING ADVANCED ECONOMETRIC TECHNIQUES

JUNE 2006

Thesis

Σχολή Ναυπηγών Μηχανολόγων Μηχανικών

Συγγραφέας: Πολυμενόπουλος Δημήτριος

Επιβλέπων: Λυρίδης Δημήτριος

ΑΘΗΝΑ

CONTENTS

TABLE OF FIGURES

Introduction

The purpose of this paper is to compare different types of time charter rate forecasting models using SAS. In this attempt, economic variables were chosen based on their relativity to the time charter market. Chosen variables include Crude Oil Purchase prices, Scrap, Second Hand and Newbuilding prices and other ship class time charter rates themselves. The thesis will try and find the statistical significance, meaning the relationship that exists between the many variables that make up the time charter market. The data used is monthly data for each ship type from October 1979 to May 2004:

- VLCC : 1 Year Tanker Time Charter Rates (250,000 DWT, 70s)
- AFRAMAX : 1 Year Tanker Timecharter Rates (95,000 SH Early 90s)
- SUEZMAX : 1 Year Tanker Timecharter Rates (140,000 SH Early 90s)
- HANDYSIZE : 1 Year Tanker Timecharter Rates (30,000 DWT)

In chapter one, the thesis will first describe the oil and time charter tanker market in order to give the reader more insight on the market itself and in sourcing explanatory variables other than these presented in the thesis. Chapter two includes the statistical methods by which both variables and explanatory variables are chosen. Chapter three will analyze the both time charter rates and their returns as well as propose an alternate data transformation that has been used in other theses as an alternative to achieving stationarity. Chapters four and five will forecast each time series using ARIMA and ARIMAX models accordingly and chapter six will use GARCH modeling for forecasting. Finally, chapter seven will compare and contrast the results of the models created in chapters four through six.

The source code is also included in a separate chapter at the end of the thesis for future reference by anyone wishing to recreate the results or continue with a more in depth exploration.

Key Words
AFRAMAX, HANDYSIZE, SUEZMAX, VLCC, TIME CHARTER RATES, ARIMA, GARCH, SAS, FORECASTING, TIME SERIES

1. The Oil and Tanker Market [1] [2]

1.1 The Economic Importance of Oil

The oil market is an environment characterized by intricate relationships between products, transportation and storage issues and environmental regulation. This importance comes not because of the large size of the oil market but because of the strategic role it plays in the economies of all countries. To get an idea about the size of the oil market, the revenues from oil producing countries account for up to 20% of the GDP while in oil consuming countries, oil imports can climb to 20% of the total import bill. These numbers have a substantial impact on national economies, especially in developing countries where energy price spikes cause adverse macroeconomic effects.

1.1.1 Crude Oil Supply

Changes in supply will have a much more dramatic effect on rates than demand. Normally, the price for the delivery of oil is higher the more immediate the need for delivery. This is true especially when oil stocks are low or are not sufficient to meet current needs. This situation is characteristic of a market in *backwardation.* On the other hand, during a time period where oil stocks are high and immediate delivery is not needed, forward prices exceed spot prices. This market situation is called *contango.* The shift from a market in backwardation to a market in contago creates the volatility that is typical of time charter rates. The reason is that, when oil stocks are high they can cover a sudden surge in demand as oil excess oil stocks can be shifted to where they are needed. Its obvious that oil stocks can't be shifted from future production to meet the immediate oil needs. This means that there's a bigger oil price increase in a market in backwardation than there is a decrease in a market in contango.

This characteristic volatility decreases the longer the forecast outlook since it is expected that supply and demand will balance out in the long run.

1.1.2 Crude Oil Demand

The primary drivers of oil tanker demand are the international trade profile, trade routes and cargo volumes. The key determinant of the trade profile is regional oil consumption, which drives crude oil demand in general. Demand is driven by *convenience yield* and seasonality. Convenience yield is directly related to the probability of a disruption in oil supplies. During times of market insecurity, industrial users may be willing to pay a premium for “immediate energy”. This is reflected in higher near-term forward prices relative to longer-term forward prices. Convenience yield is measured as the net benefit (value of uninterrupted production) minus the cost (including storage costs).

Seasonality has a large effect on heating oil. Naturally, it peaks in winter while gasoline demand is higher in the summer. Seasonal demand affects crude oil prices, although its effects are much less pronounced than other demand drivers. Demand, on the other hand, has a partial affect on the tanker rates. Even the most robust of economies grow relatively slowly and while oil consumption rises with a growing economy, the rate by which consumption increases is slow and steady. As a result, demand for oil is a relatively slow and steady process and isn’t a demand driver for profitability and thus, charter rates.

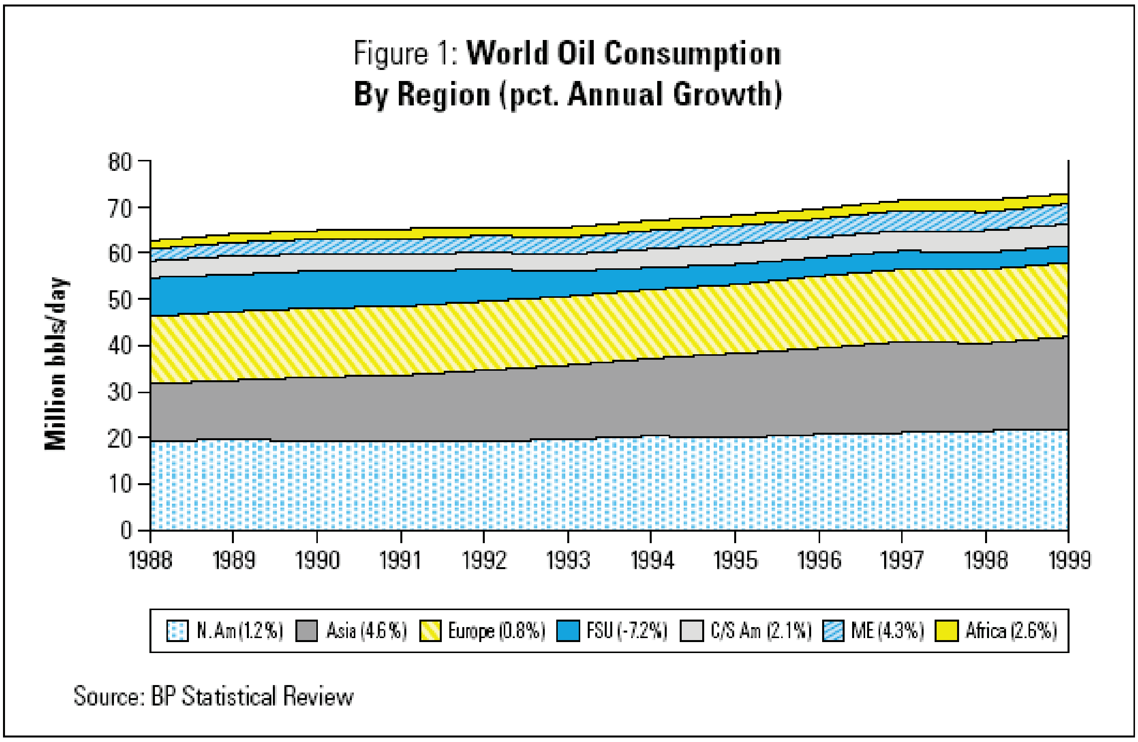

Figure 1 : World Oil Consumption By Region

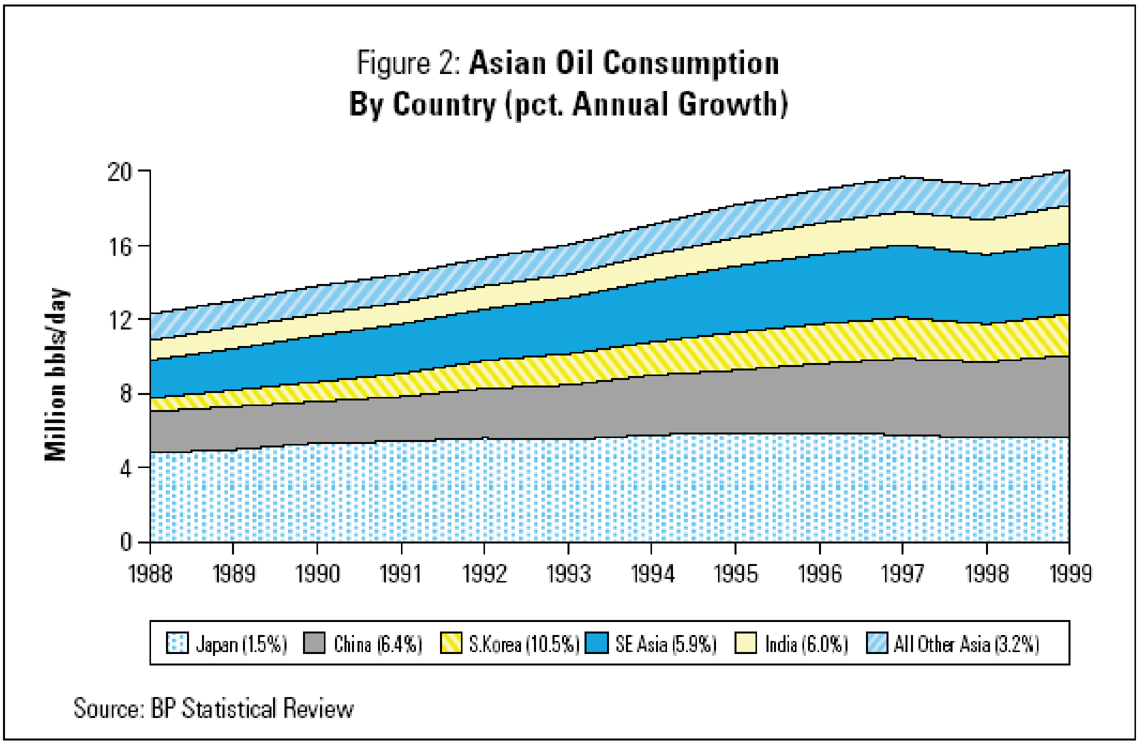

Figure 2 : Asian Oil Consumption By Country

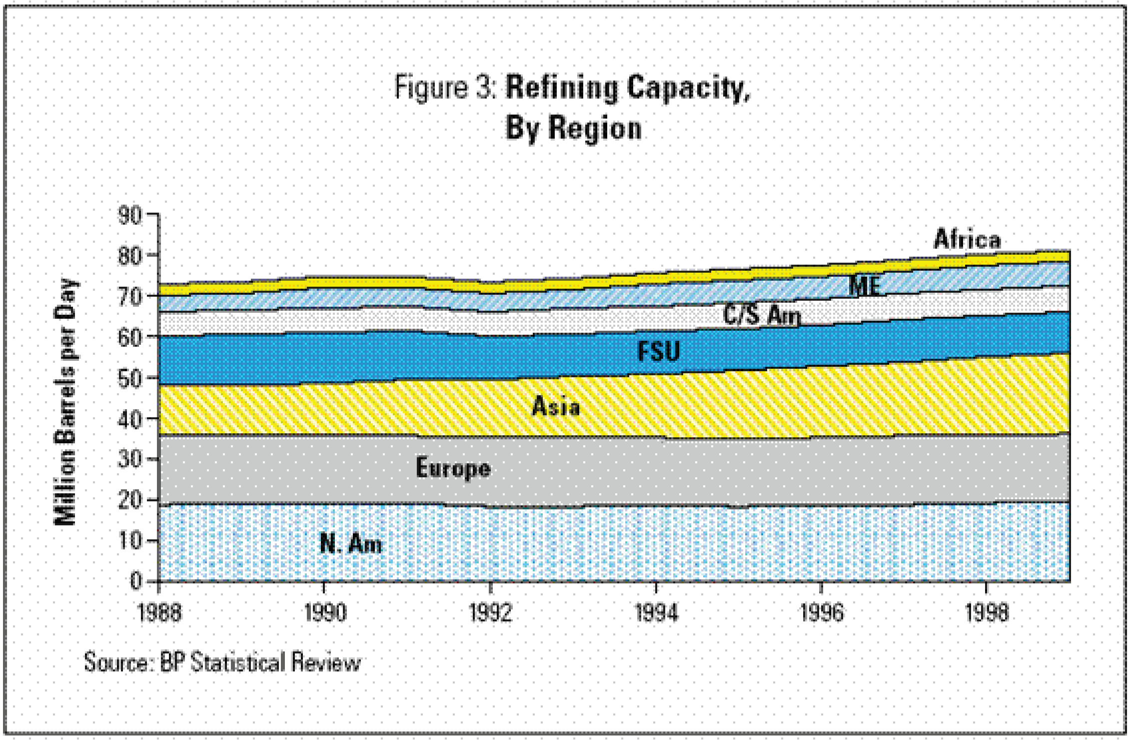

Figure 3 : Refining Capacity By Region

1.1.3 Crude Oil Refining

Another demand driver for tankers is international refining. There is motivation for building refineries and this motivation is twofold. Refineries are built to provide petroleum products to their regional markets. Refining adds significant value to oil and, in the long term, oil consuming regions benefit from refining oil locally rather than pay a mark-up on imported products. This *consumer-proximate* refining is beneficial to crude oil tanker demand since it requires crude oil to be shipped, often long distances, to its point of refinement. Still, there is also motivation by the oil producer to add value to its product. This means that an oil producing region would prefer exporting refined oil rather than using it for regional requirements. This producer-proximate refining strategy would necessarily reduce the need for crude oil tankers. Figure 3 shows that, at least for now, there is a consumer-proximate market where refining takes place close to end users (North America, Europe and Asia) instead of the oil producing sources like the Middle East.

1.1.4 Crude Oil Production

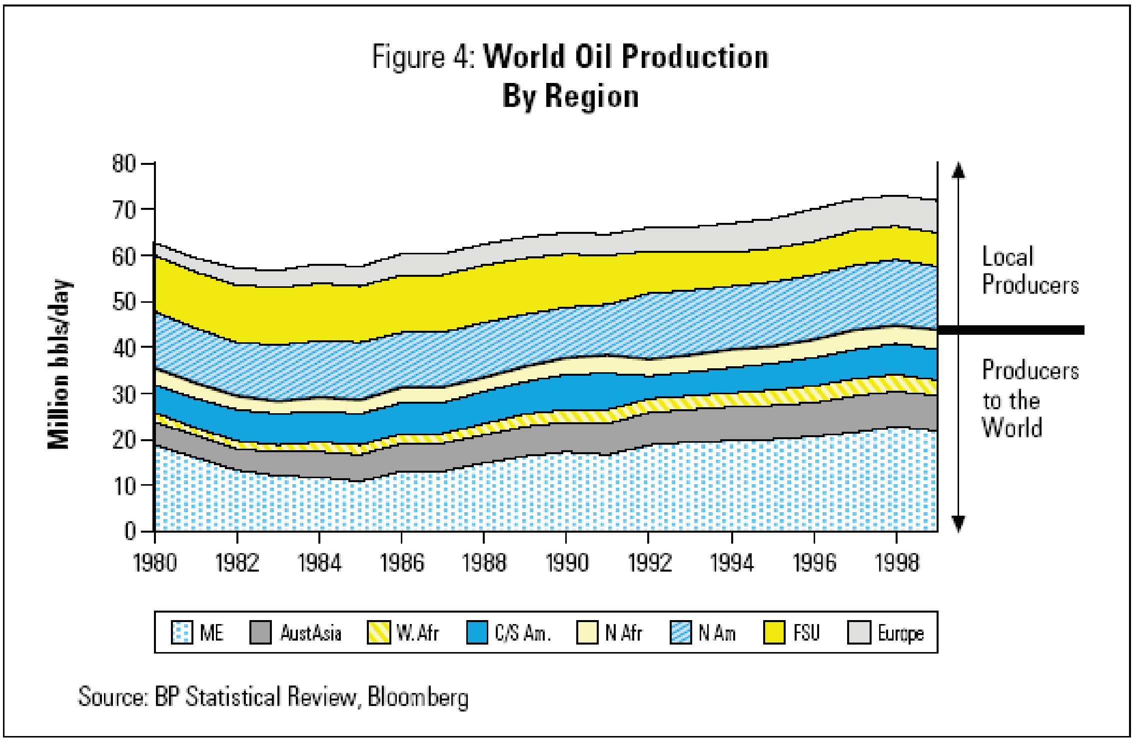

Figure 4 : World Oil Production By Region

Most crude oil comes from one of two types of producers:

Producers to the World: This is the most important group of producers from a shipping perspective. This group typically has large reserves but very little domestic use for crude oil produced; and must therefore rely on exports. This group resides in five major regions:

- Middle East (OPEC) – is the largest global producer of oil (about 1/3 of total production)
- Asia/Australia (mainly Indonesia) – with production mainly intended for Asian use.
- West Africa – small region that produces 4 % of the world total but with growing importance.
- Central and South America (Venezuela, Colombia and Ecuador) – oil is produced mainly for the US Gulf and East cost refineries.
- North Africa (Algeria and Lybia) – oil is produced mainly for Europe but is also characterized by the shorted impact on shipping since it represents the shortest-haul exporter.

Local Producers: Large reserves found and produced closer to end-users. Production goes directly to domestic refineries, or is exported mostly to neighboring or very proximate markets. There are three key regions :

- North America – US, Mexico and Canada oil is refined exclusively in local U.S. refineries. Offshore and land-based production is shipped via pipeline or through short haul trade-routes.
- Former Soviet Union – production is intended mostly for domestic and European use although there are some sea based trade-routes.
- Europe – Although European oil is geared for shipment to Northern Europe, there are significant amounts of long-haul export to the east coast of North America and to Southern Europe.

1.2 Oil Prices (1947-2003)

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 OPEC and non-OPEC supply. What's obvious is that major economic and world events alone aren't to blame for dramatic shifts in crude oil purchase price. For instance, the combination of the Israeli-Palestinian conflict (Yom Kippur War), the overthrow of the Shah in Iran and subsequent Iranian-Iraq war drove the price of crude oil through the roof. On the other hand, the Gulf war and the 9/11 attacks, while major in scale (and having major global consequences) kept the price at nominal levels. What we're seeing now, as prices again hover in record high levels, is that the market 'feeling' of whether events **may** unfold (in Iran, Iraq, Nigeria, Saudi Arabia or Venezuela) plays a role as crucial as if a crisis were actually taking place. As a result, the change in oil price has to be a combination of economic and geopolitical events as well as the result of the oil market's psychology.

Economic and geopolitical factors include:

• Production: An increase or decrease in oil production should affect tanker demand directly. The greater the output, the greater the demand for oil tankers. Production depends on the quotas OPEC and Russia set and whether it complies with them. Also, the discovery and/or further development of oil reserves (e.g. China) will also affect production.

• Geo-political events: There are several non-market factors that effect oil production and trading:

– Iran's nuclear aspirations and its use of oil as a leverage tool will undoubtedly affect oil prices in the future.
– The future of the Iraqi state (and its subsequent oil production) is also of vital importance.
– Balkan and Ural military/political situations.
– Venezuela: Chavez regime's relationship with PDVSA and the trend for energy nationalization sweeping over Latin America.

• The economic conditions of the major importers are also important.

– Will the economies of Japan and the US's grow or shrink?
– China is currently a relatively small consumer. As it becomes more industrialized, plans will be drawn up to open or upgrade refineries.

• Geographic shifts in refining patterns may occur. Such changes may have an impact on tanker tonne-mile requirements.

– In crude oil importing countries, there may be changes in the amount of oil refined domestically (producer-proximate) vs. oil refined elsewhere and imported (consumer-proximate). For instance, Japanese refiners may not continue to be as competitive in producing for domestic market. This would result in increased product import there. In the U.S. and Europe, the effect of decreased production on crude imports is a little more ambiguous, as a significant amount of crude oil feedstock comes from domestic or nearby production.
– Look for trends towards refining closer to crude oil sources. Many Long Distance Carriers (particularly in Asia) are planning expansion of domestic refinery capacity. Japan would be the likely export target. This would result in a small reduction in crude oil tanker demand overall, as the same amount of crude feedstock would be transported a shorter distance (AG-SE Asia vs. AG-Japan). Also, Middle Eastern producers are planning to build large refineries close to their crude oil sources. Economic and political considerations make it uncertain that many of these projects will be completed. If these projects are realized, this would significantly reduce crude oil tanker demand per barrel produced.

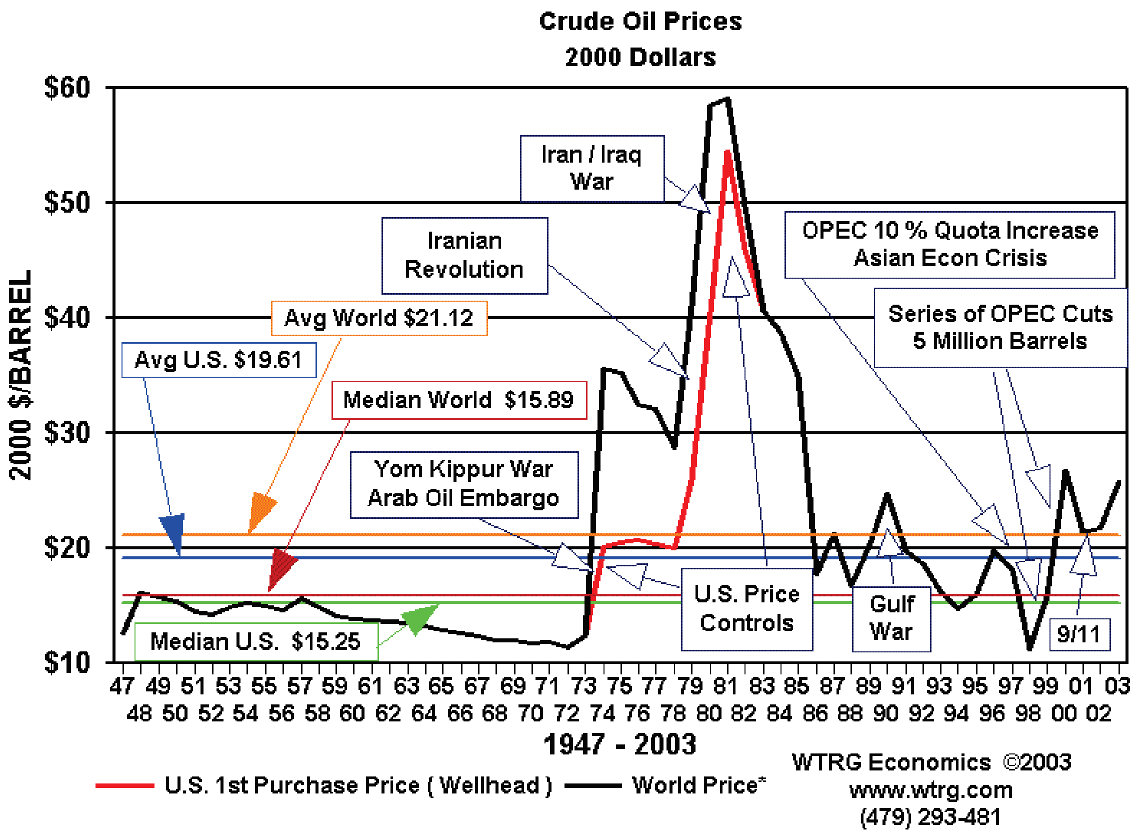

Figure 5 : Crude Oil Prices (2000 Dollars)

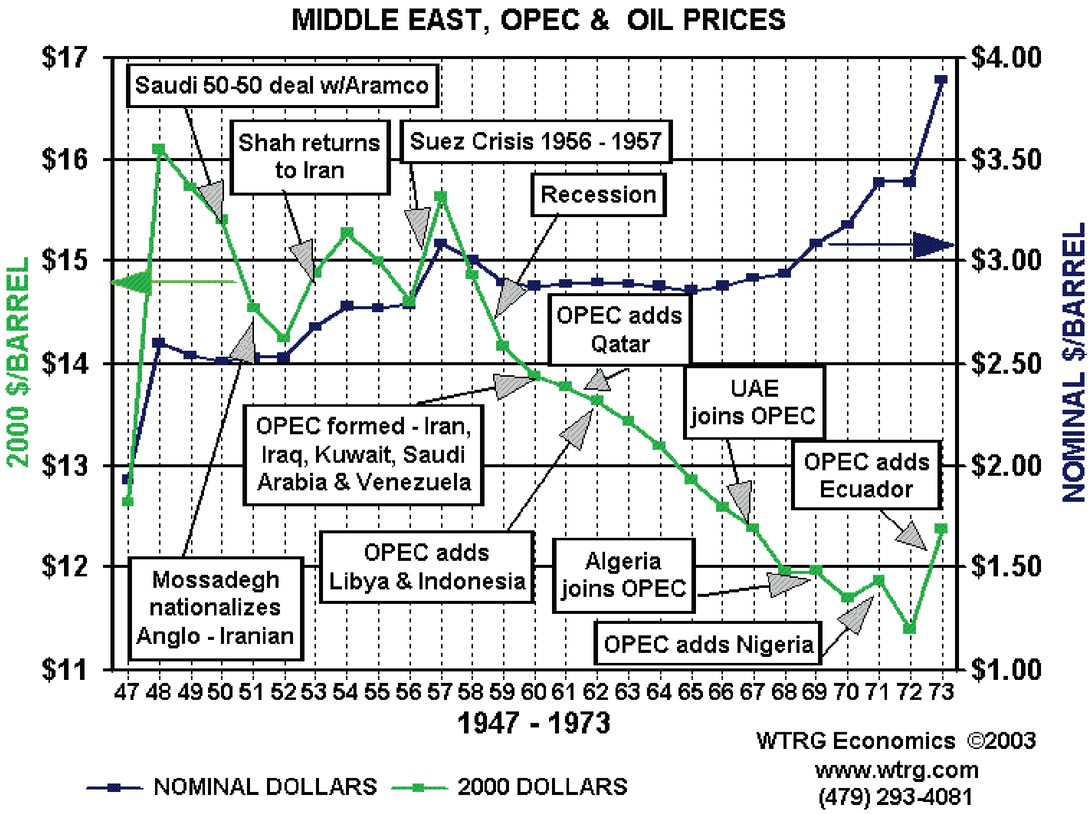

Figure 6 : Middle East, OPEC & Oil Proces

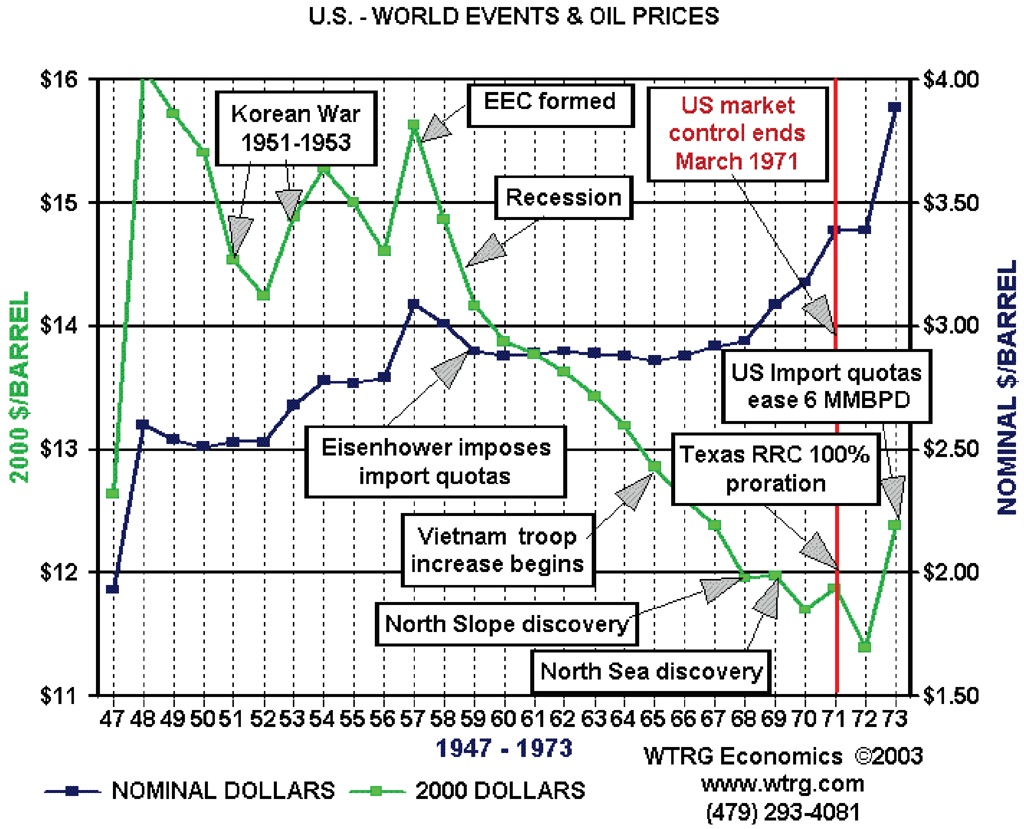

Figure 7 : U.S. - World Events & Oil Prices

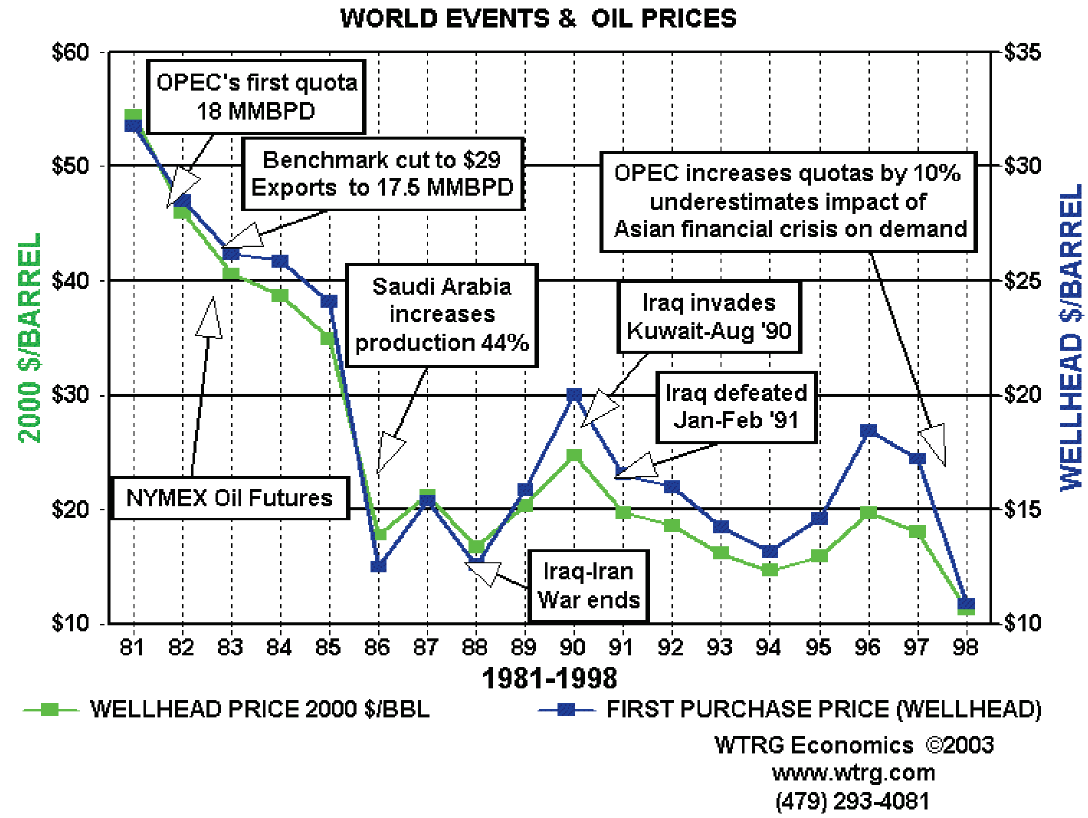

Figure 8 : World Events & Oil Prices

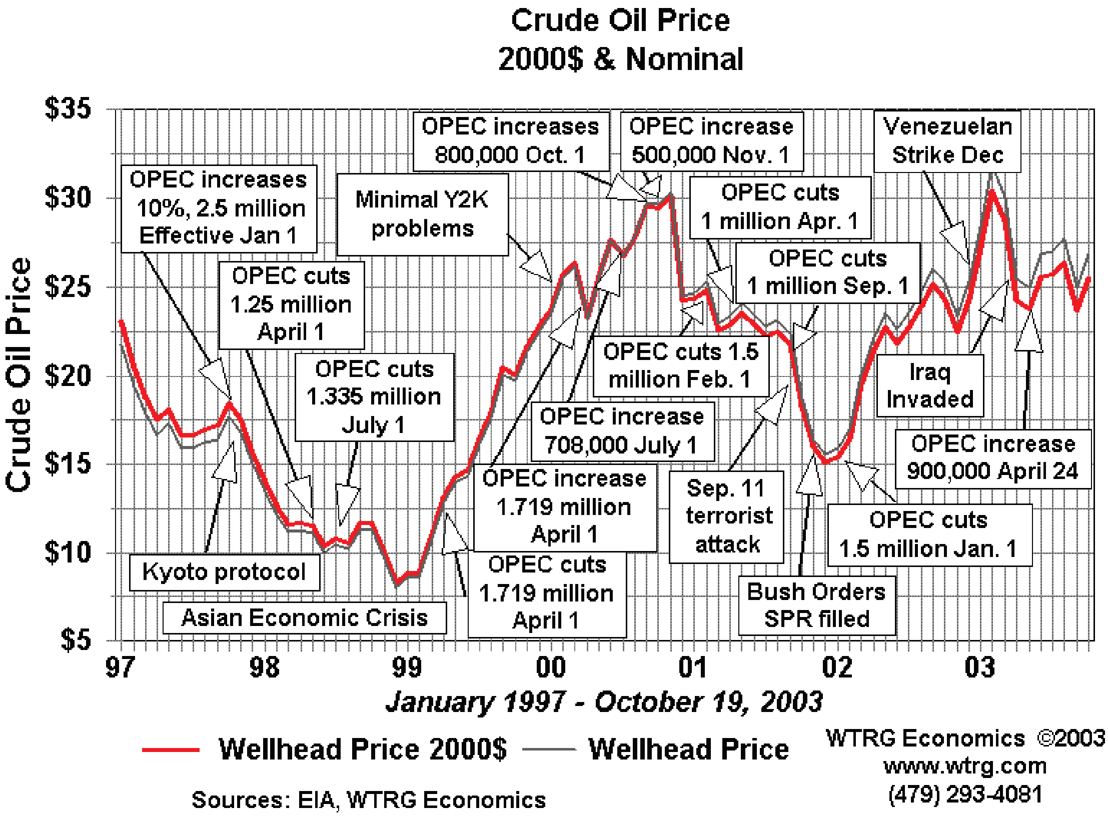

Figure 9 : Crude Oil Price $2000 & Nominal

1.3 The Tanker Market

With this brief introduction on oil, we can now turn to the market behind its shipment from oil producing countries to oil refining countries or consumers. In areas where production, refining and consumption are local, pipelines and land-based transportation is used. The majority of oil though is usually exported from oil producing countries either as crude or after having undergone refinement, by tankers that supply all five continents. Crude oil tankers are split into three main types based on their size. There are, of course, tankers other than the aforementioned but we will focus on the three basic sectors for this thesis.

The Very Large Crude Carrier (VLCC) is the largest tanker. The VLCC is dominant as far as total tonnage is concerned. Taking into account the tanker's large size, one would expect VLCC carriers to be relatively few in respect to the other classes. That is not the case, as VLCCs number just a bit less than Aframaxes. The VLCC carrier is the workhorse of the tanker fleet. VLCC's are the least specialized, commercially, of all crude oil tankers. That is why the VLCC charter rate is viewed an indication of the strength of tanker markets in general. While their time charter rates exhibit extreme volatility, modeling the VLCC carrier aims to keep into the future of the tanker market in general.

The second class size is the Aframax tanker. Basically, the Aframax is employed in the major crude oil routes that cannot be accessed by the VLCC or Suezmax due mostly to port restrictions. It is by far the tool-of-the-trade in the Caribbean to U.S. East Coast (shallow ports). Aframaxes are also quite prevalent in the Indonesia to Japan trade routes.

The next size class down is the Suezmax. In terms of fleet size, this is the smallest of the three classes. The Suezmax is almost synonymous with the West Africa to U.S. Gulf and East Coast trades. It is also popular in short-haul Cross-Mediterranean trades and for Caribbean cargoes.

The Handysize is the smallest sized tanker in the market. These vessels are a maximum of 10,000 DWT to 30,000 DWT and are more maneuverable and have shallower draft than larger vessels. As a result, they make up the majority of the world's ocean-going cargo fleet.

1.3.1 Supply and Demand

A few decades ago, most oil tankers were the property of oil companies. Today, oil companies own only 25%-30% of the world fleet while the rest of the fleet belongs to ships owners who offer their ships for hire. This change of philosophy is the result of a number of factors that have come into play in the last years.

- The oil companies now prefer to charter ships from independent ship-owners rather than spend money building them.
- Smaller, more specialized ship-owners now offer better services.
- Using chartered ships, oil companies are more flexibly to change shipping strategies.
- Independent ship-owners act as a intermediary between oil companies when there’s a disparity between crude oil production and refining capability and companies mediate an exchange of oil.
- The ability to outsource oil shipments also cuts down on much red tape.

Contrary to the oil market, the demand for tankers can shift 10-20% a year. There are several factors that play a part in this shift.

- Global Economy: Growth or shrinkage of developed and developing countries plays a part in the need for tankers.
- Global consumption and reserves.
- Sea trade: The seasonality of certain products increases the need for tankers.
- The average distance between each shipment (average haul)
- The shipping cost: Larger and more modern ships will lower the shipping cost.
- Geopolitical events.
- Alternate shipping routes (land-based).

Tanker supply, on the other hand, shifts relatively slowly. The fact that a ship requires several years to be built is in direct conflict with the needs of a dynamic market. The tanker supply/demand models exhibits this disparity.

1.3.2 The Market Cycle

The shipping market is made up of three cycles that repeat each other with differing intensity. In the first phase, there is an abundance of tonnage. Ships usually employ a technique called “slow steaming” which literally means going slower than usual in order to cut fuel costs. This of course makes the journey last longer but since demand is low and supplies are high, there’s no rush. During this phase, time charter rates fall to the point of basically being equal to the ship’s running costs. As a result, old or badly maintained ships with elevated working costs are scrapped. The low charter rates, combined with the monetary needs of ship-owners force them to sell their ships at prices that, in older ships, can fall to demolition costs.

The resulting first phase leads to a reduction of tanker availability and the supply-demand scale begins to even off. Generally, the first sign of a market comeback is that time charter rates are raised above the ship’s running cost. By now the tanker fleet has grown smaller as many ships have been scrapped and the prices of used ships slowly creep up.

In the third phase, with the world shipping fleet being at a minimum, there's a race to ship as much as possible in the least amount of time. The profits are high and as a result banks give out loans more readily, and new ships are usually commissioned in this phase. After a period of weeks or even years, the third phase will ultimately give way to the first phase and the cycle repeats itself.

1.4 The Charter Market

There are two types of charters - in a voyage charter, the ship owner charters his ship for a pre-determined price, type, amount of cargo and time-span. In a time charter the ship is chartered for a pre-determined amount of type without any restrictions on cargo size, type or even status (it can be sub-chartered to someone else). The charter rates are more or less set by the laws of a perfect market where supply and demand bring the price of the rates up or down. The owner of the ship supplies the crew and maintains the ship while the charterer is free to chose the ship's voyage, the amount of fuel, port taxes and everything else that has to do with the economic and working costs of the ship during its journey. Also, the time charter agreement includes certain ship parameters (like speed and fuel consumption) that the owner has to guarantee. In case the ship doesn't meet these parameters, the charterer is given the ability to cut or renegotiate the contract. Time chartering is very lucrative for a ship owner who wants a low risk investment. That's because, baring any unforeseeable tragedy, its easy to estimate the costs that the ship owner will have to undertake. On the other hand, the charterer doesn't have to (or may not be able to) buy a ship for his shipping needs. As a result, both ship owner and charterer's needs and interests are met with the time charter market.

1.5 Data and methods

1.5.1 Data

The brief explanation of the tanker and oil markets is meant to give the reader a general idea about the type of data that will be collected in the model. As it has been shown, the kind of data that can be collected can vary wildly. From oil, to economic or even weather data, the mathematical model isn't picky about what we can use as long as this data has certain properties, either inherent or induced. We must shift now from economic theory to a more practical introduction to the kinds of statistical properties that model data must have. In order for model data assumptions to hold, data must be transformed so that they

- Have zero mean
- No autocorrelation (stationary)

Also, the error terms of the fitted model are expected to show:

- Zero mean
- A constant and finite error variance over all x_t values.
- The errors are statistically independent of one another
- No relationship between the error and corresponding variate
- A normally distributed error term.

Mean and variance are well known statistical definitions and won't be analyzed in this thesis. On the other hand autocorrelation is not a familiar engineering term and since it plays an important role it will be analyzed here.

A distinguishing feature of the autoregressive process is its long 'memory'. Suppose that the current value t of a series z_t can be expressed as a liear function of the previous value of the series plus a random shock a_t. If we denote the previous value of the series by z_{t-1} this process is written as

$$z_t = \phi_1 z_{t-1} + a_t$$

Where φ is an autoregressive parameter that describes the effect of a unit change in z_{t-1} on z_t. The term a_t is a random shock, otherwise known as an error or white noise series and is assumed to be independently distributed with zero moean and constant variance. Suppose we successively eliminate the lagged z_t's. That is, substitute

$$z_{t-1} = \phi_1 z_{t-2} + a_{t-1}$$

into the previous equation and get

$$z_t = \phi_1^2 z_{t-2} + a_t + \phi_1 a_{t-1}$$

then substitute

$$z_{t-2} = \phi_1 z_{t-3} + a_{t-2}$$

into the previous equation and get

$$z_t = \phi_1^3 z_{t-3} + a_t + \phi_1 a_{t-1} + \phi_1^2 a_{t-2}$$

And so on until we get the *error-shock* form

$$z_t = a_t + \phi_1 a_{t-1} + \phi_1^2 a_{t-2} + \phi_1^3 a_{t-3} + \ldots$$

Thus the autoregressive model is rewritten as the sum of the current error and an infinte number of error terms. Therefore the current observation, z_t, is still influenced by shocks, a_t, that occurred in the distant past. If the process is *stationary*, then $|\phi_1| < 1$ and the effect of the shock will gradually dissipate. If this is not the case and past shocks influence present values then the process is *non-stationary* or *autocorrelated*. The graph that describes this stationarity is called the *autocorrelation plot*.

One quick way to discern whether a series is stationary or not, is by the way that the autocorrelation function decreases slowly or not. A quick decrease of the autocorrelation function indicates a stationary series. A lingering autocorrelation plot indicates a non-stationary series (see figures 10 and 11 below).

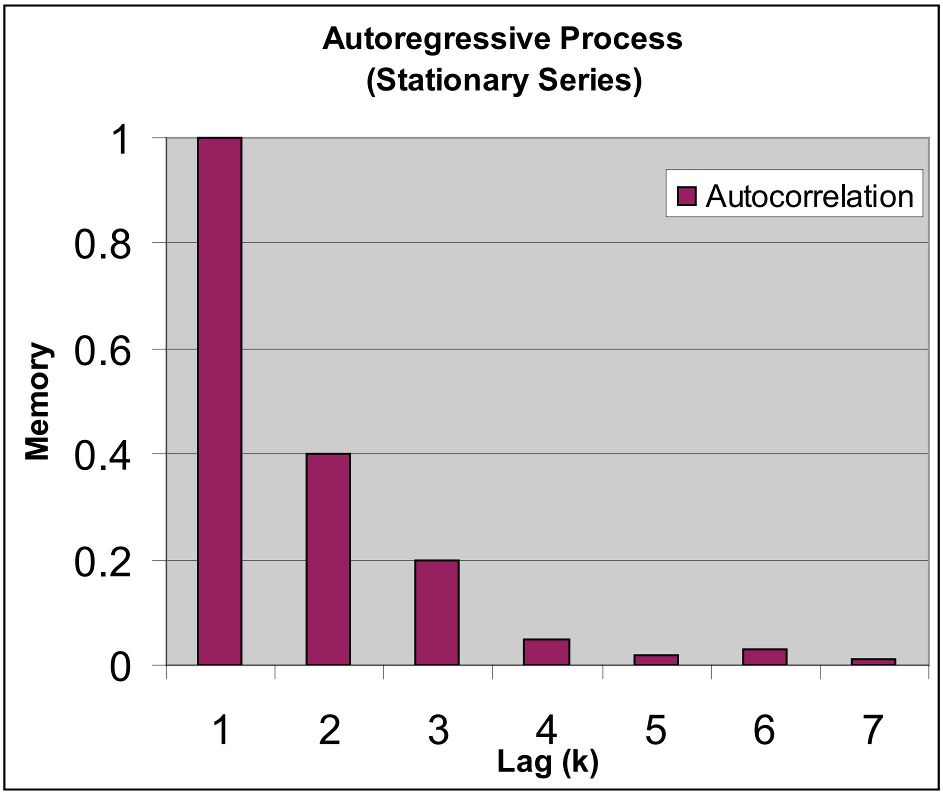

Figure 10 : ACF of a stationary series

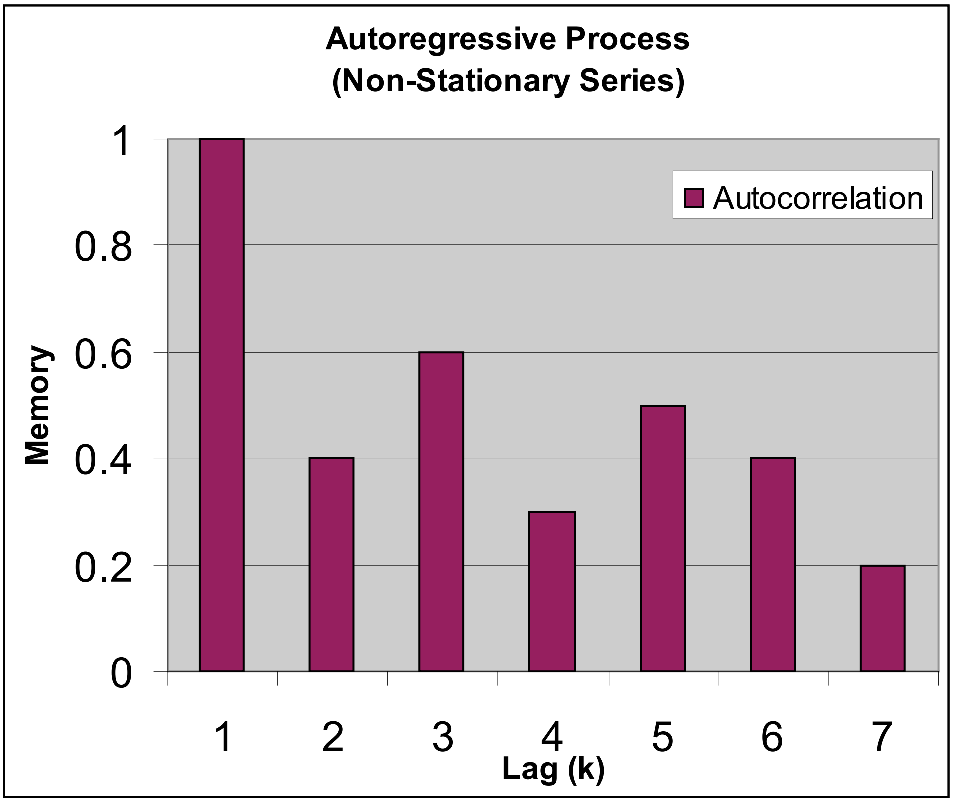

Figure 11 : ACF of Non-Stationary Series

1.5.2 Choice of Explanatory Variables

The demand for shipping services is driven by the sea trade market. As trade increases, there is a higher demand for shipping services which translates into a higher freight rate. There can be many explanatory variables that can be used to model economic this activity. Models can be based on data from GDP growth or on commodity production levels. Unfortunately, these kinds of indicators have serious weaknesses. The data and the sources that give out this information is non-constant and production levels are usually prone to adjustment at any time period after their initial announcement. For instance, GDP levels can be readjusted many times in a year or even some years after their publication. Commodity production levels can also suffer from the same kinds of adjustments since countries may have to understate or overstate their production numbers in order to meet certain economic and/or political objectives. For instance, many OPEC members produce more oil to increase their profits but understate the official quantity so that they will comply with OPEC production quotas that have been pre-agreed upon. This means that variables indicating production quotas may not be accurate and in fact might officially show an decrease when in fact there's an increase in the production level.

On the other hand, a reliable indicator is the commodity *price* which is not a 'figure' but is in fact determined by forces outside the producers' control – by supply and demand. These indicators describe the market better and thus are more reliable indicators. Unfortunately, many commodity prices and most notably oil, are influenced by unforeseen geopolitical events or technical problems which influence production supply. Such events can be depicted in our model as dummy variables and are expected to have either a positive or negative sign, depending on the situation.

1.5.3 The Worldscale Rate

In tankers, prices are quoted in Worldscale rates. Worldscale is an abbreviation of the Worldwide tanker nominal freight scale. The creation of the Worldscale rate arose from the need to make fixtures in different types of trades and vessel classes directly comparable. The Worldscale is a cost based schedule and is re-calculated on an annual basis for a full cargo for the standard vessel based upon a round voyage from loading port to discharging port and return. This means that when changes occur in the bunker prices, the port dues or the exchange rate of the currency of the States included in this route, WS100 for this year will be different in dollar terms than WS100 for the same route the previous year. Consequently, discrepancies will occur between the WS and the actual rates obtained by the vessels over the years thus distorting the final results. Furthermore, while in Worldscale terms one year may look better than another, a different story may appear if the Worldscale flat rate (WS100) is adjusted for increases in voyage costs. While the WS rate may be higher, the actual dollar per tonne rate may be lower, thus depicting worse rather than improving market conditions. For this reason it was decided that dollar per tonne rates are a more reliable indicator than WS rates and so WS rates, though available, were not used.

1.5.4 Forecasts

As was explained beforehand, the forecasting ability of the timecharter models is far better than that of a model based on spot rates. This is attributed to the fact that freight rates are inherently stochastic while the WS rate is more deterministic, making its forecast a difficult task.

Using CLARKSON's database, a collection of various types of data concerning the VLCC carriers has been obtained. Each monthly time series has its own starting point due to the fact that data from earlier years may be either scarce or unreliable. As a result, a common starting point for all time series has been set at 1979. Another factor in data selection process is the length of time between each observation. Some variables change on a monthly basis while others are slower to fluctuate. This poses a problem when trying to combine the two types of variables into a common model. For instance, the average economic life of a ship is about 25 years, only a small proportion of the fleet is scrapped each year, so the pace of adjustment to changes in the market is measured in years not months. On the other hand the price of oil can fluctuate significantly.

1.5.5 Crude Oil Purchase Price vs. Time Charter Rates

By modifying the data to fit the graph, we can see a rough sketch of the time charter rates with respect to crude oil prices. This simple graphical approach shows that time charter rates sometimes follow crude oil prices, but sometimes don't.

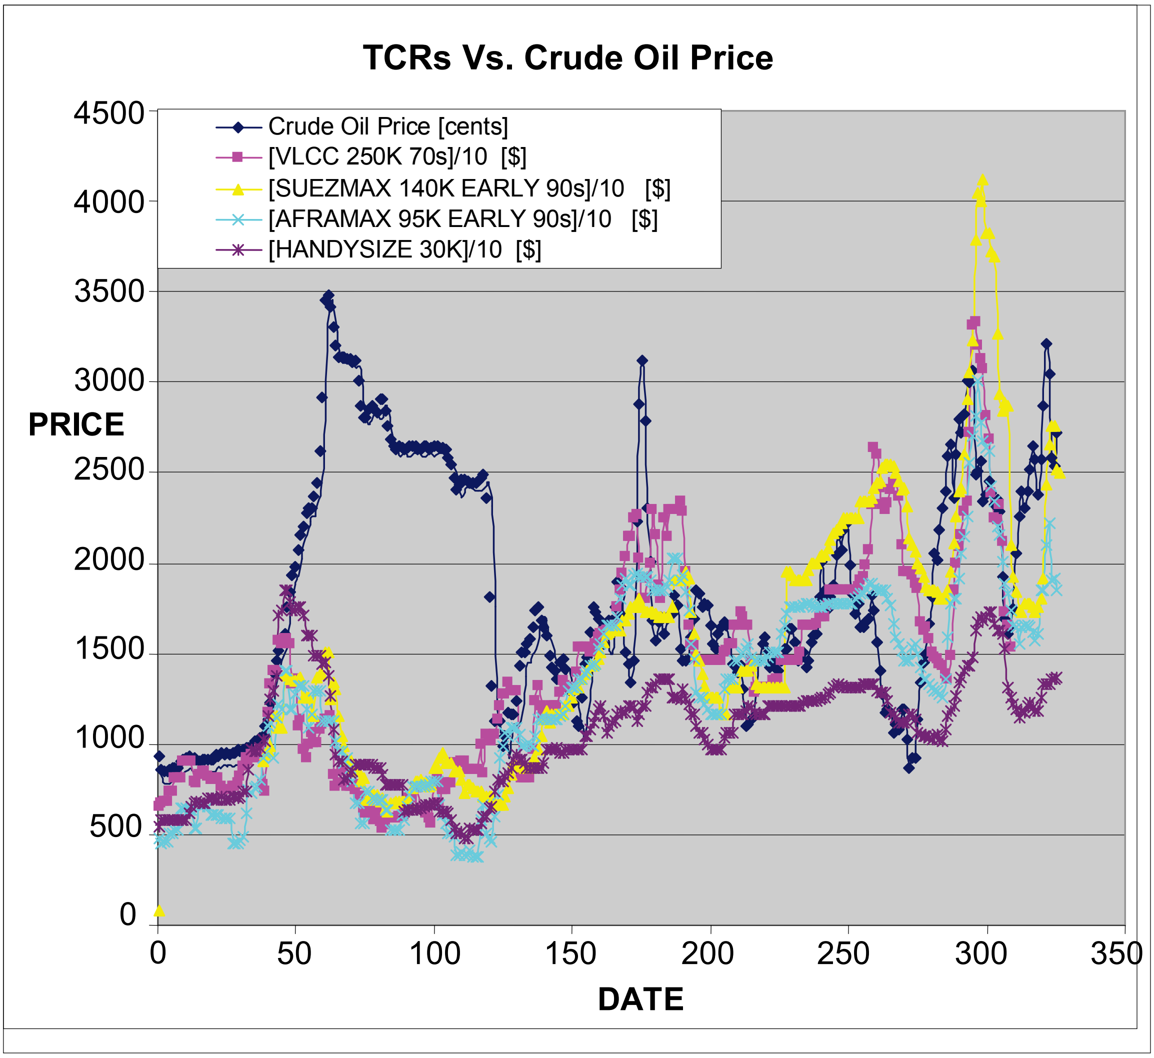

Figure 12 : Time Charter Rates & Crude Oil Price

2. Statistical Methods

2.1 Ordinary Least Squares Method

The simplest statistical method is the Ordinary Least Squares (OLS) method. This method implies a simple linear relationship between dependent (X) and independent (Y) variables.

$$Y_i = \beta_0 + \beta_1 X_i + \varepsilon_1 \quad \text{where i = 1…n}$$

ε_1 is a random variable with zero mean and σ^2 variance.

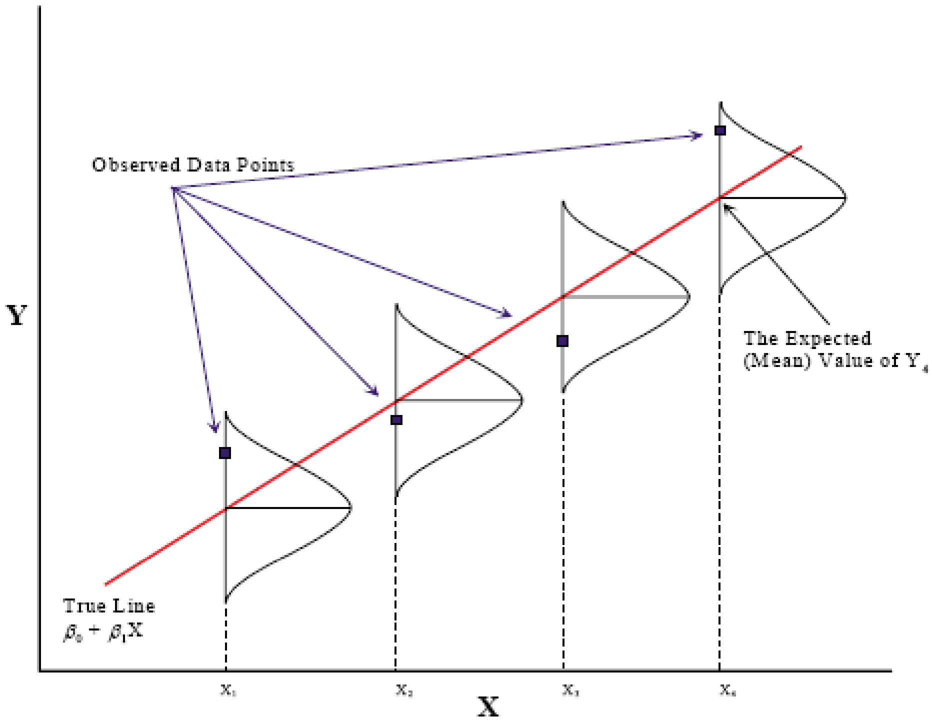

As "dependent" or "controlling" variable we chose the Crude Oil Purchase Price (COPP). The independent variables are considered as being "controlled" by the dependent variables. As independent variables, we chose the VLCC, AFRAMAX, HANDYSIZE and SUEZXMAX time charter rates. This will give us an initial image for the relationship between COPP and time charter rates:

```
                        Dependent Variable    VLCC

                    Ordinary Least Squares Estimates

     SSE                 8221397319    DFE                     266
     MSE                   30907509    Root MSE               5559
     SBC                 5391.79013    AIC              5384.60816
     Regress R-Square        0.1205    Total R-Square       0.1205
     Durbin-Watson           0.0459


                                           Standard                 Approx
Variable          DF      Estimate          Error    t Value    Pr > |t|

Intercept          1         20578           1098      18.73      <.0001
COPP               1     -331.4603        54.9083      -6.04      <.0001
                                                                              8.11
```

```
                        Dependent Variable    SUEZMAX

                    Ordinary Least Squares Estimates

     SSE                 1.40796E10    DFE                     266
     MSE                   52930745    Root MSE               7275
     SBC                 5535.97017    AIC              5528.78819
     Regress R-Square        0.0147    Total R-Square       0.0147
     Durbin-Watson           0.0210


                                           Standard                 Approx
Variable          DF      Estimate          Error    t Value    Pr > |t|

Intercept          1         17791           1438      12.38      <.0001
COPP               1     -143.1395        71.8555      -1.99      0.0474
                                                                              8.22
```

```
                        Dependent Variable    AFRAMAX

                    Ordinary Least Squares Estimates

     SSE                 7135870475    DFE                     266
     MSE                   26826581    Root MSE               5179
     SBC                 5353.83974    AIC              5346.65776
     Regress R-Square        0.0707    Total R-Square       0.0707
     Durbin-Watson           0.0236


                                           Standard                 Approx
Variable          DF      Estimate          Error    t Value    Pr > |t|

Intercept          1         17537           1023      17.14      <.0001
COPP               1     -230.1675        51.1551      -4.50      <.0001
                                                                              8.33
```

Dependent Variable HANDYSIZE

Ordinary Least Squares Estimates

SSE	2426109770	DFE	266
MSE	9120713	Root MSE	3020
SBC	5064.70923	AIC	5057.52725
Regress R-Square	0.0237	**Total R-Square**	**0.0237**
Durbin-Watson	0.0189		

Variable	DF	Estimate	Standard Error	t Value	Approx Pr > |t|
Intercept	1	11854	596.7228	**19.86**	**<.0001**
COPP	1	-75.7741	29.8277	**-2.54**	**0.0116**

8.34

The R-squared statistic of each model is shown in bold. It measures how close the variable relationship by expressing the percentage of the variation of the dependent variable that is explained by the variation of the independent variable(s). R-squared is a number between 0 and 1. The R-Square statistic is $R^2 = 1 - \frac{SSE}{SST}$ where *SST* is the sum of squares for the original response variable corrected for the mean and SSE is the final error sum of squares. This means that a low R^2 would indicated a large SSE which in turn indicates poor model fit.

In bold and underlined are the significance tests of the dependent variables which are also important in creating an adequate model. The Pr > |t| column is the probability that an observation from a Student's t distribution with degrees of freedom DF is greater in absolute value than the absolute value of the observed statistic *t*. In simple English, the Pr > |t| can tell is if the explanatory variable contributes to the model. The best case scenario is a Pr > |t| of <.0001 which indicates that the variable is statistically significant in the model. It is generally acceptable that a Pr > |t| statistic be significant up to the 5% level (Pr > |t| <= 0.05).

Although the significance tests on all the parameters indicate that they are significant in their importance in an OLS model, we see that in all three cases, the R-square statistic is very low. This negates the statistical significance of each parameter since the "significant" parameters explain only miniscule variations (very small R-square values). This is a common error in many regression models where only the statistical significance of the dependent variables is taken into consideration. This error can be shown graphically by taking a theoretical 100 data points and fitting a statistically significant correlation that explains a tiny 4% of the variation in y (R-square = 0.2).

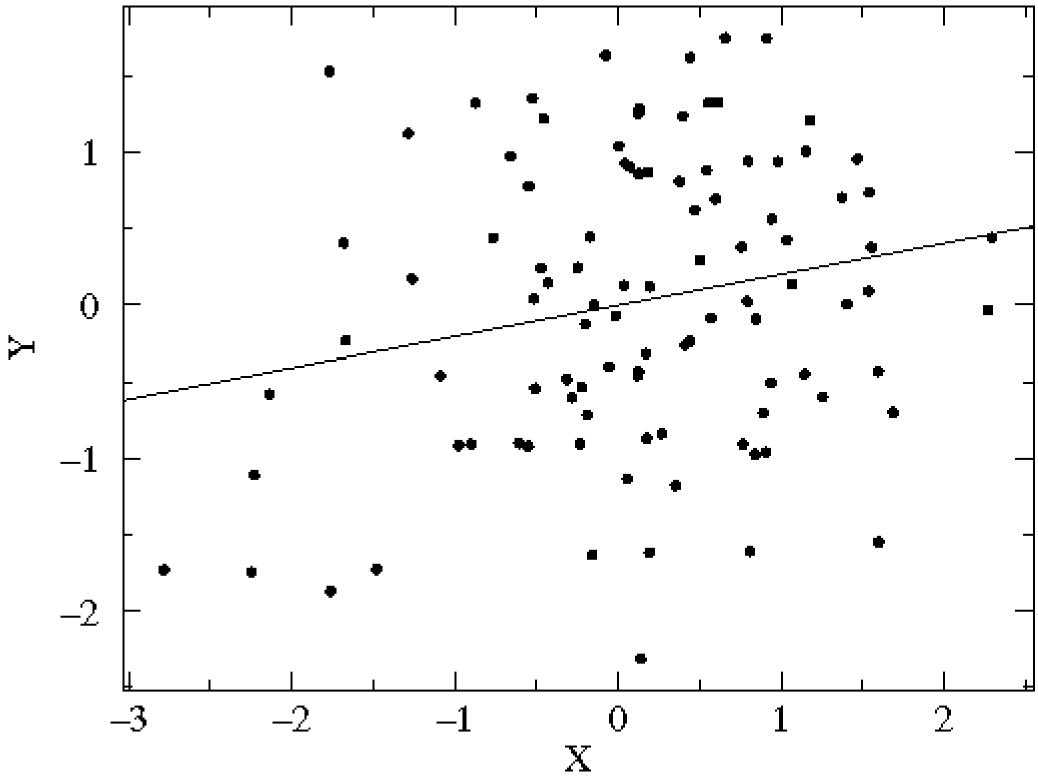

While the correlation may be statistically significant for X, it is undoubtedly not significant in explaining the variation of Y's value. Concluding, a strong statistical correlation by itself is not important unless it is also followed by a strong R-square statistic.

Other given statistics of the OLS Regression output are:

- **Variable** - This column shows the predictor variables. The first variable represents the constant or the *Y intercept*. This is the height of the regression line when it crosses the Y axis. In other words, this is the predicted value of each time charter rate when all other variables are 0.
- **DF** - This column give the degrees of freedom associated with each independent variable. All continuous variables have one degree of freedom. The total variance has N-1 degrees of freedom where N is the number of observations. The model degrees of freedom

corresponds to the number of predictors (plus the intercept) minus 1 (K-1). The Residual degrees of freedom is the total DF minus the model DF (N - 1) – (K – 1) .

- **Parameter Estimates** - These are the values of the regression equation that are used for predicting the dependent variable from the independent variable.
- **Standard Error** - These are the standard errors associated with the coefficients. The standard error is used for testing whether the parameter is significantly different from 0 by dividing the parameter estimate by the standard error to obtain a t-value (see the column with t-values and p-values). The standard errors can also be used to form a confidence interval for the parameter, as shown in the last two columns of this table.

Having analyzed each time charter rate's association with the Crude Oil Purchase Price and found it lacking in model fit, we can perform an analysis on the relationship between time charter rates.

Dependent Variable VLCC

Ordinary Least Squares Estimates

SSE	1535465473	DFE	266
MSE	5772427	Root MSE	2403
SBC	4942.11116	AIC	4934.92919
Regress R-Square	0.8357	**Total R-Square**	**0.8357**
Durbin-Watson	0.2196		

Variable	DF	Estimate	Standard Error	t Value	Approx Pr > |t|
Intercept	1	999.8356	389.4707	**2.57**	**0.0108**
AFRAMAX	1	1.0086	0.0274	**36.79**	**<.0001**

8.5

Dependent Variable VLCC

Ordinary Least Squares Estimates

SSE	5061742851	DFE	266
MSE	19029108	Root MSE	4362
SBC	5261.80227	AIC	5254.6203
Regress R-Square	0.4585	**Total R-Square**	**0.4585**
Durbin-Watson	0.0683		

Variable	DF	Estimate	Standard Error	t Value	Approx Pr > |t|
Intercept	1	597.3929	949.3091	**0.63**	**0.5297**
HANDYSIZE	1	1.3133	0.0875	**15.01**	**<.0001**

8.6

```
                         Dependent Variable    VLCC

                      Ordinary Least Squares Estimates

     SSE                 1713427616    DFE                      266
     MSE                    6441457    Root MSE                2538
     SBC                 4971.50064    AIC               4964.31867
     Regress R-Square        0.8167    Total R-Square        0.8167
     Durbin-Watson           0.1487

                                         Standard                  Approx
Variable           DF     Estimate          Error    t Value    Pr > |t|

Intercept           1         3258       355.4897       9.17      <.0001
SUEZMAX             1       0.7309         0.0212      34.43      <.0001

                                                                             8.77
```

We can see that the R-squared and t-value for an ordinary least squares regression is quite high, indicating good model fit with all the time charter rates (SUEZMAX, AFRAMAX and HANDYSIZE) significant in explaining the VLCC time charter rate. We can also come to the conclusion that the time charter rates of the smaller class vessels are better suited to explaining a VLCC time charter rate, than the Crude Oil Purchase Price.

2.2 The Cross Correlation Function

Another question that is also of great importance is whether two series influence each other in future points in time. The cross correlation function measures the degree of associaton between an explanatory variable and the dependent variable at various time lags. Because the cross correlation function is not symmetric, cross correlations are calculated for both positive lags and negative lags (leads). The large lag cross correlations are an indication that current y_t is related to past values of the explanatory variable. Large lead cross correlations are an indication that y_t is a predictor of x_t. As an example, let's take the cross correlation function of the VLCC time charter rate with itself, albeit shifted 5 months backwards in time. This would mean that the shifted VLCC series (explanatory variable) will "influence" the normal VLCC series (dependent variable) in t = +5 periods. From the output below, the cross correlation function shows the obvious relationship between the dummy series and the actual value.

Since we are dealing with returns, this cross correlation function shows that a shift in the dummy series will cause an analogous shift in the actual series in t = +5 periods. The correlation is *positive* indicating that the series moves in tandem (which is expected since it is the same series).

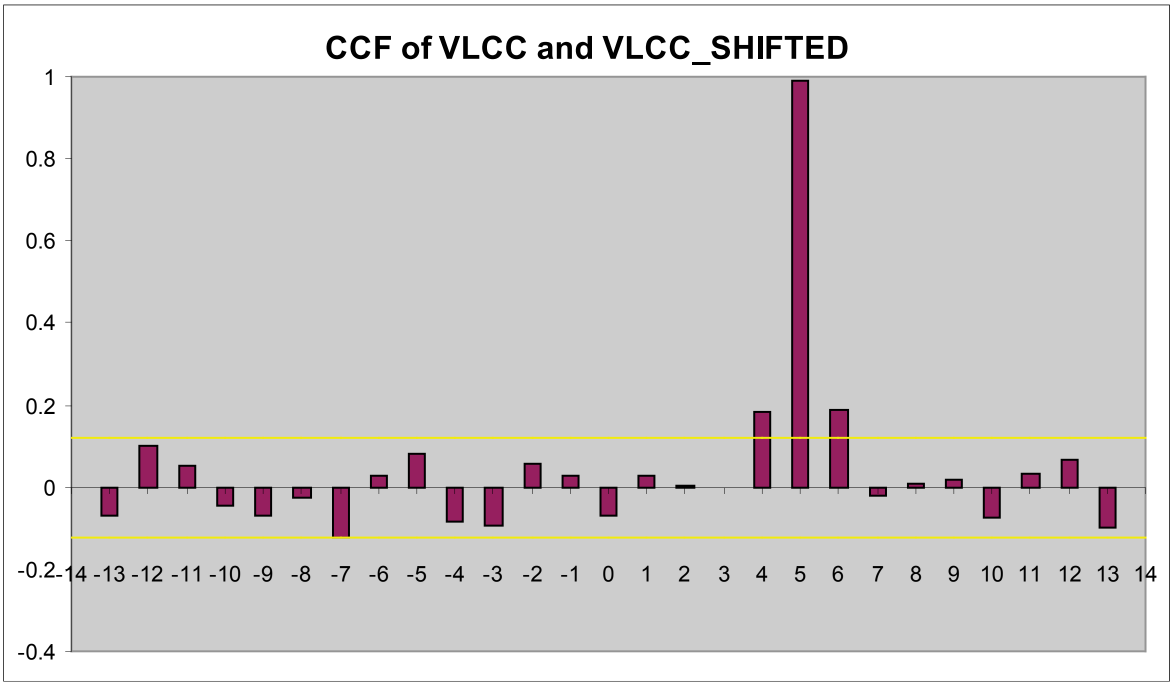

Figure 13 : CCF of VLCC and VLCC_SHIFTED

We can use the cross correlation function to look at all the explanatory variables, checking for any relationships between them and the VLCC time charter rate. While we don't expect a high order of correlation, we can check whether any of the time charter rates (AFRAMAX, SUEZMAX, HANDYSIZE) affect the VLCC time charter rate in the future.

In the first cross correlation we compare the 1^{st} differenced time charter rate at time t, with 1^{st} differenced oil data at a time t-i using the cross correlation command. This would show whether an any noticeable spike in crude oil prices at time *t* has any impact on future time charter rates. Note that from now on only differenced series will be used because as it has been described, model data sets must follow certain principles (stationarity and independence of the data) which are not met by ordinary raw data. A cross correlation function of time charter rates and crude oil purchase price follows.

2.2.1 VLCC

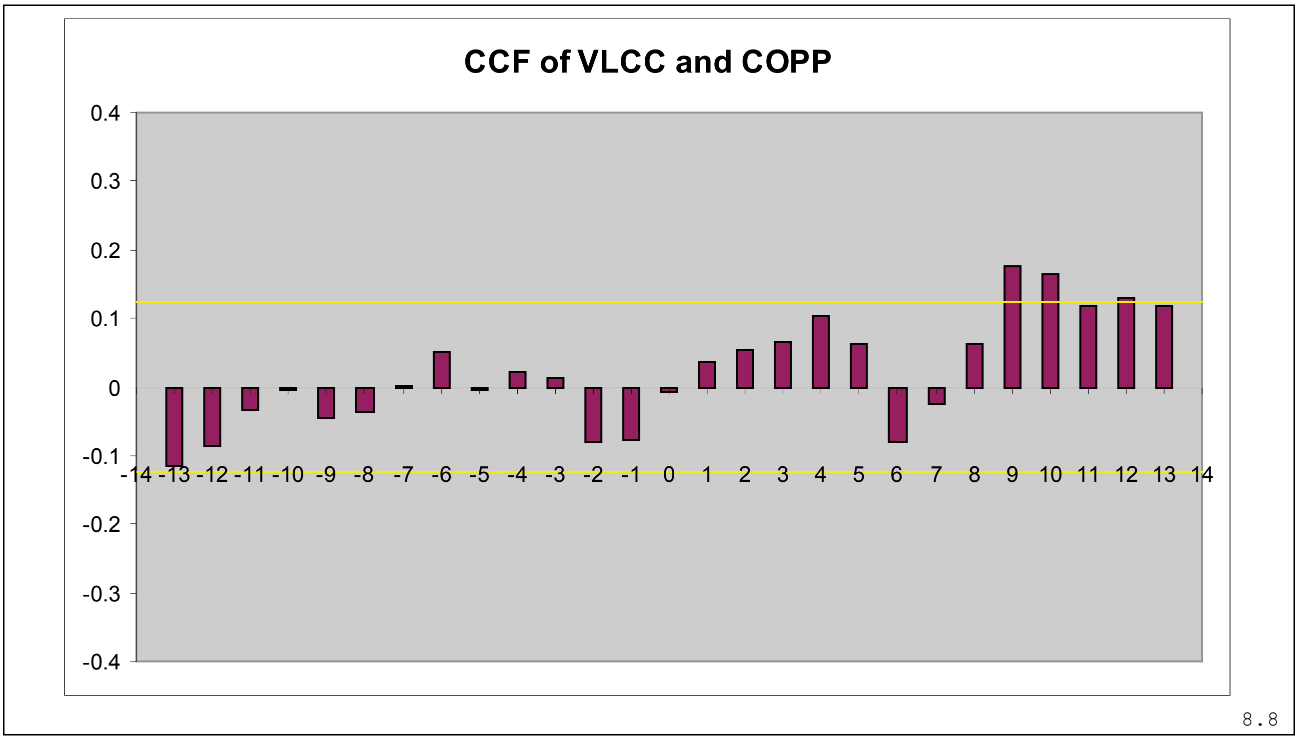

Figure 14 : CCF of VLCC and Crude Oil Purchase Price

We can see that the crude oil purchase price (COPP) has the biggest effect on the VLCC time charter rate (VLCC TCRs) at around t = +9 and +10. This is expected since the volatility of the VLCC TCRs is due to its long voyages and long reaction times. In this case, the value is positively correlated meaning that a rise in the COPP at month t will mean a rise in time charter rates beginning at the 9th and 10th months after the initial price change, onwards.

2.2.2 AFRAMAX

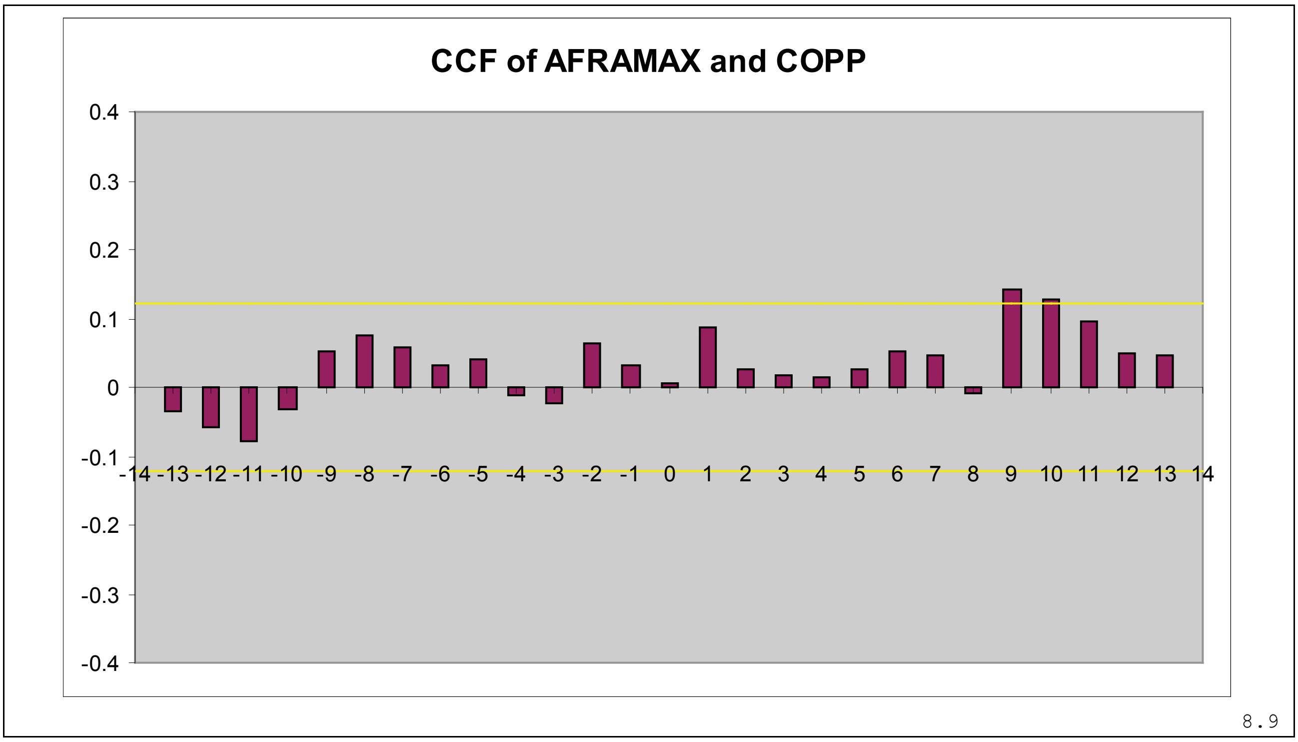

Figure 15 : CCF of AFRAMAX and Crude Oil Purchase Price

As can be shown from the cross correlation diagram between the AFRAMAX time charter rate (AFRAMAX TCR) and the Crude Oil Purchase Price (COPP), the COPP has the biggest effect on the AFRAMAX time charter rate (AFRAMAX TCRs) at around t = +9 and +10. It is more or less the same characteristic found in the VLCC time charter rate albeit less pronounced. This indicates that both AFRAMAX and VLCC time charter rates are influenced by crude oil prices within the same time frame.

2.2.3 HANDYSIZE

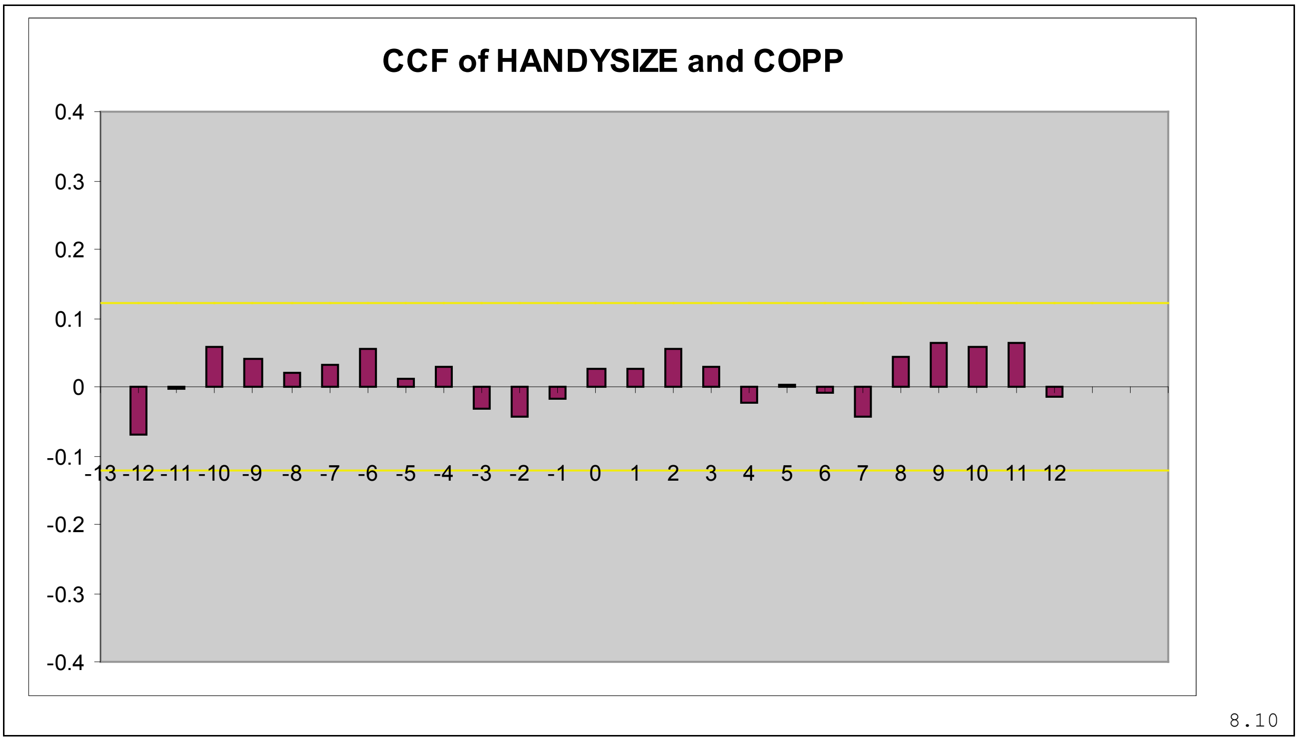

Figure 16 : CCF of HANDYSIZE and Crude Oil Purchase Price

As can be shown from the cross correlation diagram between HANDYSIZE time charter rate (HANDYSIZE TCR) and the Crude Oil Purchase Price (COPP), there isn't any significant correlation (positive or negative) between the data. This indicates that a shift in the COPP now has doesn't have an effect on the HANDYSIZE TCR in the future or vice versa. This is expected since the HANDYSIZE TCR much less volatility because it carries much less cargo for shorter voyages and can adapt more smoothly to sudden changes in COPP.

2.2.4 SUEZMAX

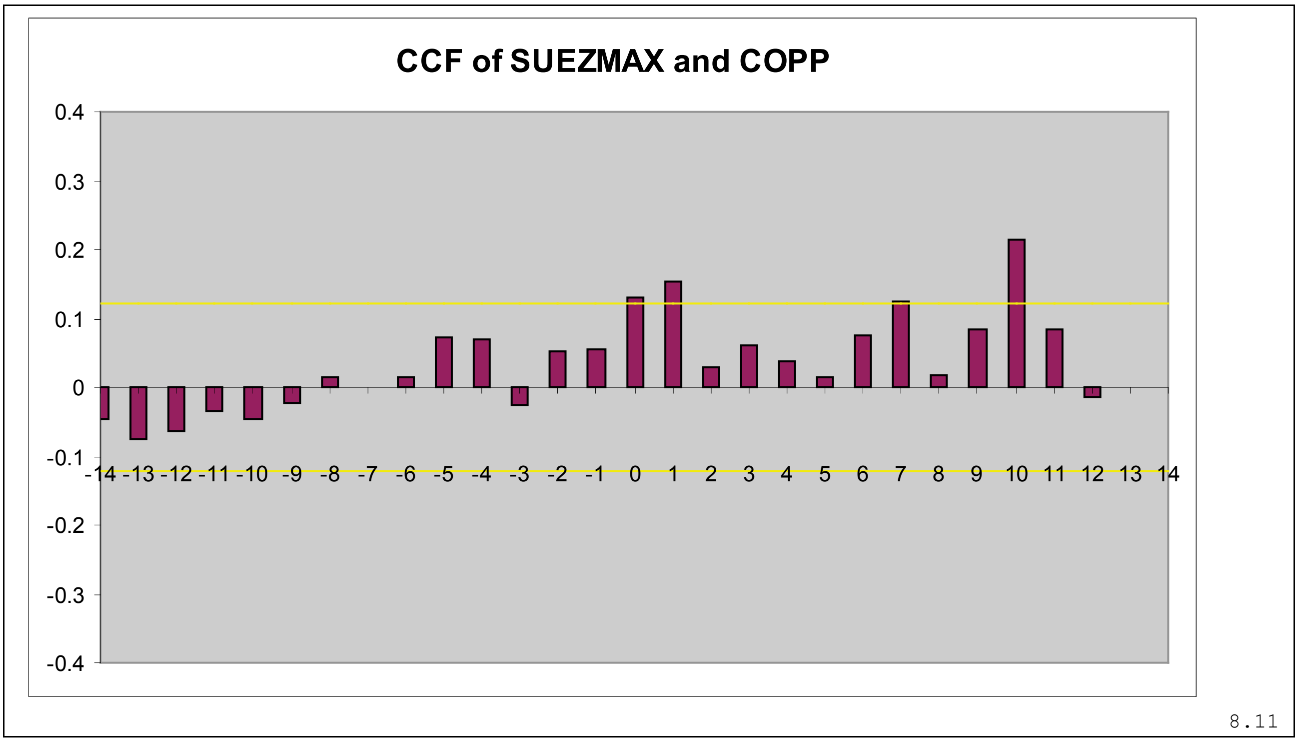

Figure 17 : CCF of Suezmax and Crude Oil Purchase Price

We can see that the HANDYSIZE TCRs have different relationship with COPP than AFRAMAX and VLCC TCRs. The statistically significant lags is are at t= +1, t=+7 and a significant t=+10 which indicates that SUEZMAX TCRs are influenced by COPP in both the short and mid and long term. As with the previous two TCRs, there seems to be a statistically significant lag at t = +10 which in this case is quite significant. It is the author's advice that any future thesis considering to model TCRs can use the COPP lagged 9 or 10 periods.

2.3 Alternate Explanatory Variables

Other than the COPP, we can also check for explanatory variables in the VLCC category that may influence the VLCC time charter rate. These include:

- VLCC Scrap Prices
- 5 Yr VLCC Second-Hand Prices
- Crude Energy Materials
- Newbuilding Prices

The returns are used for these series cross correlation function for each of these explanatory variables with the VLCC charter rate is shown below:

2.3.1 VLCC Scrap Prices

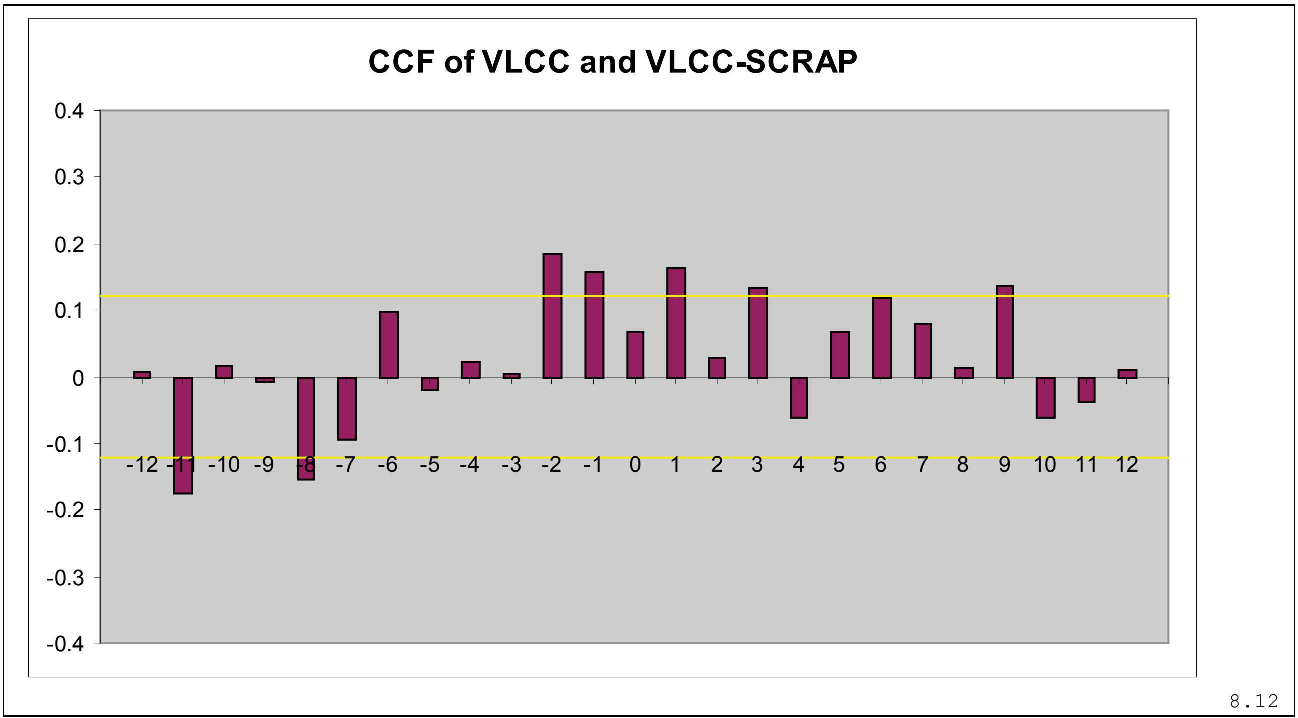

Figure 18 : CCF of VLCC and VLCC Scrap Prices

As we can see from the diagram, there is a positive correlation between the VLCC time charter rate returns and VLCC Scrap Price returns through the whole range of the diagram. Unlike the other time charter rates or the Crude Oil Purchase price, we can say that the VLCC Time Charter rate and the explanatory variable may influence each other throughout the year. This can be seen in the diagram where there is no single conclusive point in time where a change in scrap prices will effect the VLCC time charter rate or vice versa. As before, we can see that Scrap Prices are positively correlated with the VLCC Time Charter Rate at t=+1,t=+3, t=+6 and t=+9. This means that the Scrap Prices have an influence on the VLCC TCR at t=+1, t=+3 , t=+6 and t=+9 months in the future. We can also see the effect that the VLCC TCR has on Scrap Prices (the negative lags). At t= -11 and t= -8 months in the past, the VLCC TCR influences Scrap Prices negatively (Drop in the VLCC Time Charter Rate = Rise In Scrap Prices) while the short term past influence from 2 to 1 months before indicates a positive influence on the VLCC time charter rate (Drop in the Scrap Prices = Drop in Charter Rate).

2.3.2 5 Yr VLCC Second-Hand Prices

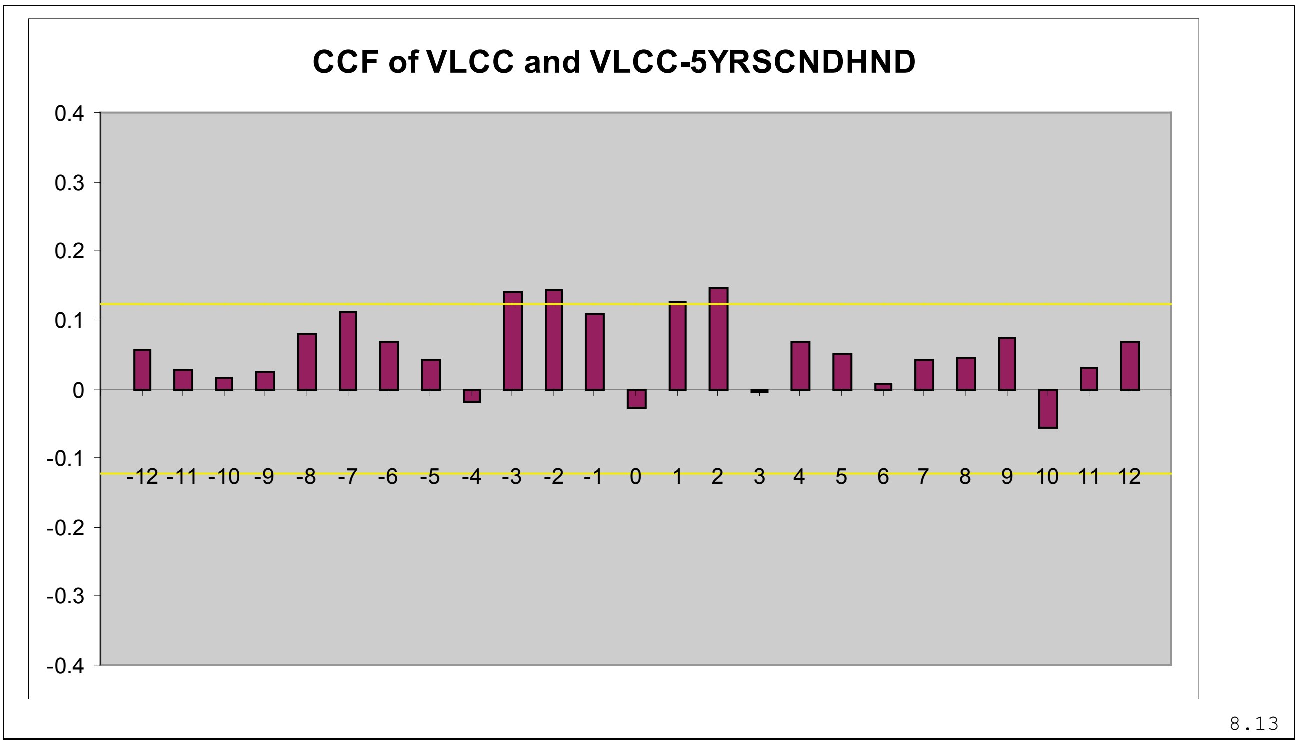

Figure 19 : CCF of VLCC and 5yr Second Hand Prices

There is a short term increase as t approaches 0 from both sides of the correlation function meaning that VLCC prices and 5 Yr Second Hand VLCC prices influence each other only in the short term although for time t = 0 the relationship seems to disappear hinting that the relationship between the rates is not instantaneous. Five Year Second Hand Prices seem to slightly influence the VLCC time charter rate at t=+1 and t=+2, that is, one and two months in the future while 5 Year Second Hand Prices influence the VLCC Time Charter Rate at t = -2 and t= -3.

2.3.3 Crude Energy Materials

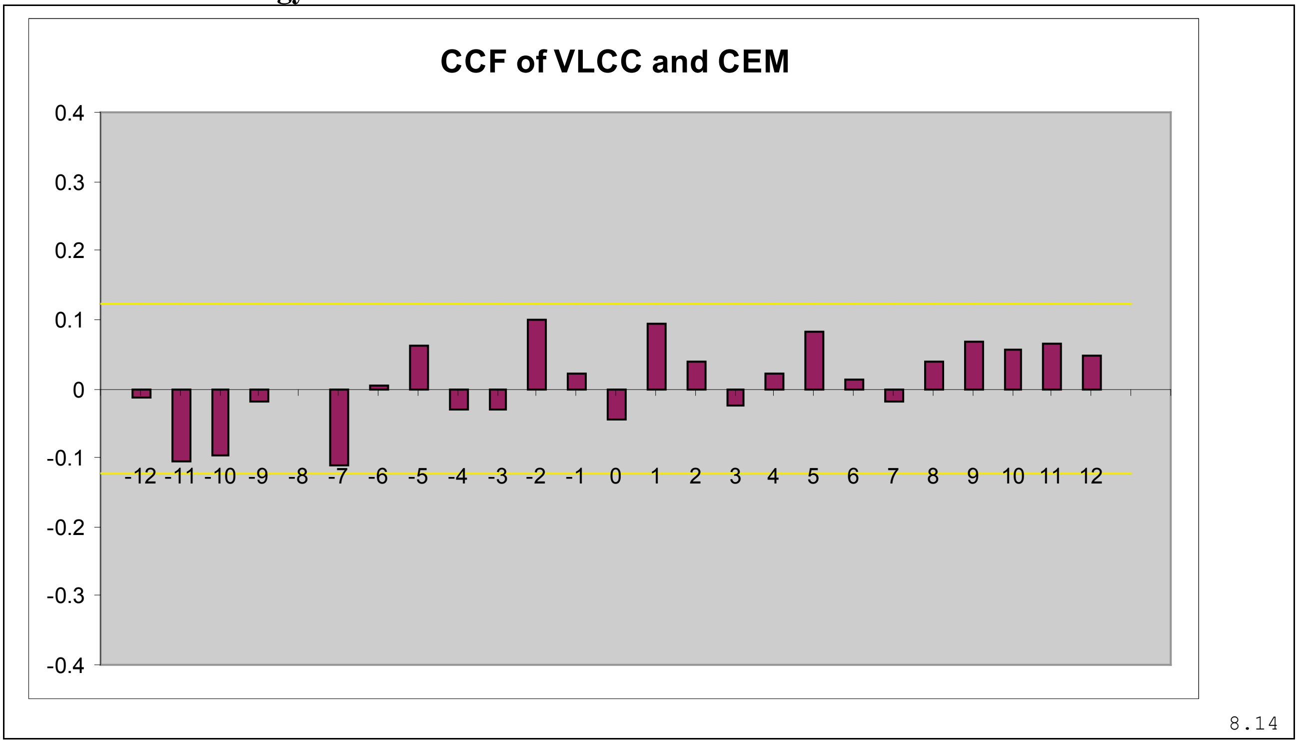

Figure 20 : CCF of VLCC and Crude Energy Materials

VLCC tankers carry oil and the US Department of Labor's Producer Price Index for Crude Energy Materials shows that there is absolutely no major relationship between a shift in the index and a shift in the VLCC time charter rate. The spikes at lags t=-11, t=-10 and t=-7 indicate that VLCC Time Charter rates might have a negatively correlated influence on the CEM Price Index 7,10 and 11 months later. Unfortunately the lag amounts are within the standard error and can only be considered borderline significant.

2.3.4 Newbuilding Prices

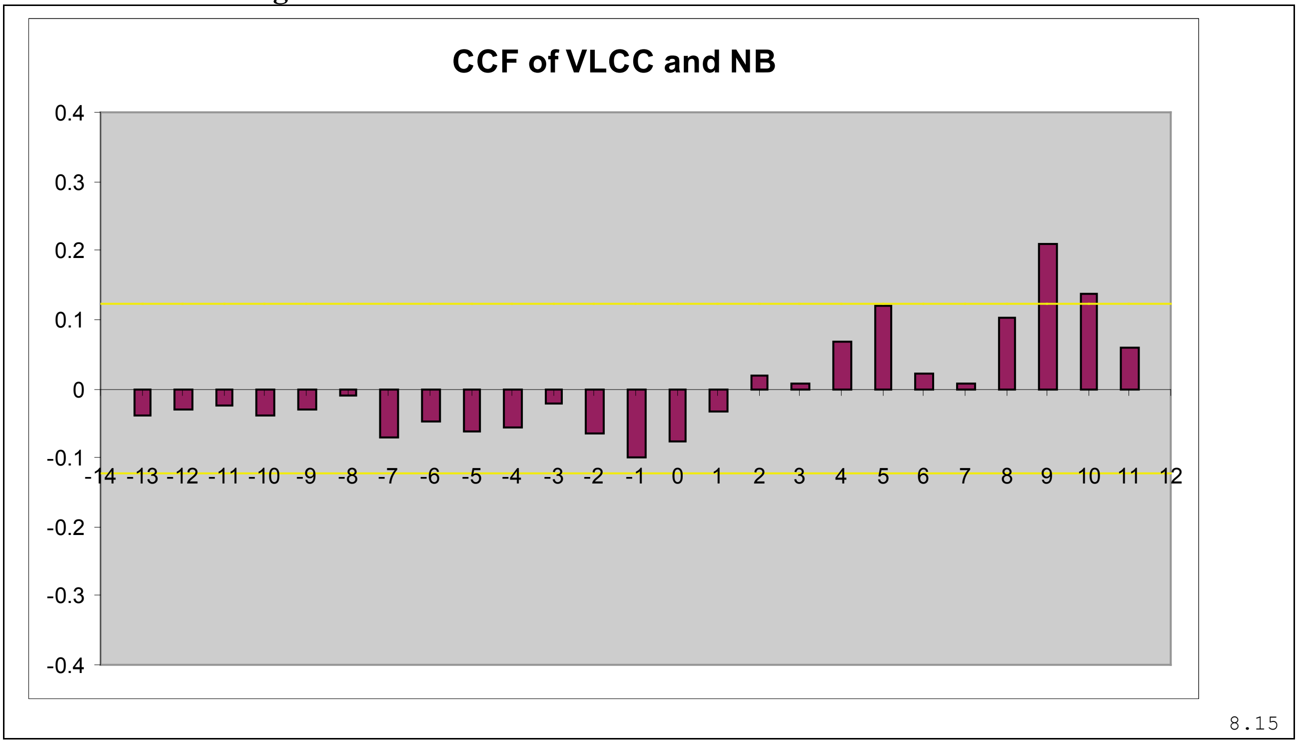

Figure 21 : CCF of VLCC and NewBuilding Prices

The spikes exhibited for lags +5, t=+9 and t=+10 in this CCF leads us to conclude that Newbuilding Prices (NB) influence future VLCC TCRs 5,9 and certainly 10 months in the future.

2.3.5 Newbuilding and Arab Group Oil Production

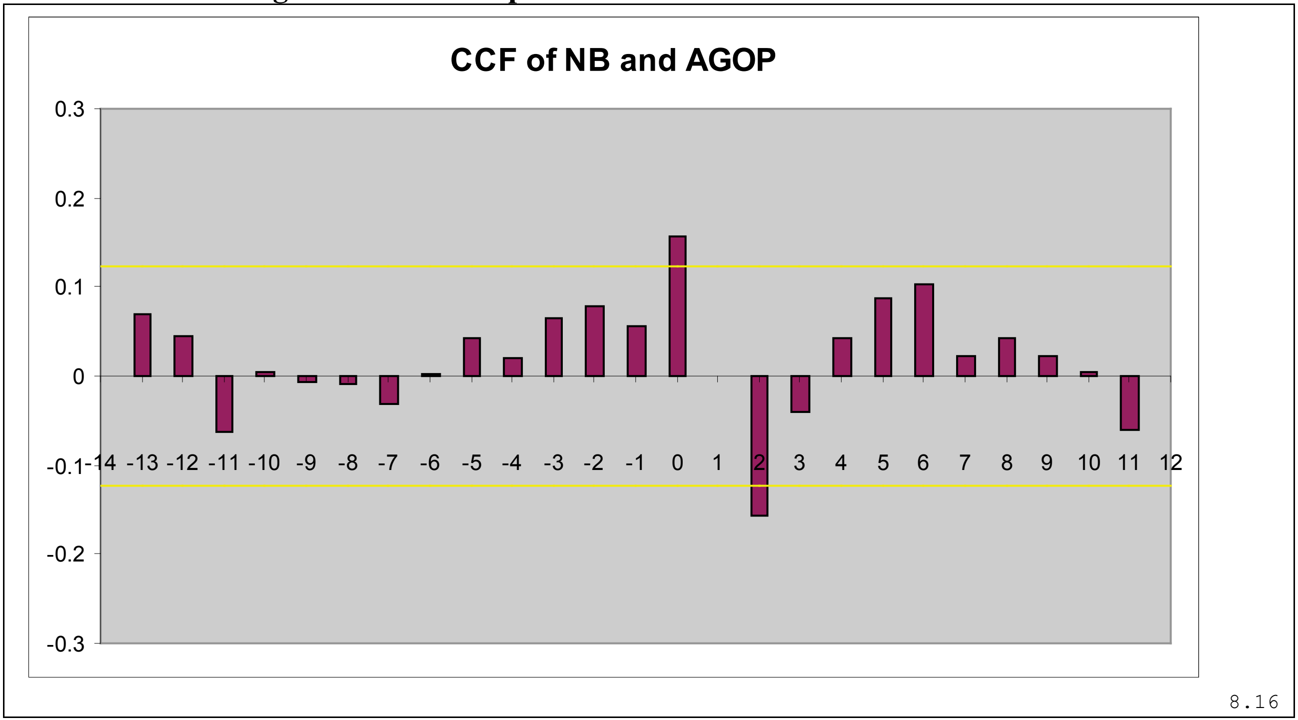

Figure 22 : CCF of Newbuilding Prices and Arab Group Oil Production

As a final CCF between two explanatory variables, the Newbuilding Prices and Arab Group Oil Production shows that Arab Group Oil Production is negatively correlated at the t=+2 lag. This means that Arab Group Oil Production negatively influences Newbuilding Prices two months in the future.

2.4 Time Charter Rate Cross Correlations

Clarkson's Research (2004) kindly provided time series of monthly period charter rates and vessel values for the period from Oct 1979 to May 2004. We explored the following shipping sectors: early 70s 250,000 DWT VLCC tanker, Early 90s 140,000 SH DWT Suezmax tanker, Early 90s 95,000 DWT Aframax tanker, 30,000 DWT Handysize tanker. These ships a broad cross-section of the international tanker market in terms of ship size and ship type.

VLCC	SUEZMAX	AFRAMAX	HANDY
1 Year Tanker Timecharter Rates - 250,000 70s $/Day	**1 Year Tanker Timecharter Rates - 140,000 SH Early 90s $/Day**	**1 Year Tanker Timecharter Rates - 95,000 SH Early 90s $/Day**	**1 Year Tanker Timecharter Rates - 30,000 $/Day**

Running a cross correlation between the time charter rates (not the returns) we can see that all the time charter rates show correlation with each other to a high extent. This means that when rates for one time charter rate are high, they are high for the others as well. What we are interested in though is the *rate* in which this change occurs. By running a cross correlation function on the returns of the rates, we will be able to examine if the time charter rates react the same way to the various, and unknown events that drive them. Naturally, when we convert the time charter rate to returns, the correlation will drop off significantly since not all time charter rates react the same way. The cross correlations of the returns of the VLCC time charter rate with each explanatory time charter rate follow.

2.4.1 VLCC Returns and AFRAMAX Returns

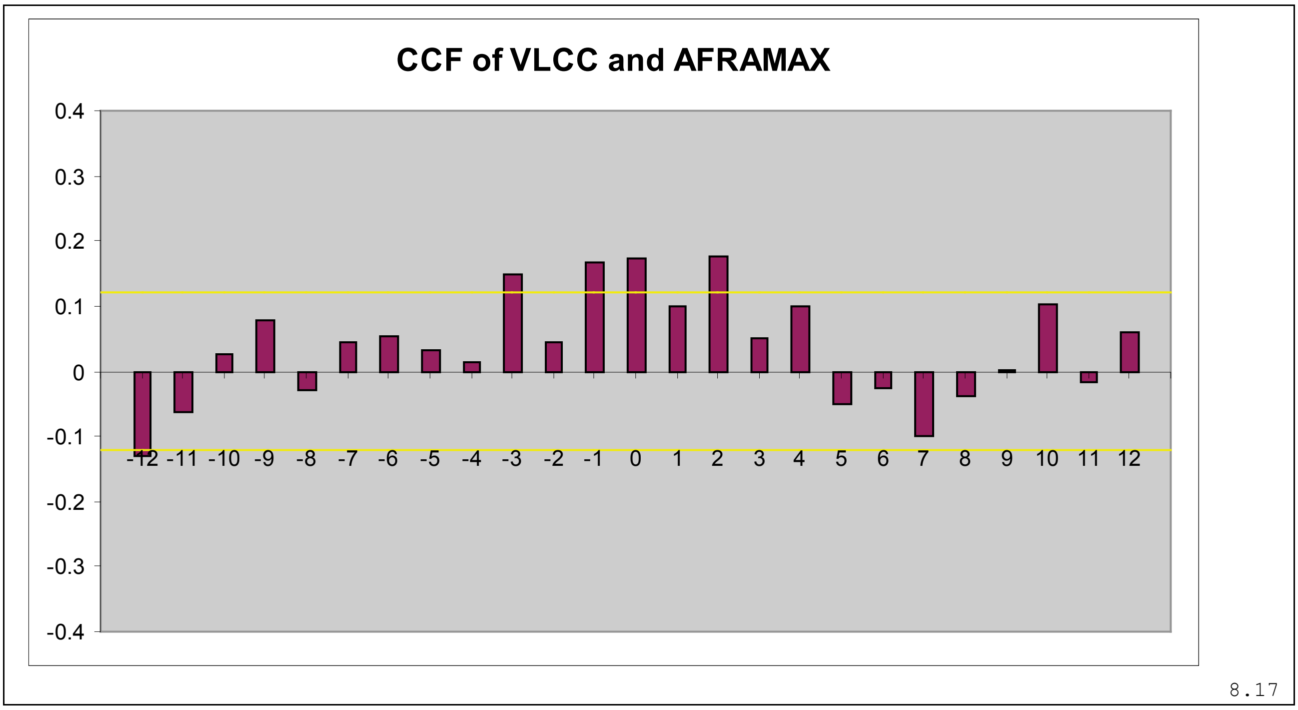

Figure 23 : CCF of VLCC and AFRAMAX

The CCF shows that the AFRAMAX TCR is likely to move in tandem with the VLCC TCR at time t=+0 t=+2.

2.4.2 VLCC Returns and SUEZMAX Returns

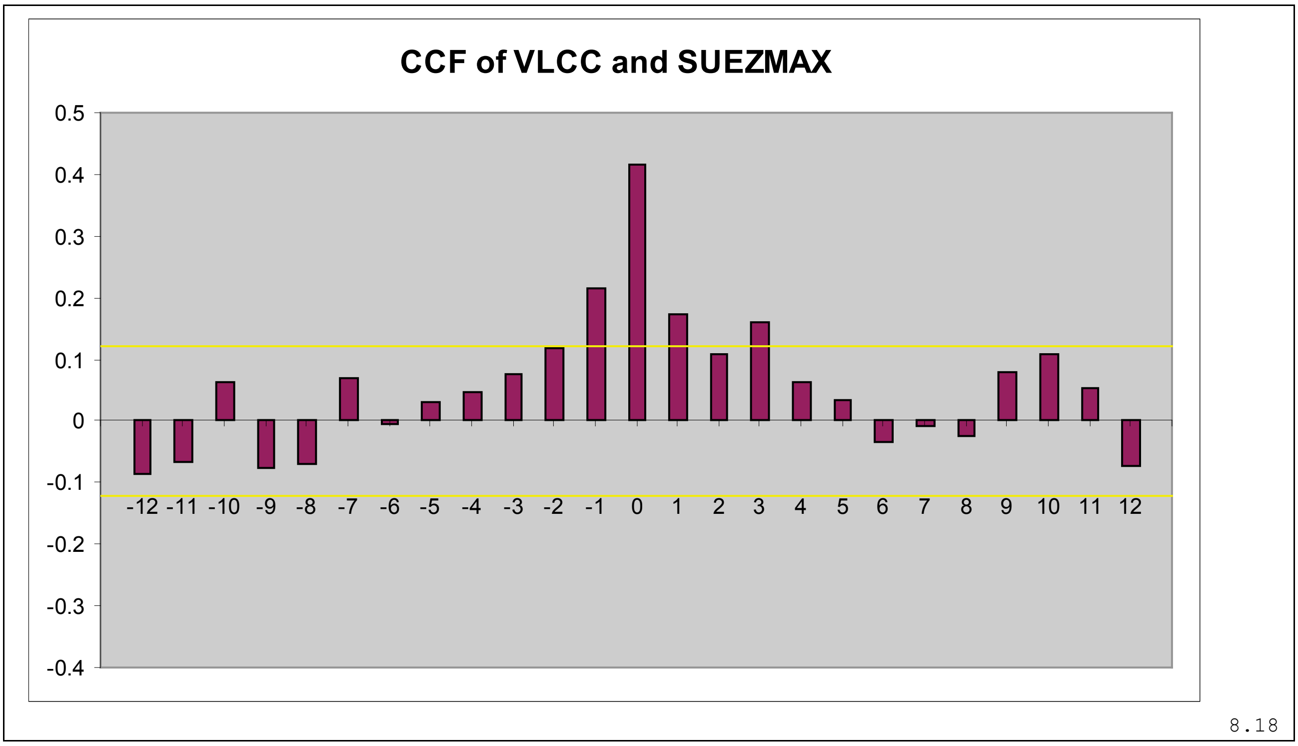

Figure 24 : CCF of VLCC and SUEZMAX

We can see here a direct correlation between the VLCC and SUEZMAX time charter rate returns which shows that they are closely linked in their behavior. Their behavior is one of two time charter rates following each other closely and one can assume that both time charter rates are influenced by similar processes. The lag at t=+1 and t=+3 indicate a tendency for the SUEZMAX to influence the VLCC time charter rate one and three months later. From the CCF it is obvious though that the VLCC TCR and SUEZMAX move in tandem. While a strong case could be made for the VLCC TCR influencing the SUEZMAX TCR at t = -1.

2.4.3 VLCC Returns and HANDYSIZE Returns

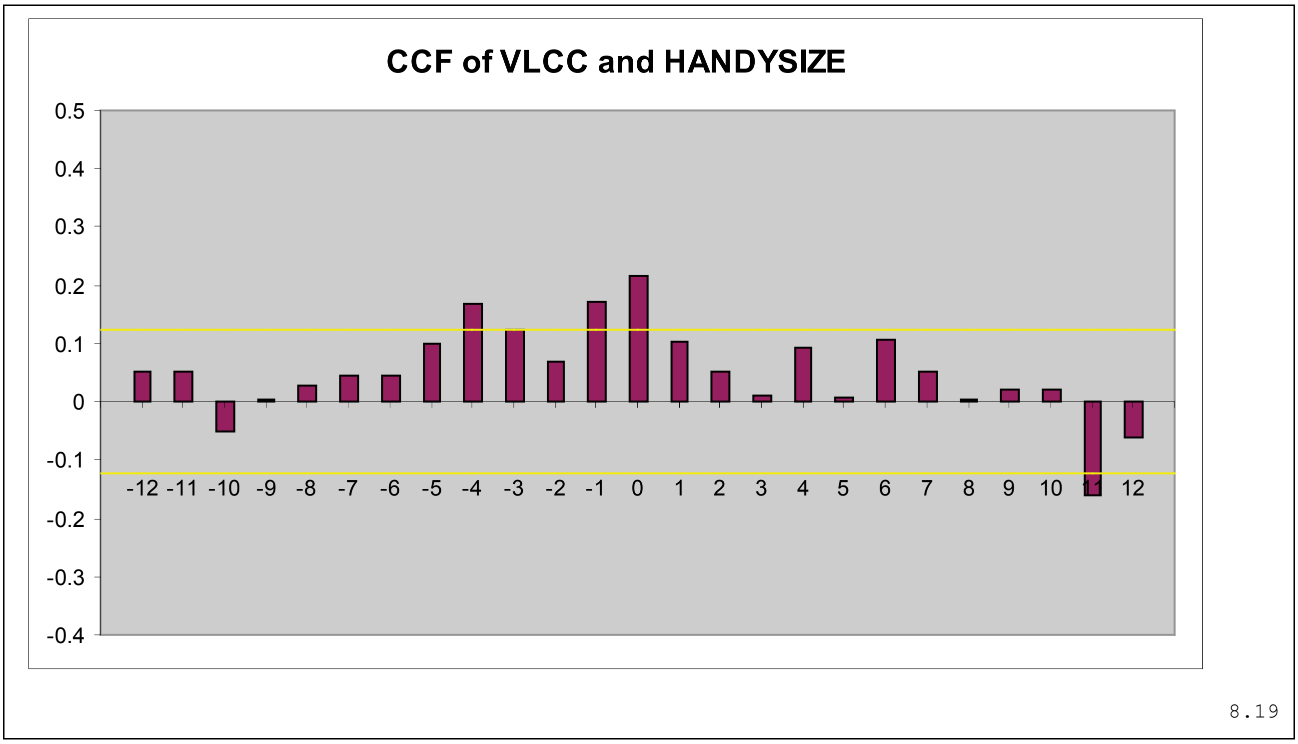

Figure 25 : CCF of VLCC and HANDYSIZE

The HANDYSIZE TCR shows a behavior somewhere between the AFRAMAX and SUEZMAX TCR. It may not by as closely in sync with the SUEZMAX TCR, but there does seem to be an association with the VLCC time charter rate at lags t=0 and t=+11. There is also a case to be made concerning the influence of the HANDYSIZE TCR by the VLCC TCR at t= -1 and t= -4.

3. Empirical Results & Discussion

General

The first step in a time series analysis is to plot the observations over time. Features such as a trend, seasonality or outliers will be immediately visible. In the case of our time charter rates we see that there isn't a visible recurring pattern (seasonality). On the other hand we do not know for certain that there isn't some sort seasonality present in time charter rates. It might turn out that the time charter rates might be a *random walk* – a model in which the changes are brought about by a white noise series. These types of series can swing wildly as a result of the combined effects of shocks that drive the series. To get a better understanding of the time charter rates, we check each series' autocorrelation, partial autocorrelation functions and stationarity. If differencing is needed, returns for time charter rates will be used because returns difference and de-trend the series at the same time. Logarithimic differencing is needed in order to remove a possible trend and stabilize the variance thus making the series stationary.

3.1 Analysis of VLCC Time Charter Rates

Using statement (8.20) we plot the data for the VLCC time charter rate.

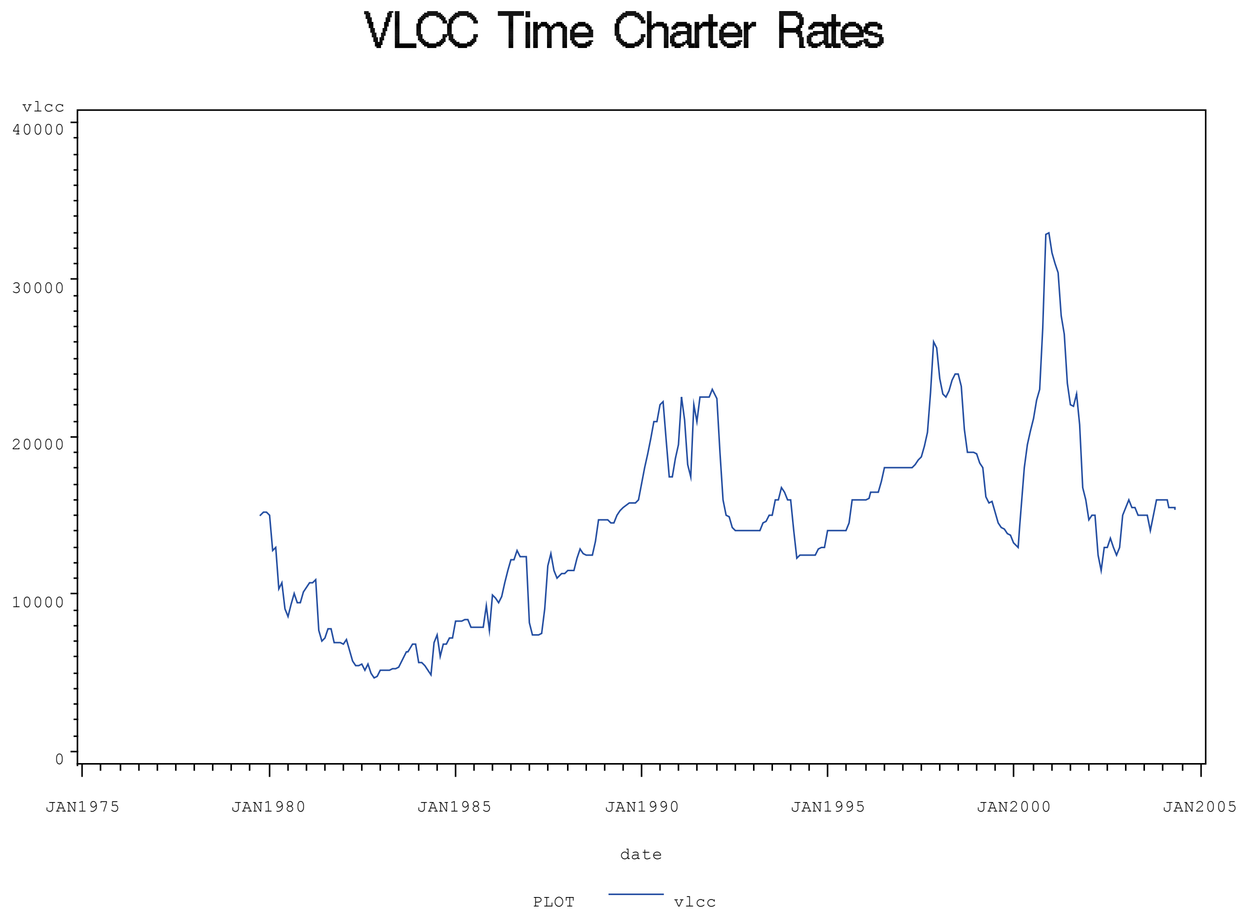

Figure 26 : VLCC Time Charter Rates

First, we use statement (8.21) to print the descriptive statistics for the VLCC Time Charter Rate.

```
Mean of Working Series    14239.01
Standard Deviation        5755.258
Number of Observations         286
```

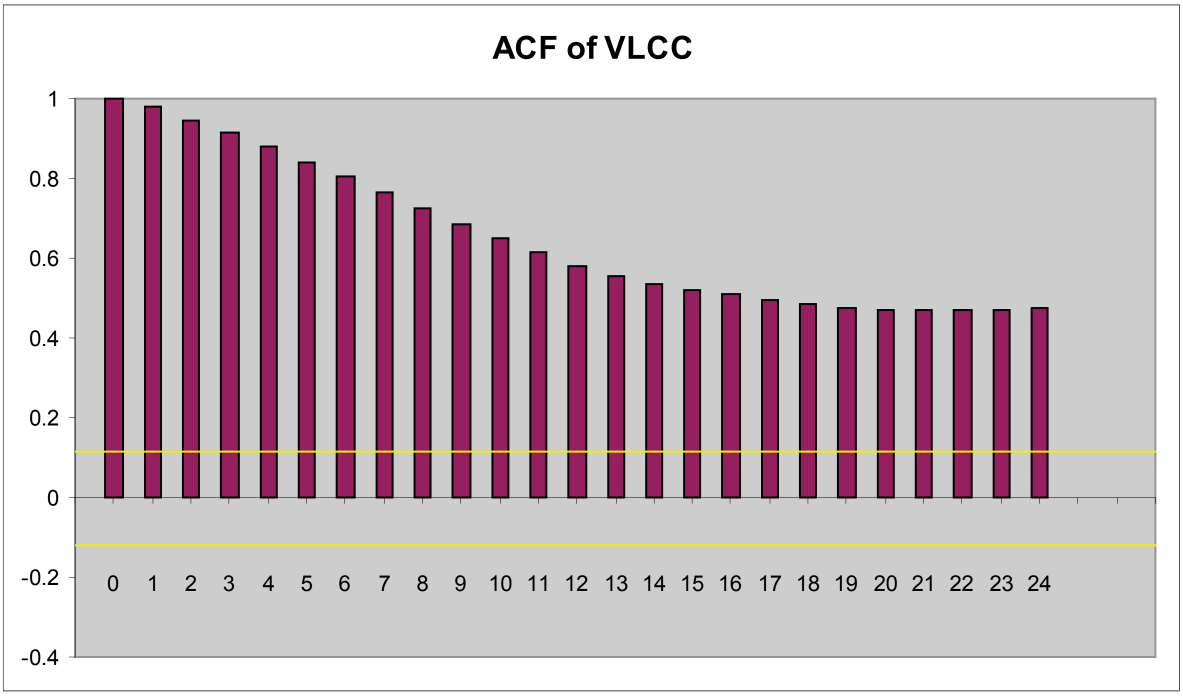

8.21

Figure 27 : ACE of VLCC Time Charter Rate

A distinguishing feature of a *non-stationary* series is its long memory. What the autocorrelation function shows is that the time charter rate at time t ($VLCC_t$) is influenced by shocks (a_t) that occurred in the distant past. The autocorrelation plot shows how values of the series are correlated with past values of the series. For example, the value 0.97987 in the "Correlation" column for the Lag 1 row of the plot means that the correlation between VLCC and the VLCC value for the previous period is .97987. The rows of asterisks show the correlation values graphically. These plots are called autocorrelation functions because they show the degree of correlation with past values of the series as a function of the number of periods in the past (that is, the lag) at which the correlation is computed. We can see that the VLCC time charter rate autocorrelation plot decays slowly. This is typical of a *non-stationary* series. This means that the VLCC time charter rate should be made stationary in order for us to be able to analyze it. After differencing, the autocorrelation plot should quickly fall down to nominal values.

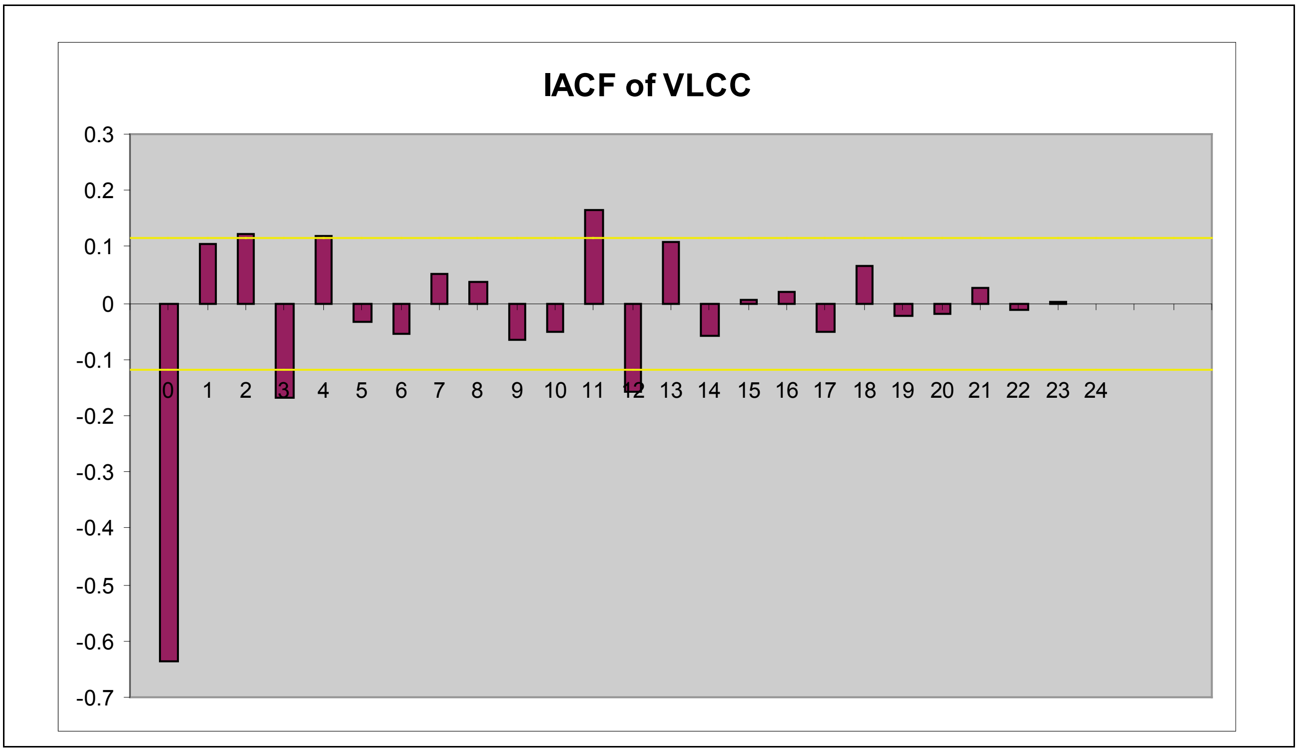

Figure 28 : IACF of VLCC Time Charter Rate

Inverse autocorrelation function (IACF) can detect over-differencing. If the data comes from a nonstationary or nearly nonstationary model, the IACF has the characteristics of a noninvertible moving average. Likewise, if the data come from a model with a noninvertible moving average, then the SIACF has nonstationary characteristics and, therefore, decays slowly. In particular, if the data have been over-differenced, the SIACF looks like a SACF from a nonstationary process. In this case where we only have raw data and no differencing to speak of, we can see that the slow decay of the IACF implies that a moving average exists.

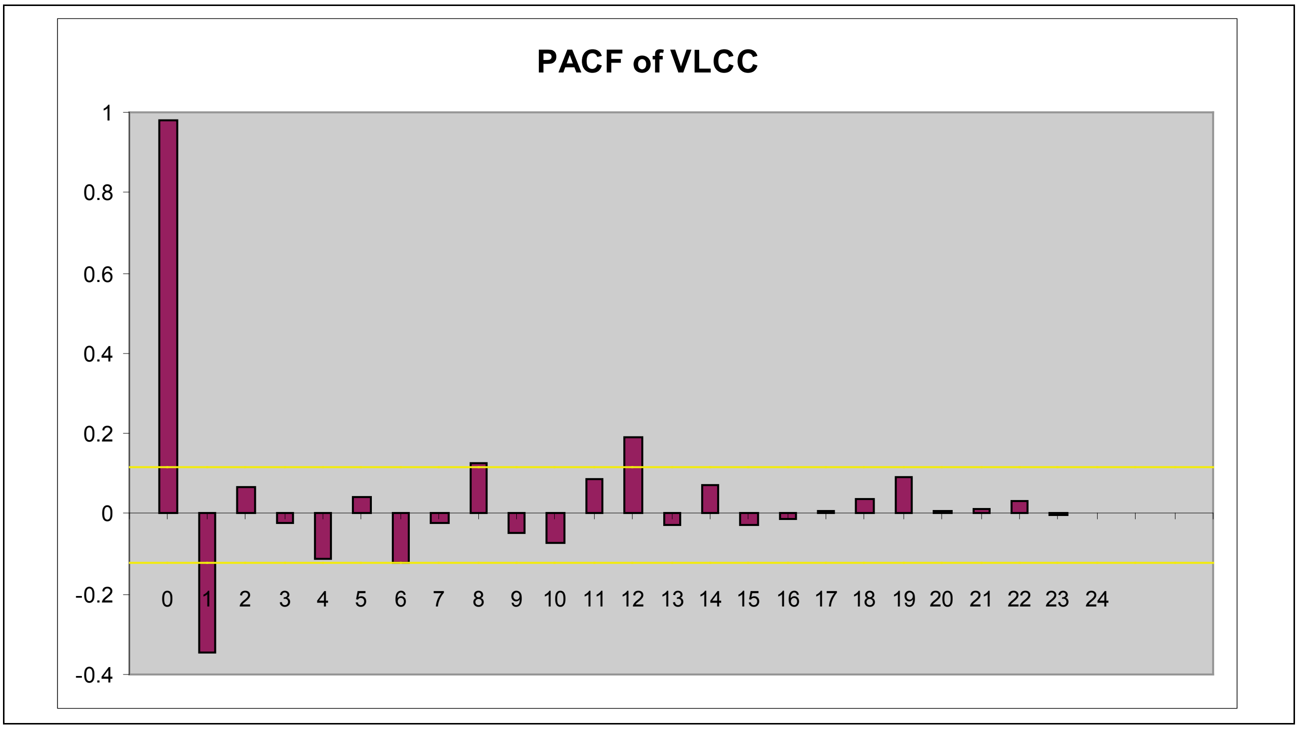

Figure 29 : PACF of VLCC Time Charter Rate

Partial autocorrelations summarize all the information in the autocorrelation function of an autoregressive process in a small number of nonzero statistics (lags). If the time series can be modeled using a simple autoregressive process AR=1 then the partial autocorrelation function of that series will die down immediately after nlag=1. In this case, the VLCC time charter rate's partial autocorrelation function does not die out until 13 meaning that a simple autoregressive p rocess will not be sufficient to model the process and that the series will require multiple autoregressive parameters going back more than one period.

By examining these plots, we can get a first estimate on whether the series is *stationary* or *nonstationary*. In stationary series, the autocorrelation function decays rapidly.

Stationarity

The noise (or residual) series for an ARMA model must be *stationary*, which means that both the expected values of the series and its autocovariance function are independent of time. The standard way to check for nonstationarity is to plot the series and its autocorrelation function. You can visually examine a graph of the series over time to see if it has a visible trend or if its variability changes noticeably over time. If the series is nonstationary, its autocorrelation function will usually decay slowly. Sationarity is important in modeling because if two variables are trending over time, a regression of one on the other could have a high *R*-squared even if the two are totally unrelated. Also the usual "*t*-ratios" will not follow a *t*-distribution, so the hypothesis tests used for the regression parameters will not be valid. Therefore, before we start building a model, we must first test for stationarity of the time series. Most time series are nonstationary and must be transformed to a stationary series before the ARIMA modeling process can proceed. If the series has a trend over time, seasonality, or some other nonstationary pattern, the usual solution is to take the difference of the series from one period to the next and then analyze this differenced series. Sometimes a series may need to be differenced more than once or differenced at lags greater than one period. (If the trend or seasonal effects are very regular, the introduction of explanatory variables may be an appropriate alternative to differencing.) In this case, a visual inspection of the autocorrelation function plot indicates that the VLCC time series is nonstationary, since the ACE decays very slowly. In our case, the time charter rates are converted to returns which is the difference between the log of time charter rate values t and t+1.

$$\mathrm{Re}\,turns = \log(VLCC_{t+1}) - \log VLCC(t)$$

Durbin Watson Test

The Durbin-Watson (DW) is a test for first order autocorrelation. The conditions which must be fulfilled for the DW test to be valid are:

- Constant term in regression
- Regressors are non-stochastic
- No lags of dependent variable

White Noise Test

The white noise test is an approximate statistical test of the hypothesis that none of the autocorrelations of the series up to a given lag are significantly different from 0. If this is true for all lags, then there is no information in the series to model, and no ARIMA model is needed for the series. The autocorrelations are checked in groups of 6, and the number of lags checked depends on the NLAG= option. For the VLCC time series, the white noise test returns:

To Lag	Chi-Square	DF	Pr > ChiSq	--------------------Autocorrelations--------------------					
6	1391.49	6	<.0001	0.980	0.946	0.910	0.873	0.835	0.795
12	2163.22	12	<.0001	0.753	0.709	0.669	0.634	0.598	0.563
18	2613.64	18	<.0001	0.537	0.515	0.500	0.488	0.474	0.462
24	3007.46	24	<.0001	0.455	0.454	0.456	0.459	0.463	0.467

In this case, the white noise hypothesis is rejected very strongly, which is expected since the series is highly correlated (nonstationary) because it has not been converted to returns. The *p* value for the test of the first six autocorrelations is printed as <0.0001, which means the *p* value is less than .0001. In order for the series to become stationary, we will convert it to returns (logged differences). We can check the stationarity with the Dickey Fuller test specified:

Augmented Dickey-Fuller Unit Root Tests

Type	Lags	Rho	Pr < Rho	Tau	Pr < Tau	F	Pr > F
Zero Mean	0	-0.7329	0.5214	-0.58	0.4657		
	1	-1.5669	0.3853	-0.87	0.3393		
	2	-1.4129	0.4071	-0.82	0.3614		
Single Mean	0	-5.6444	0.3742	-1.68	0.4422	1.41	0.7111
	1	-11.6651	0.0876	-2.40	0.1429	2.88	0.3339
	2	-10.9287	0.1051	-2.29	0.1742	2.63	0.3964
Trend	0	-11.4782	0.3371	-2.51	0.3242	3.17	0.5415
	1	-23.6746	0.0295	-3.51	0.0403	6.18	0.0565
	2	-23.3118	0.0320	-3.41	0.0516	5.85	0.0741

The Pr < Rho, Pr < Tau and Pr > F probabilities show that the series is non-stationary (has a unit root) although from the slowly decaying autocorrelation plot it is obvious.

3.2 VLCC Returns

Now we can analyze the return series by using statement (8.22):

```
Name of Variable = vlcc_r

Mean of Working Series     0.000235
Standard Deviation          0.08179
Number of Observations      285
```

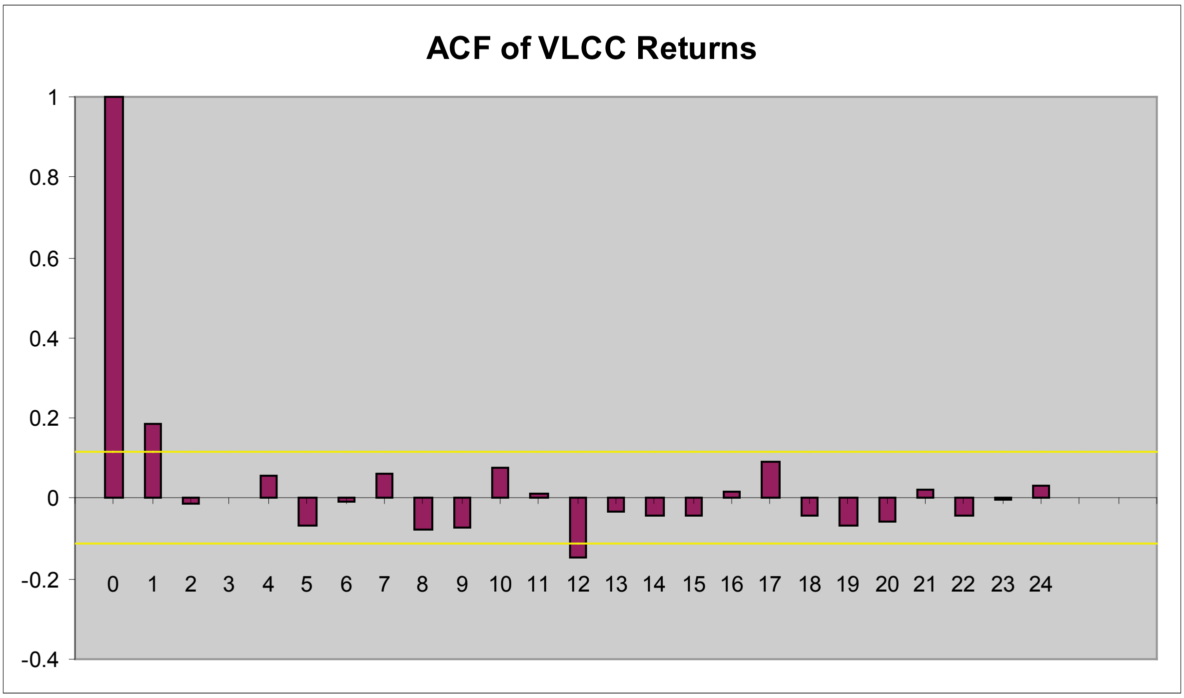

Figure 30 : ACF of VLCC Returns

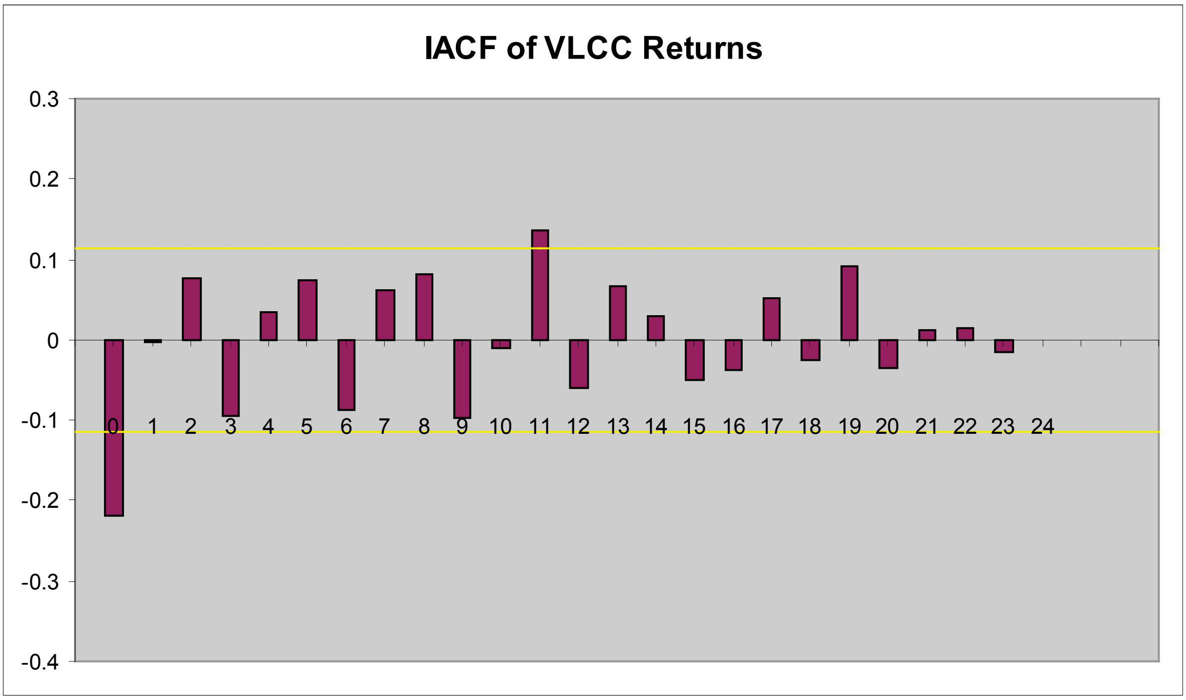

Figure 31 : IACF of VLCC Returns

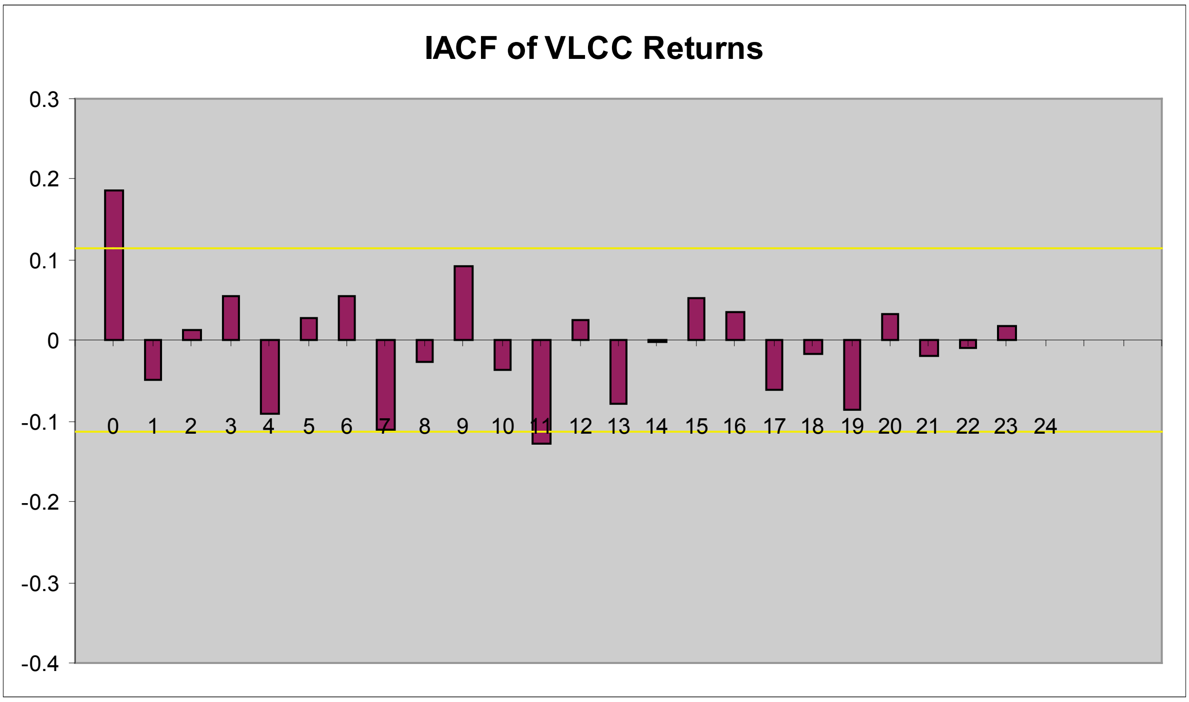

Figure 32 : PACF of VLCC Returns

To Lag	Chi-Square	DF	Pr > ChiSq	--------------------Autocorrelations--------------------					
6	11.41	6	0.0766	0.187	0.019	0.029	0.028	-0.048	-0.014
12	25.17	12	0.0140	0.066	-0.087	-0.104	0.079	-0.002	-0.131
18	31.00	18	0.0288	-0.015	-0.073	-0.047	0.026	0.069	-0.078
24	34.44	24	0.0772	-0.063	-0.065	-0.006	-0.046	0.025	0.013

The Chi-square test statistics for the residuals series indicate whether the residuals are uncorrelated (white noise) or contain additional information that might be utilized by a more complex model. In this case, the test statistics reject the no-autocorrelation hypothesis at a high level of significance. (p=0.0488 for the first six lags). This means that the residuals are not white noise, and so an AR(1) model is not a fully adequate model for this series.

Augmented Dickey-Fuller Unit Root Tests

Type	Lags	Rho	Pr < Rho	Tau	Pr < Tau	F	Pr > F
Zero Mean	0	-230.763	0.0001	-13.92	<.0001		
	1	-237.957	0.0001	-10.87	<.0001		
	2	-218.173	0.0001	-8.82	<.0001		
Single Mean	0	-230.764	0.0001	-13.89	<.0001	96.50	0.0010
	1	-237.960	0.0001	-10.85	<.0001	58.86	0.0010
	2	-218.180	0.0001	-8.81	<.0001	38.78	0.0010
Trend	0	-230.929	0.0001	-13.88	<.0001	96.28	0.0010
	1	-238.371	0.0001	-10.84	<.0001	58.76	0.0010
	2	-218.831	0.0001	-8.80	<.0001	38.72	0.0010

The stationarity test of the return series shows that the hypothesis of the existence of a unit root is rejected and that the series is stationary. The same procedure is followed for the other three time series, AFRAMAX, HANDYSIZE and SUEZMAX.

Alternate Data Transformation

When previous attempts at using the logged differences function to make a series stationary failed, an alternate data transformation was recommended. Let X_t be monthly time charter rate and $E(X_j)$ be the mean time charter rate of the same month for the entire data set. We define a new series Y_t as

$$Y_t = X_t - E(X_j)$$

We further difference the series and define a new series Q_t as

$$Q_t = div(Y_t) = Y_{t+1} - Y_t$$

The autocorrelation function of the new series is

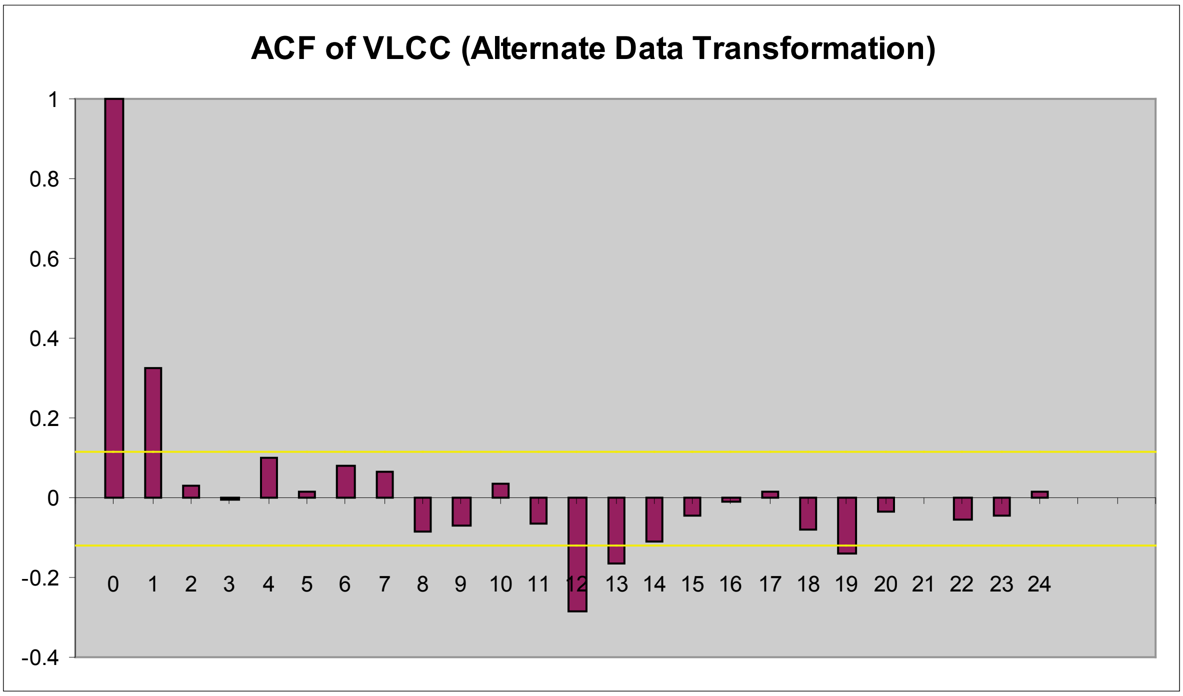

Figure 33 : Autocorrelation Function of VLCC (Alternate Data Transformation)

We can safely assume the original transformation yields a more stationary series for this data.

3.3 Analysis of AFRAMAX Time Charter Rates

Using statement (8.23) we plot the data for the VLCC time charter rate.

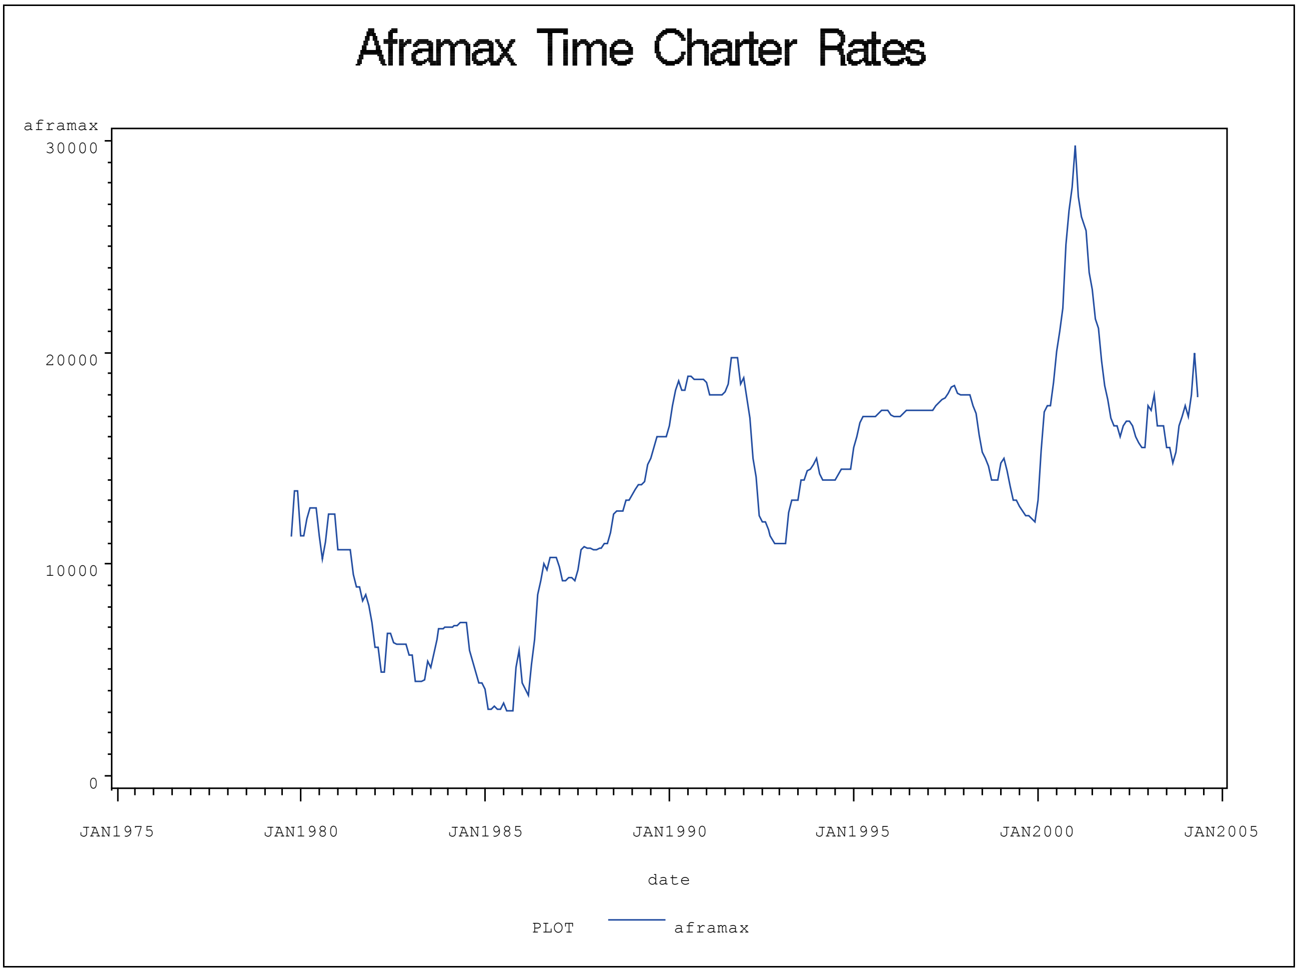

Figure 34 : AFRAMAX Time Charter Rates

We use statement (8.24) to print the descriptive statistics for the AFRAMAX Time Charter Rate.

```
Name of Variable = aframax

Mean of Working Series    13406.33
Standard Deviation        5289.604
Number of Observations         286
```

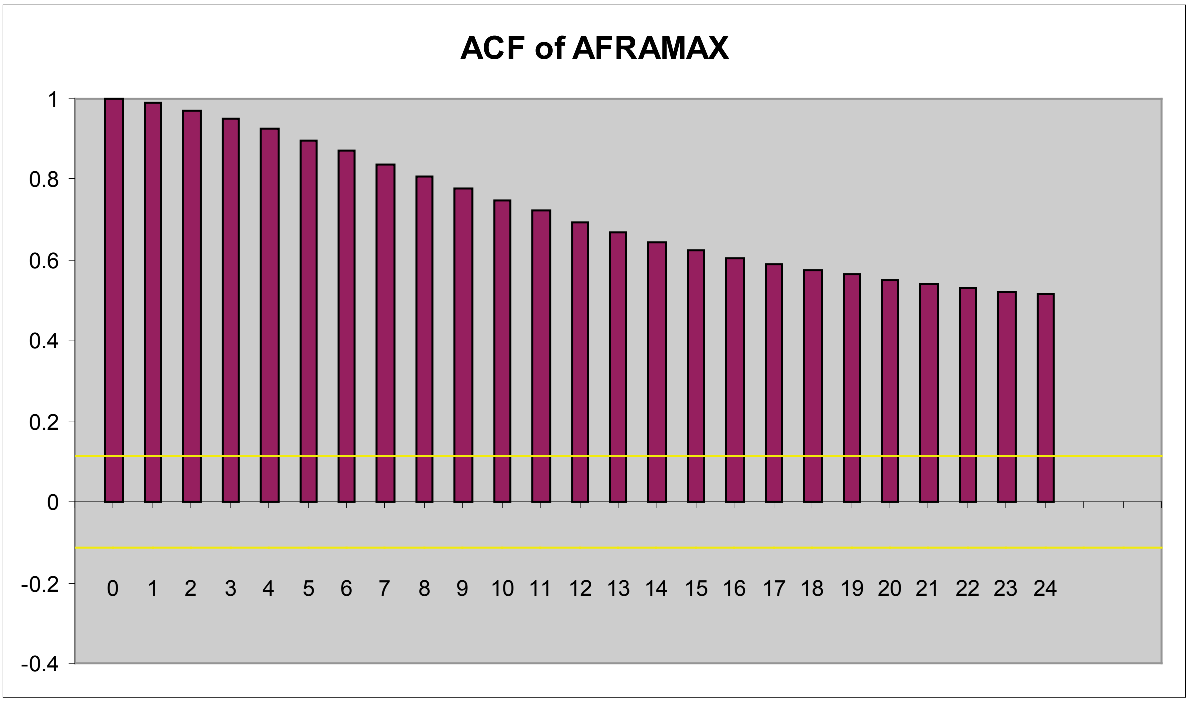

Figure 35 : ACE of AFRAMAX Time Charter Rates

We can see that the AFRAMAX time charter rate autocorrelation plot decays slowly. This is typical of a *non-stationary* series. This means that the AFRAMAX time charter rate should be made stationary in order for us to be able to analyze it. After differencing, the autocorrelation plot should quickly fall down to nominal values.

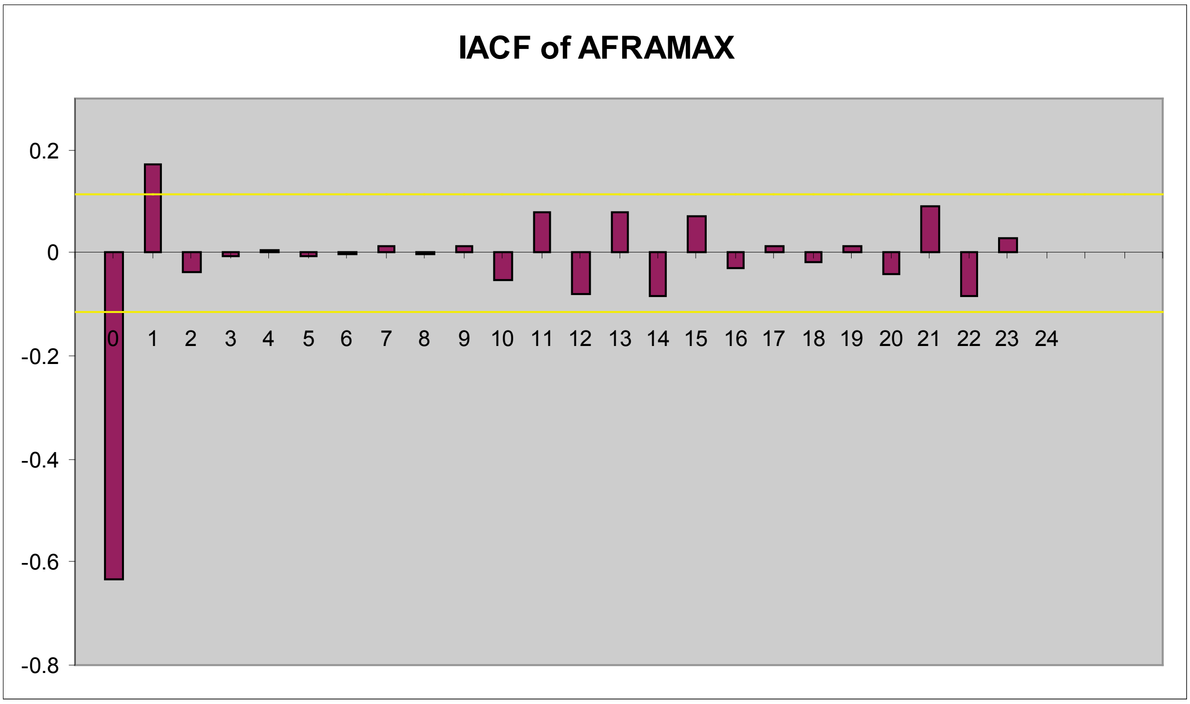

Figure 36 : IACF of AFRAMAX Time Charter Rate

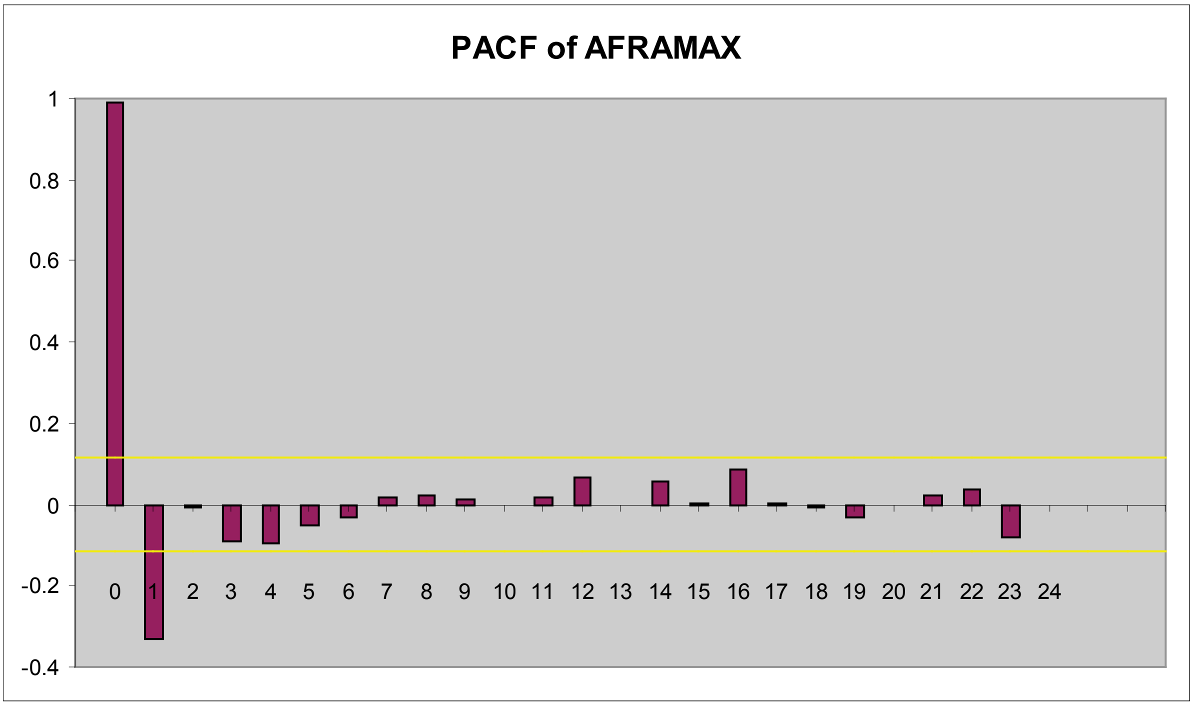

Figure 37 : PACF of AFRAMAX Time Charter Rates

If the time series can be modeled using a simple autoregressive process AR=1 then the partial autocorrelation function of that series will die down immediately after nlag=1. In this case, the AFRAMAX's time charter rate's partial autocorrelation function dies out quickly indicating that a simple autoregressive process will be sufficient to model the process.

Autocorrelation Check for White Noise

To Lag	Chi-Square	DF	Pr > ChiSq	Autocorrelations					
6	1518.79	6	<.0001	0.989	0.970	0.948	0.922	0.893	0.863
12	2538.24	12	<.0001	0.832	0.799	0.767	0.737	0.708	0.680
18	3215.42	18	<.0001	0.655	0.632	0.613	0.596	0.583	0.571
24	3760.45	24	<.0001	0.561	0.551	0.543	0.536	0.529	0.521

For the Aframax time charter rate, the white noise hypothesis is rejected very strongly, which is expected since the series is nonstationary.

Type	Lags	Rho	Pr < Rho	Tau	Pr < Tau	F	Pr > F
Zero Mean	0	0.0372	0.6910	0.04	0.6960		
	1	-0.5863	0.5512	-0.46	0.5131		
	2	-0.6495	0.5381	-0.48	0.5075		
Single Mean	0	-2.6808	0.6946	-1.11	0.7123	0.73	0.8834
	1	-5.6790	0.3712	-1.64	0.4606	1.36	0.7239
	2	-6.6813	0.2938	-1.77	0.3956	1.58	0.6666
Trend	0	-6.9855	0.6615	-1.96	0.6234	1.95	0.7874
	1	-15.5282	0.1608	-3.00	0.1342	4.61	0.2503
	2	-17.8504	0.1015	-3.09	0.1114	4.83	0.2072

The series is, of course, nonstationary as shown by the Dickey-Fuller Unit Root Test.

3.4 AFRAMAX Returns

We convert the AFRAMAX time charter rate to returns and perform the same statistical analysis on the AFRAMAX returns as in the VLCC returns (8.25).

```
Name of Variable = aframax_r

Mean of Working Series     0.00153
Standard Deviation        0.076351
Number of Observations         285
```

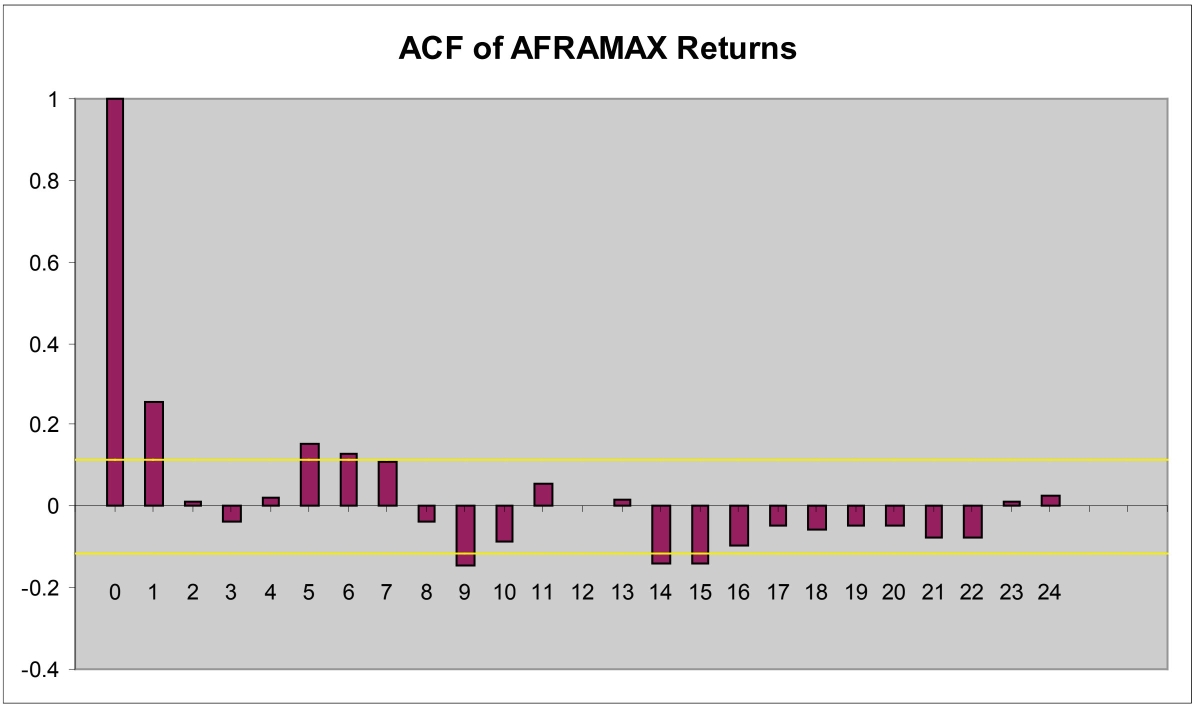

Figure 38 : ACE of AFRAMAX Returns

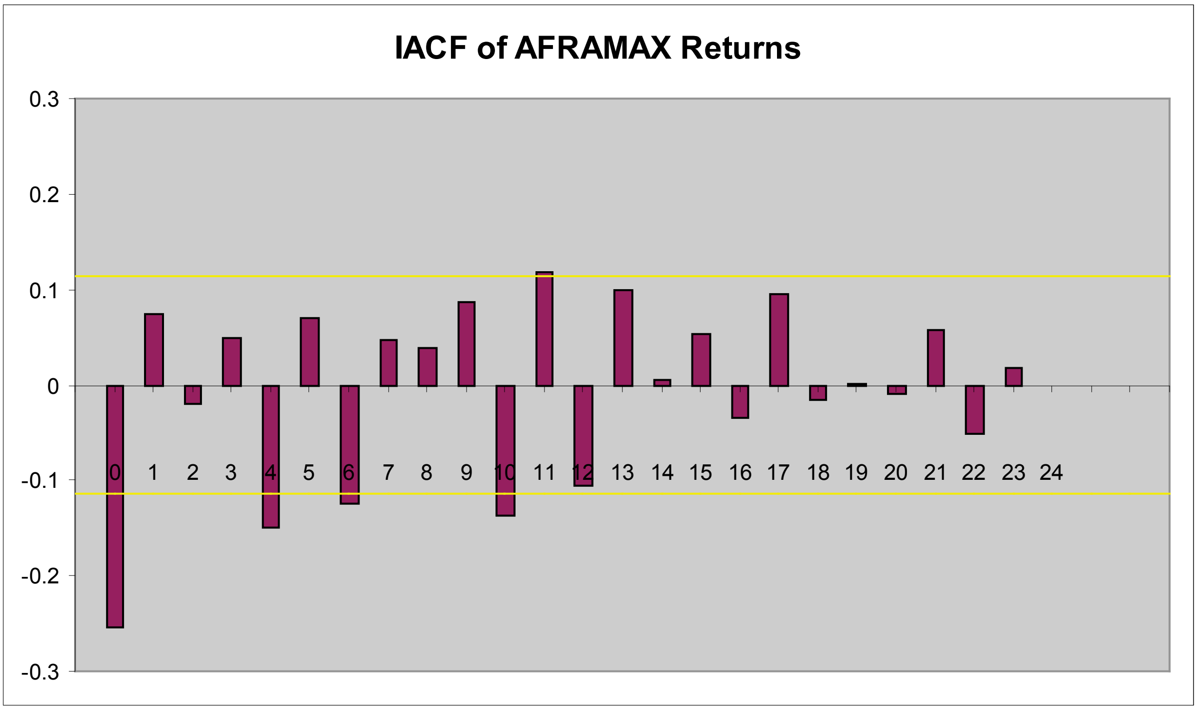

Figure 39 : IACF of AFRAMAX Returns

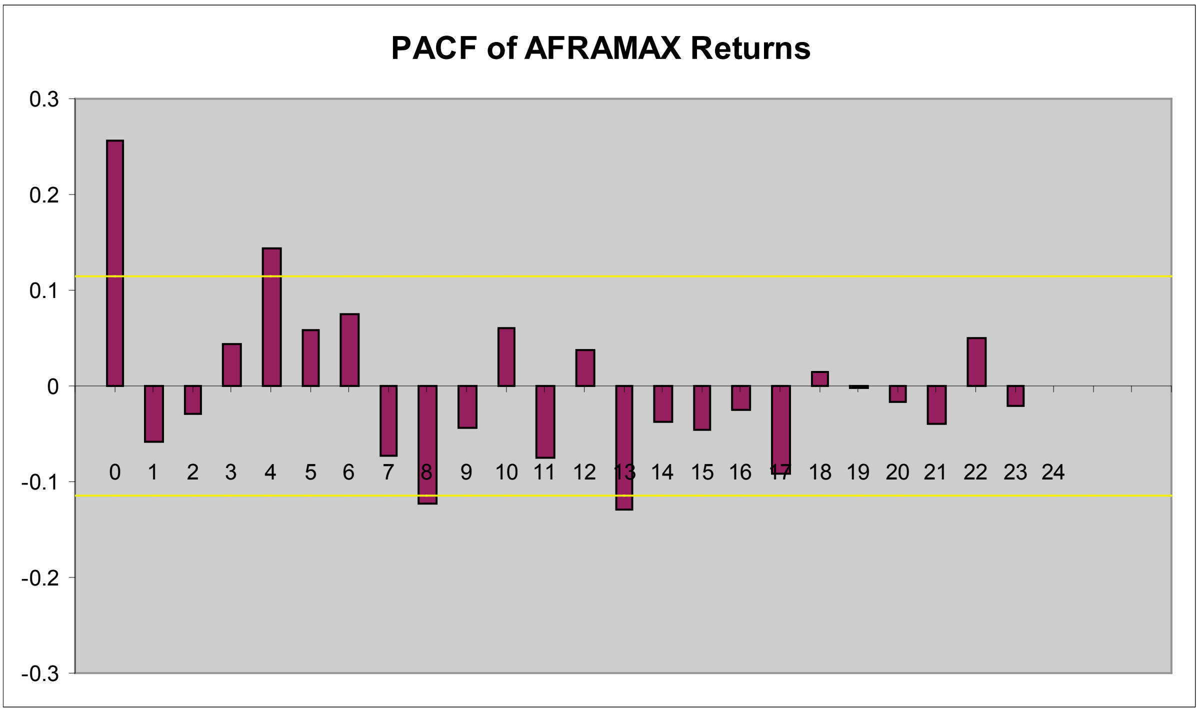

Figure 40 : PACF of AFRAMAX Returns

Unlike the untransformed series, the AFRAMAX's time charter rate's returns partial autocorrelation function do not die out quickly indicating that a simple autoregressive process will not be sufficient to model the process. As with the previous time charter rate we run the Autocorellation and Unit Root tests:

Autocorrelation Check for White Noise

To Lag	Chi-Square	DF	Pr > ChiSq	Autocorrelations					
6	31.55	6	<.0001	0.259	0.012	-0.027	0.007	0.138	0.148
12	45.51	12	<.0001	0.116	-0.037	-0.138	-0.104	0.050	0.005
18	63.61	18	<.0001	0.007	-0.135	-0.156	-0.098	-0.055	-0.067
24	68.79	24	<.0001	-0.053	-0.049	-0.078	-0.069	0.011	0.023

Augmented Dickey-Fuller Unit Root Tests

Type	Lags	Rho	Pr < Rho	Tau	Pr < Tau	F	Pr > F
Zero Mean	0	-210.409	0.0001	-13.02	<.0001		
	1	-234.789	0.0001	-10.83	<.0001		
	2	-239.497	0.0001	-9.18	<.0001		
Single Mean	0	-210.450	0.0001	-13.00	<.0001	84.52	0.0010
	1	-234.926	0.0001	-10.81	<.0001	58.44	0.0010
	2	-239.927	0.0001	-9.17	<.0001	42.04	0.0010
Trend	0	-210.679	0.0001	-12.99	<.0001	84.44	0.0010
	1	-235.540	0.0001	-10.81	<.0001	58.40	0.0010
	2	-240.602	0.0001	-9.15	<.0001	41.93	0.0010

The stationarity test of the return series shows that the hypothesis is rejected (no unit root) and that the series is stationary.

Alternate Data Transformation

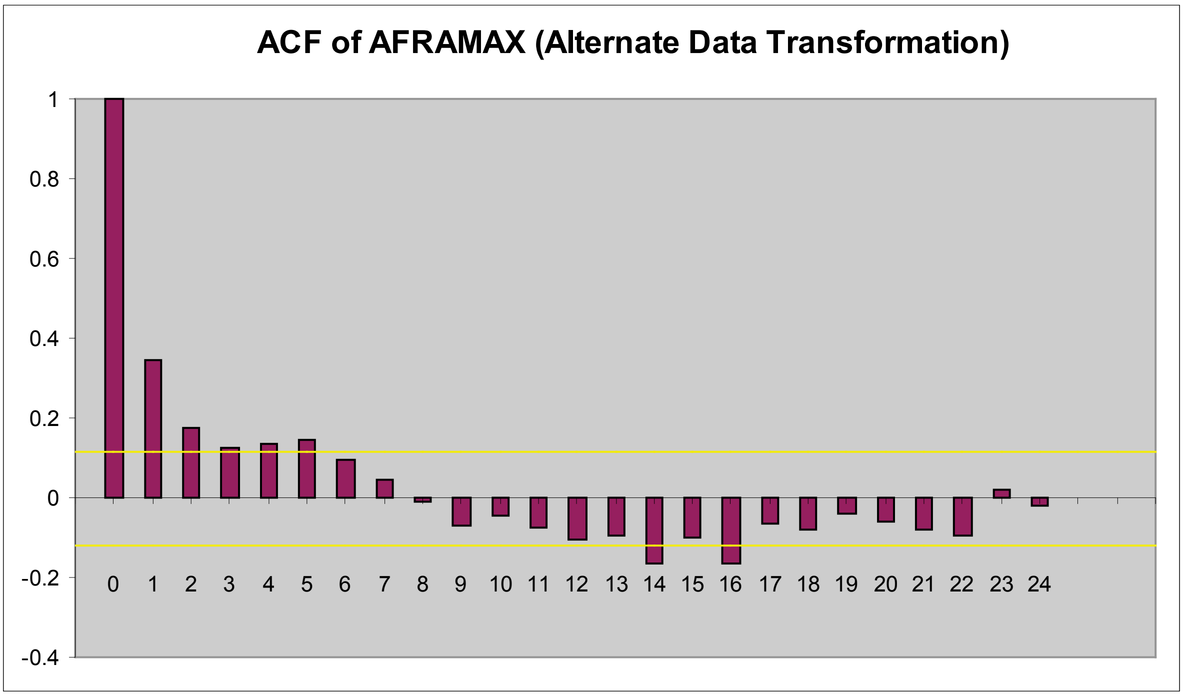

Figure 41 : Autocorrelation Function of AFRAMAX (Alternate Data Transofrmation)

There is little difference between the original and alternate data transformation – in order to preserve the homogeneity of the data transformations, the original returns will be used.

3.5 Analysis of HANDYSIZE Time Charter Rates

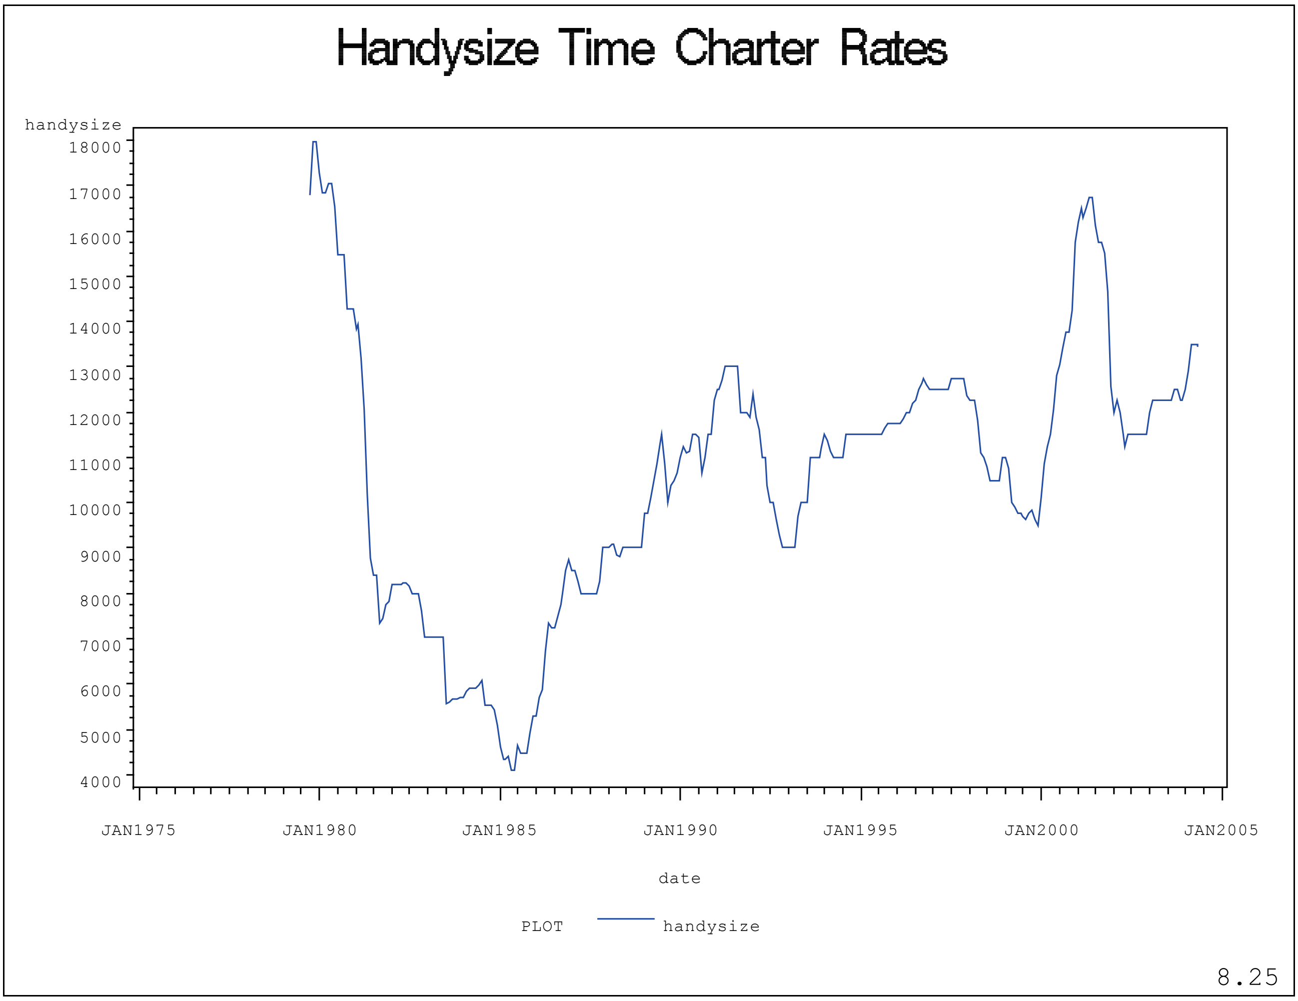

Figure 42 : Handysize Time Charter Rates

Using (8.25) we can analyze the HANDYSIZE Time Charter Rate

```
Mean of Working Series    10512.09
Standard Deviation        2980.418
Number of Observations         286
```

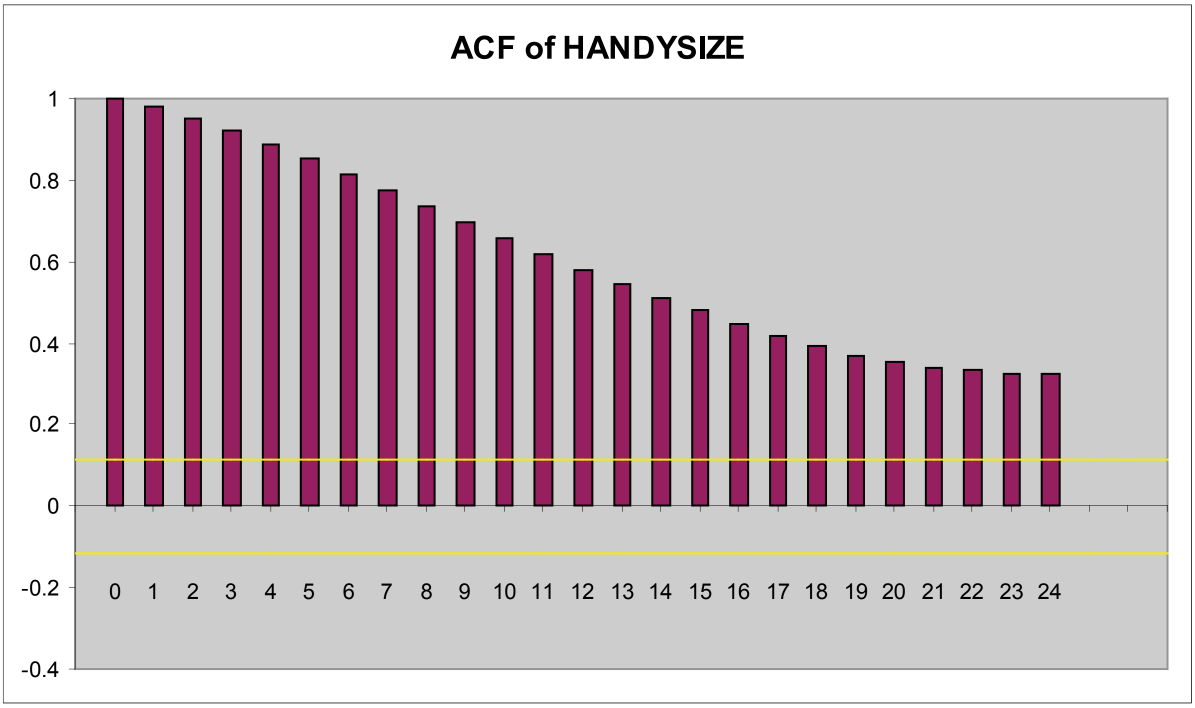

Figure 43 : ACE of HANDYSIZE Time Charter Rates

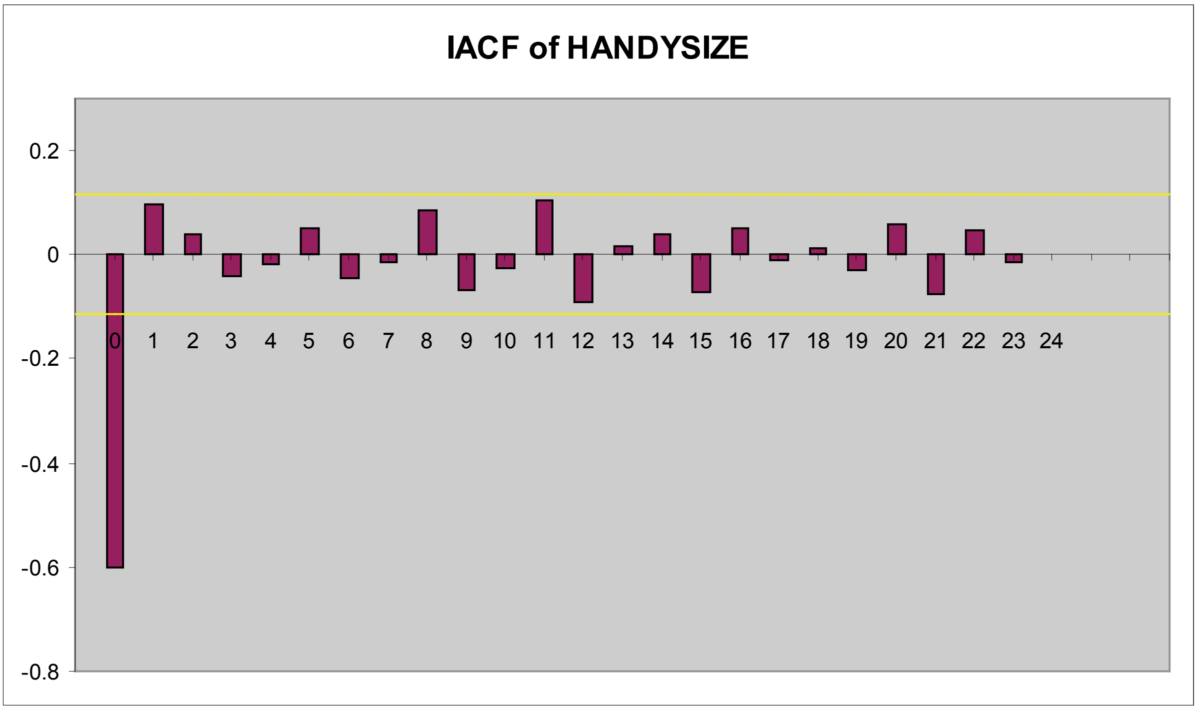

Figure 44 : IACF of HANDYSIZE Time Charter Rates

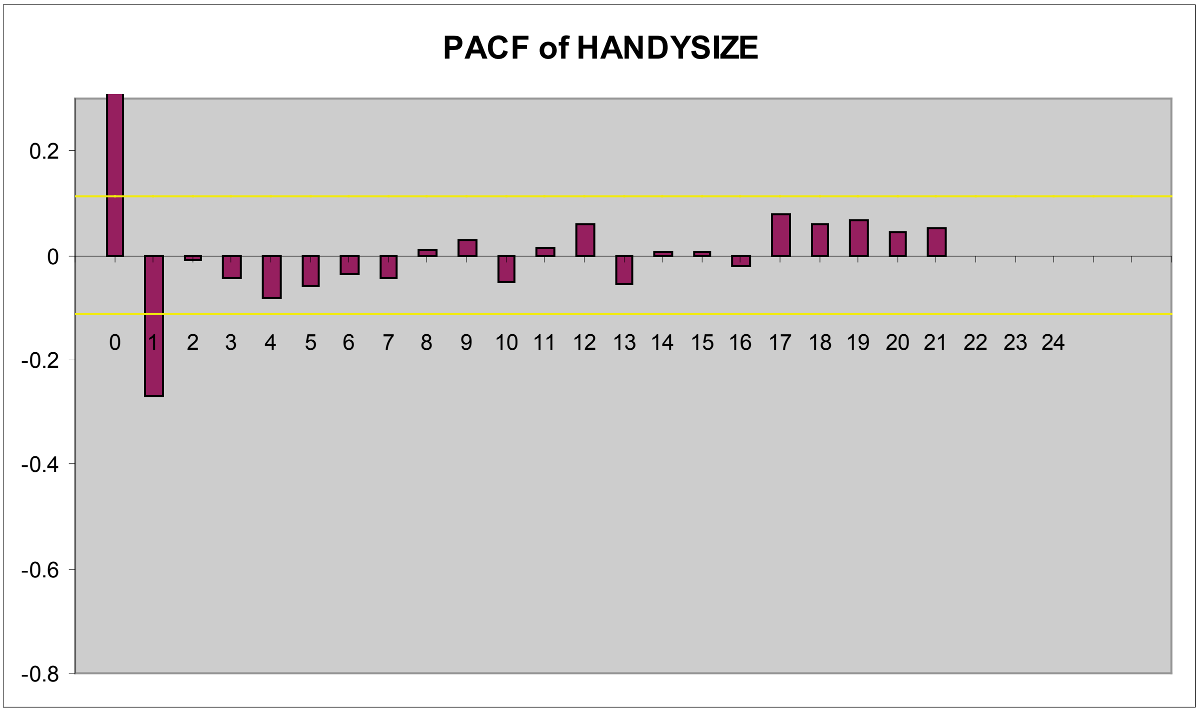

Figure 45 : PACF of HANDYSIZE Time Charter Rates

The HANDYSIZE time charter rate's partial autocorrelation function dies out quickly indicating that a simple autoregressive process will be sufficient to model the process.

The Autocorrelation and Unit Root Tests show:

Autocorrelation Check for White Noise

To Lag	Chi-Square	DF	Pr > ChiSq	Autocorrelations					
6	1424.86	6	<.0001	0.981	0.953	0.921	0.887	0.851	0.812
12	2233.76	12	<.0001	0.770	0.728	0.688	0.650	0.611	0.571
18	2625.83	18	<.0001	0.537	0.505	0.473	0.444	0.416	0.390
24	2838.55	24	<.0001	0.369	0.354	0.340	0.328	0.318	0.312

Augmented Dickey-Fuller Unit Root Tests

Type	Lags	Rho	Pr < Rho	Tau	Pr < Tau	F	Pr > F
Zero Mean	0	-0.6008	0.5481	-0.96	0.2990		
	1	-1.1019	0.4552	-1.17	0.2212		
	2	-1.0236	0.4683	-1.04	0.2680		
Single Mean	0	-4.3007	0.5043	-1.89	0.3378	1.89	0.5870
	1	-9.2515	0.1584	-2.68	0.0796	3.70	0.1243
	2	-9.3433	0.1548	-2.55	0.1056	3.32	0.2213
Trend	0	-8.0165	0.5781	-3.26	0.0754	8.33	0.0011
	1	-14.6169	0.1914	-4.06	0.0081	10.37	0.0010
	2	-14.9090	0.1810	-3.92	0.0124	9.48	0.0010

For the Handysize time charter rate, the white noise hypothesis is rejected very strongly, which is expected since the series is nonstationary. The *p* value for the test of the first six autocorrelations is printed as <0.0001, which means the *p* value is less than .0001.

3.6 HANDYSIZE Returns

We convert the HANDYSIZE time charter rates into returns and perform the same statistical analysis using 8.28:

```
Name of Variable = handysize_r

Mean of Working Series    -0.00074
Standard Deviation        0.040485
Number of Observations         285
```

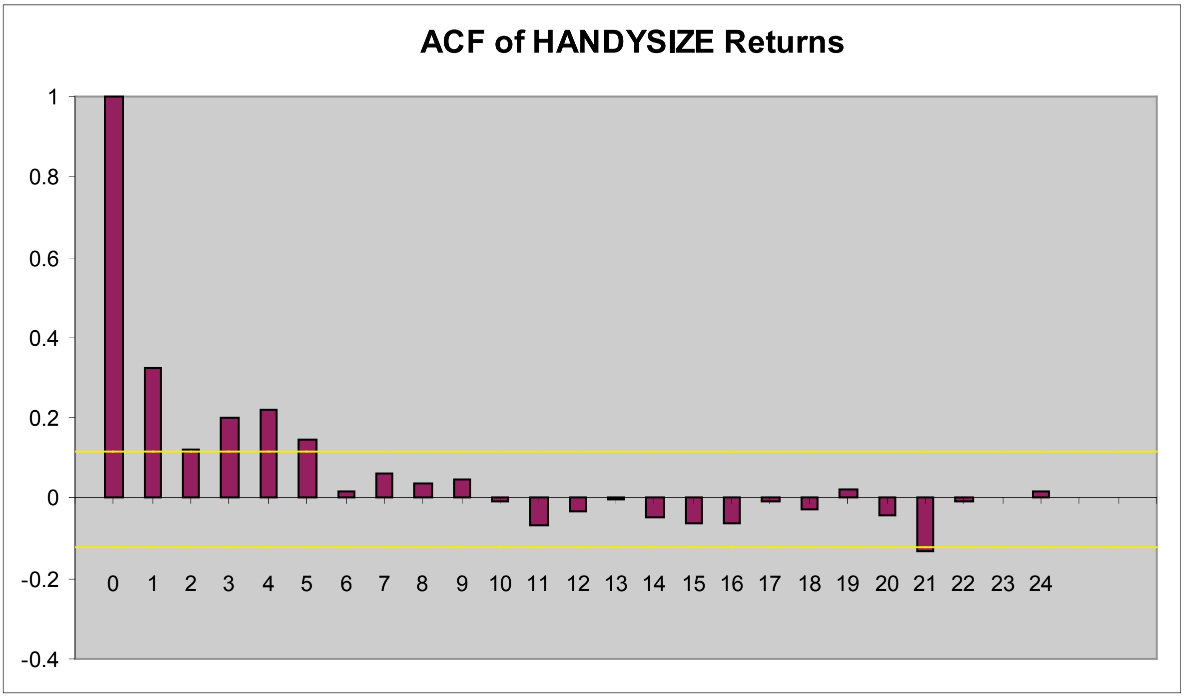

Figure 46 : ACF of Handysize Returns

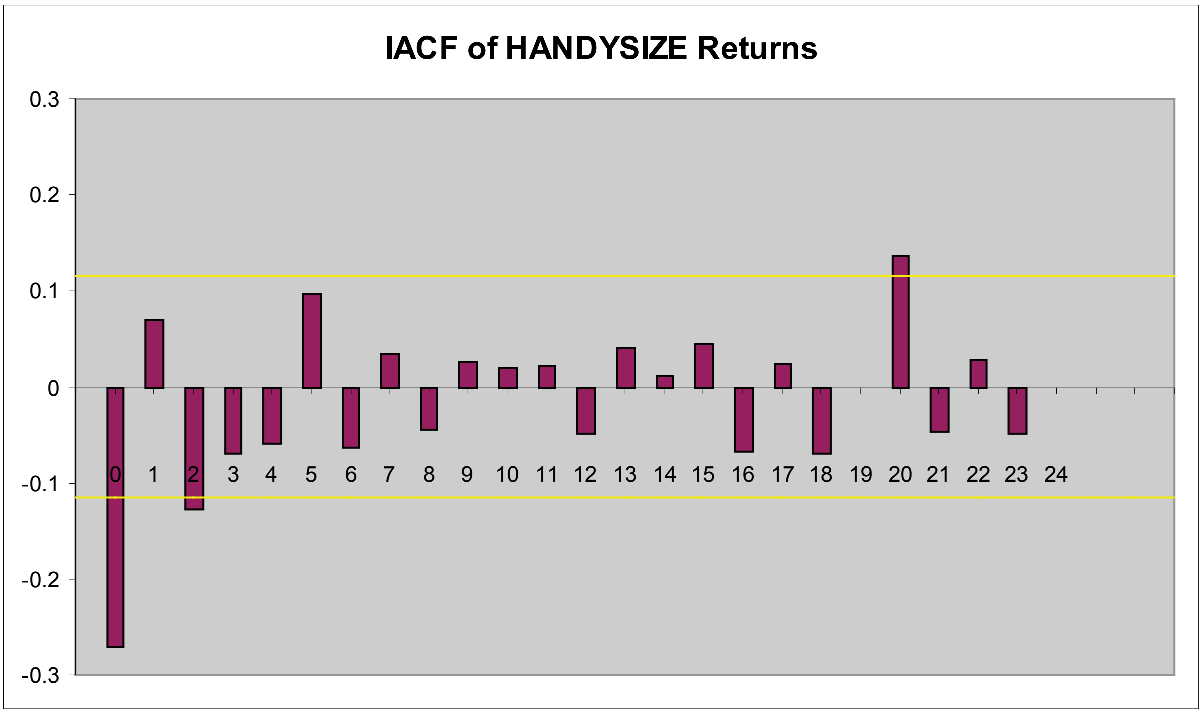

Figure 47 : IACF of HANDYSIZE Returns

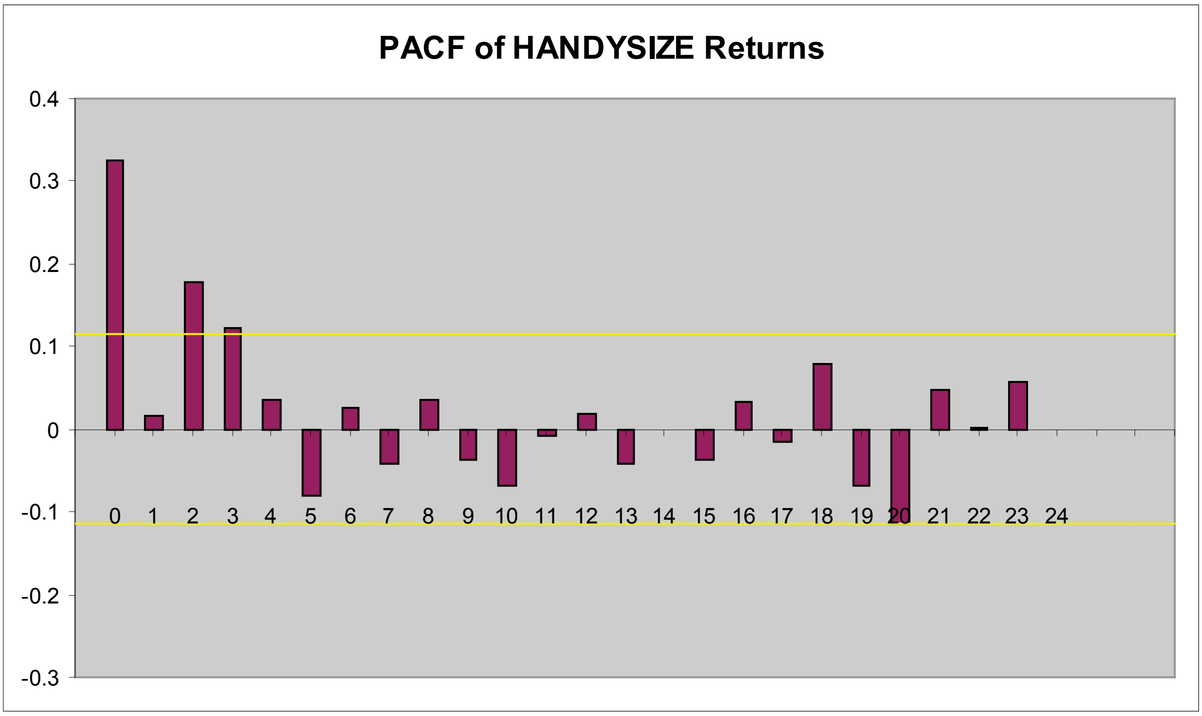

Figure 48 : PACF of HANDYSIZE Returns

Unlike the untransformed series, the Handysize's time charter rate's returns partial autocorrelation function do not die out quickly indicating that a simple autoregressive process will not be sufficient to model the process. As with the previous time charter rate we run the Autocorrelation and Unit Root Tests.

To Lag	Chi-Square	DF	Pr > ChiSq	Autocorrelations					
6	73.38	6	<.0001	0.344	0.156	0.205	0.211	0.150	0.046
12	77.93	12	<.0001	0.055	0.022	0.056	0.000	-0.053	-0.077
18	82.79	18	<.0001	-0.037	-0.050	-0.062	-0.082	-0.002	-0.040
24	87.86	24	<.0001	0.002	-0.043	-0.117	-0.010	-0.027	0.008

Augmented Dickey-Fuller Unit Root Tests

Type	Lags	Rho	Pr < Rho	Tau	Pr < Tau	F	Pr > F
Zero Mean	0	-185.946	0.0001	-11.78	<.0001		
	1	-169.437	0.0001	-9.17	<.0001		
	2	-112.607	0.0001	-6.81	<.0001		
Single Mean	0	-186.048	0.0001	-11.76	<.0001	69.20	0.0010
	1	-169.612	0.0001	-9.16	<.0001	41.93	0.0010
	2	-112.731	0.0001	-6.79	<.0001	23.09	0.0010
Trend	0	-189.245	0.0001	-11.92	<.0001	71.10	0.0010
	1	-175.329	0.0001	-9.32	<.0001	43.41	0.0010
	2	-118.025	0.0001	-6.92	<.0001	23.98	0.0010

The stationarity test of the return series shows that the hypothesis is rejected and that the series is stationary.

Alternate Data Transformation

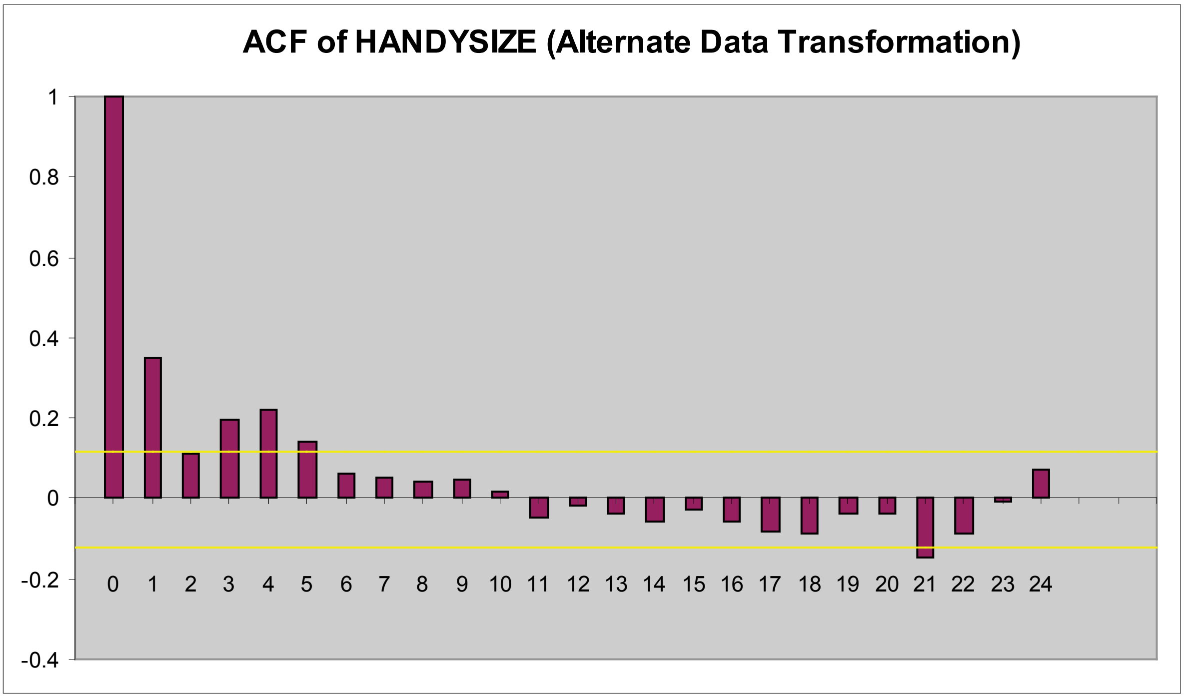

Figure 49 : Autocorrelation Function of HANDYSIZE (Alternate Data Transformation)

There is little difference between the original and alternate data transformation – in order to preserve the homogeneity of the data transformations, the original returns will be used.

3.7 Analysis of SUEZMAX Time Charter Rates

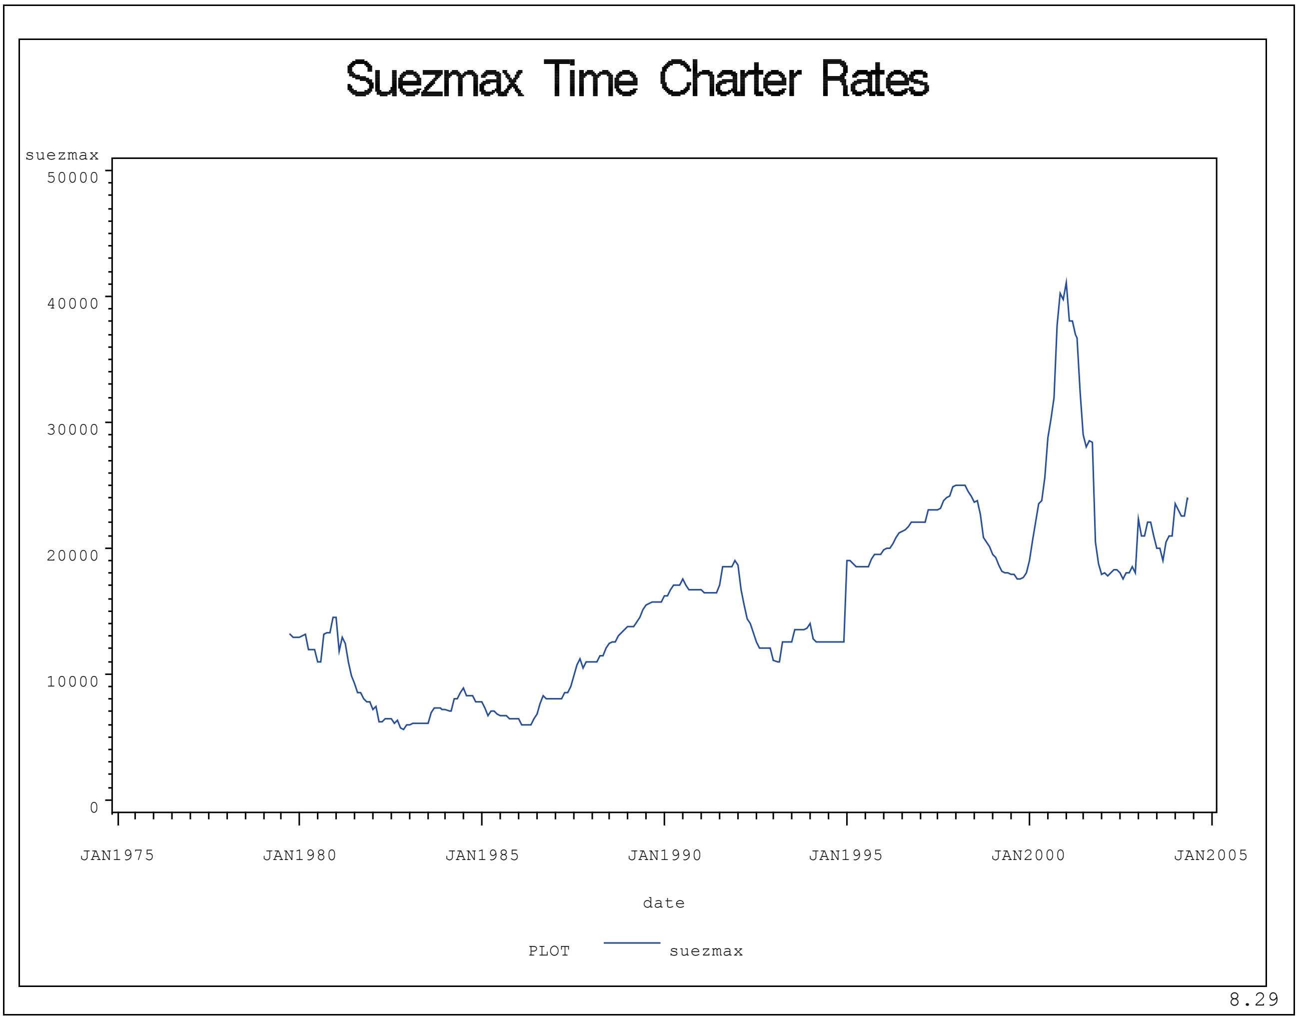

Figure 50 : SUEZMAX Time Charter Rates

Using (8.30) we obtain:

```
Name of Variable = suezmax

Mean of Working Series    15425.28
Standard Deviation        7274.987
Number of Observations         286
```

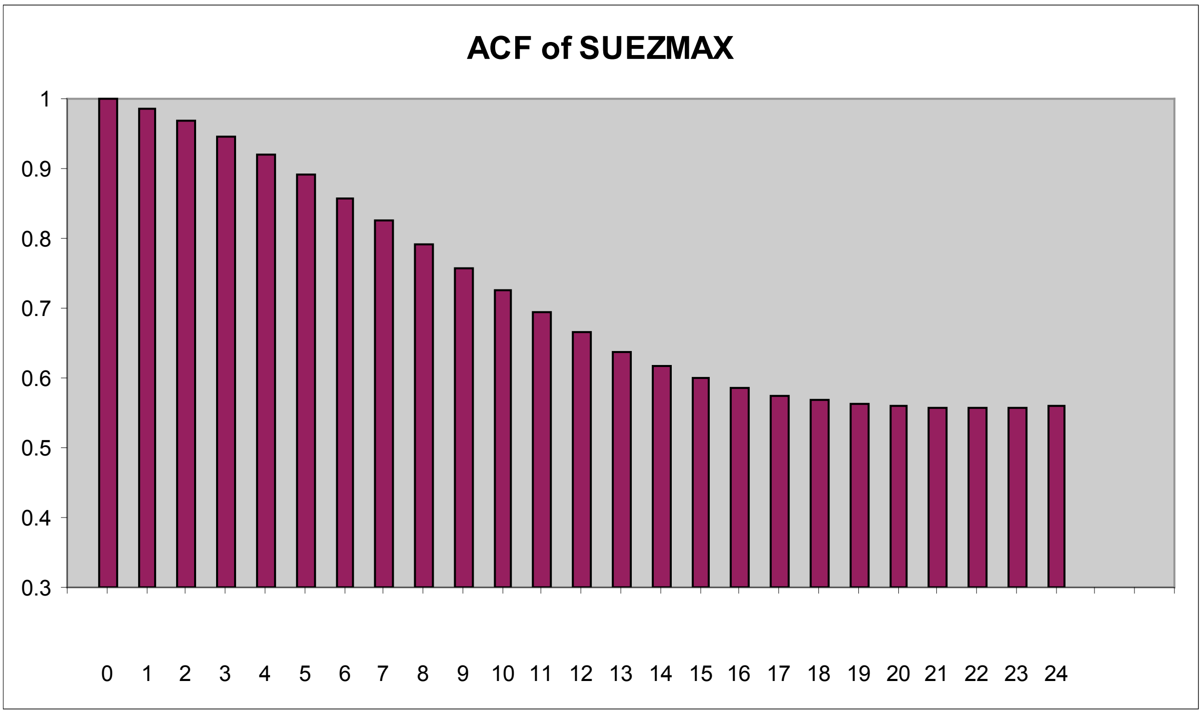

Figure 51 : ACE of SUEZMAX Time Charter Rates

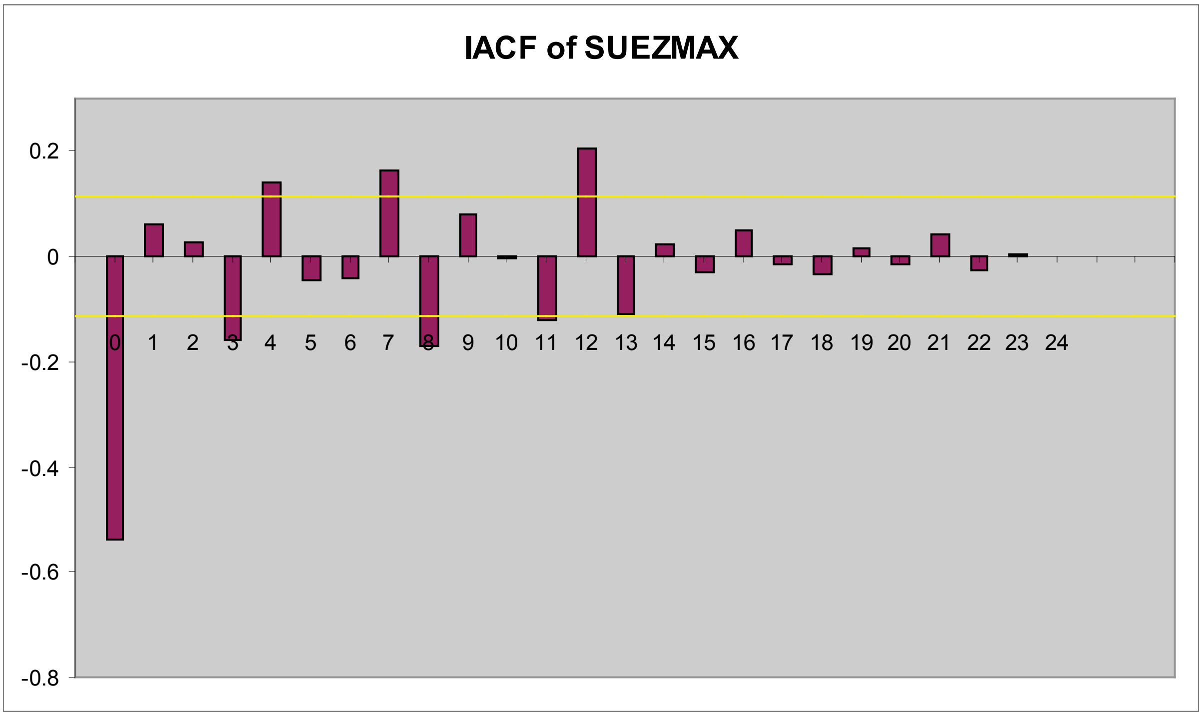

Figure 52 : IACF of SUEZMAX Time Charter Rates

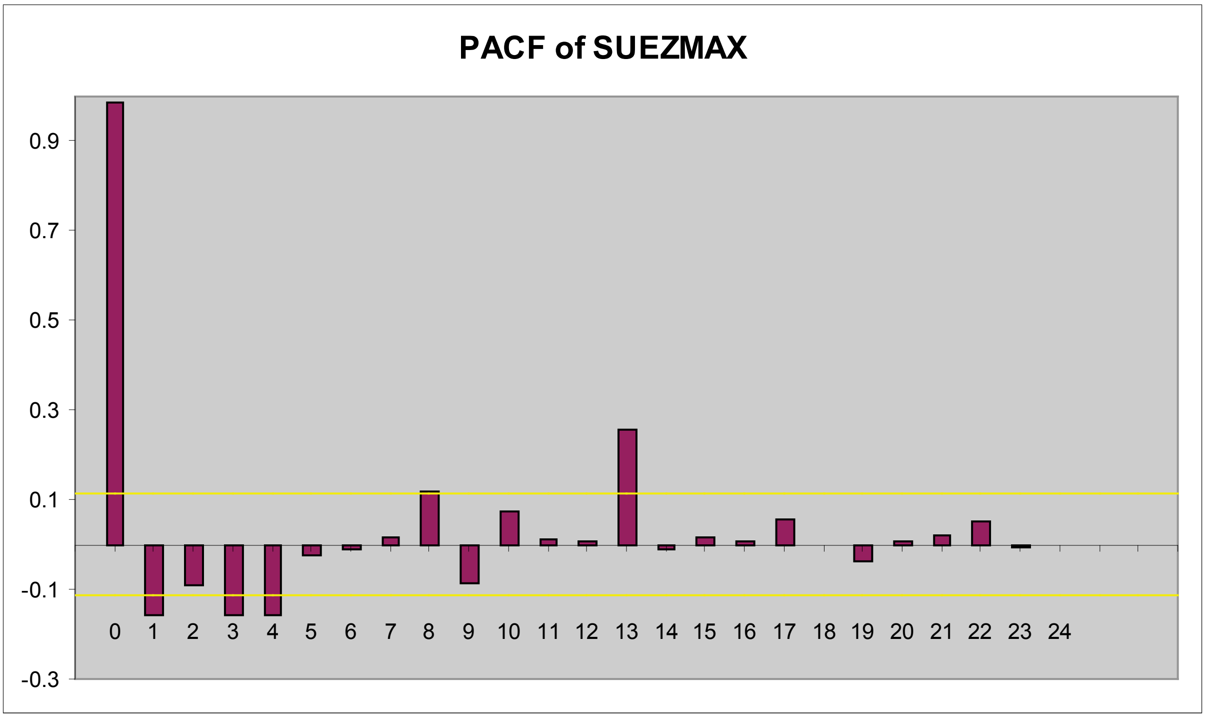

Figure 53 : PACF of SUEZMAX Time Charter Rates

The SUEZMAX time charter rate's partial autocorrelation function does not die out quickly indicating that a simple autoregressive process will not be sufficient to model the process. The Autocorrelation and Unit Root Tests follow:

Autocorrelation Check for White Noise

To Lag	Chi-Square	DF	Pr > ChiSq	Autocorrelations					
6	1486.02	6	<.0001	0.987	0.966	0.941	0.912	0.877	0.840
12	2407.44	12	<.0001	0.802	0.764	0.731	0.697	0.666	0.638
18	3031.63	18	<.0001	0.614	0.597	0.584	0.575	0.570	0.567
24	3641.88	24	<.0001	0.567	0.569	0.572	0.572	0.573	0.575

Augmented Dickey-Fuller Unit Root Tests

Type	Lags	Rho	Pr < Rho	Tau	Pr < Tau	F	Pr > F
Zero Mean	0	0.0526	0.6946	0.05	0.6984		
	1	-0.5991	0.5485	-0.39	0.5424		
	2	-0.8266	0.5035	-0.49	0.5008		
Single Mean	0	-2.5638	0.7091	-1.03	0.7410	0.68	0.8967
	1	-6.0936	0.3373	-1.68	0.4389	1.48	0.6920
	2	-7.4283	0.2459	-1.85	0.3546	1.77	0.6187
Trend	0	-9.6527	0.4547	-2.39	0.3851	2.98	0.5786
	1	-19.0481	0.0794	-3.18	0.0910	5.09	0.1545
	2	-23.8268	0.0285	-3.47	0.0449	6.05	0.0635

For the Suezmax time charter rate, the white noise hypothesis is rejected very strongly, which is expected since the series is nonstationary. The p value for the test of the first six autocorrelations is printed as <0.0001, which means the p value is less than .0001.

3.8 SUEZMAX Returns

Using (8.31) we convert the SUEZMAX time charter rate to returns in order to achieve stationarity and perform a statistical analysis on the returns:

```
  Name of Variable = suezmax_r

Mean of Working Series    0.001948
Standard Deviation        0.057972
Number of Observations         285
```

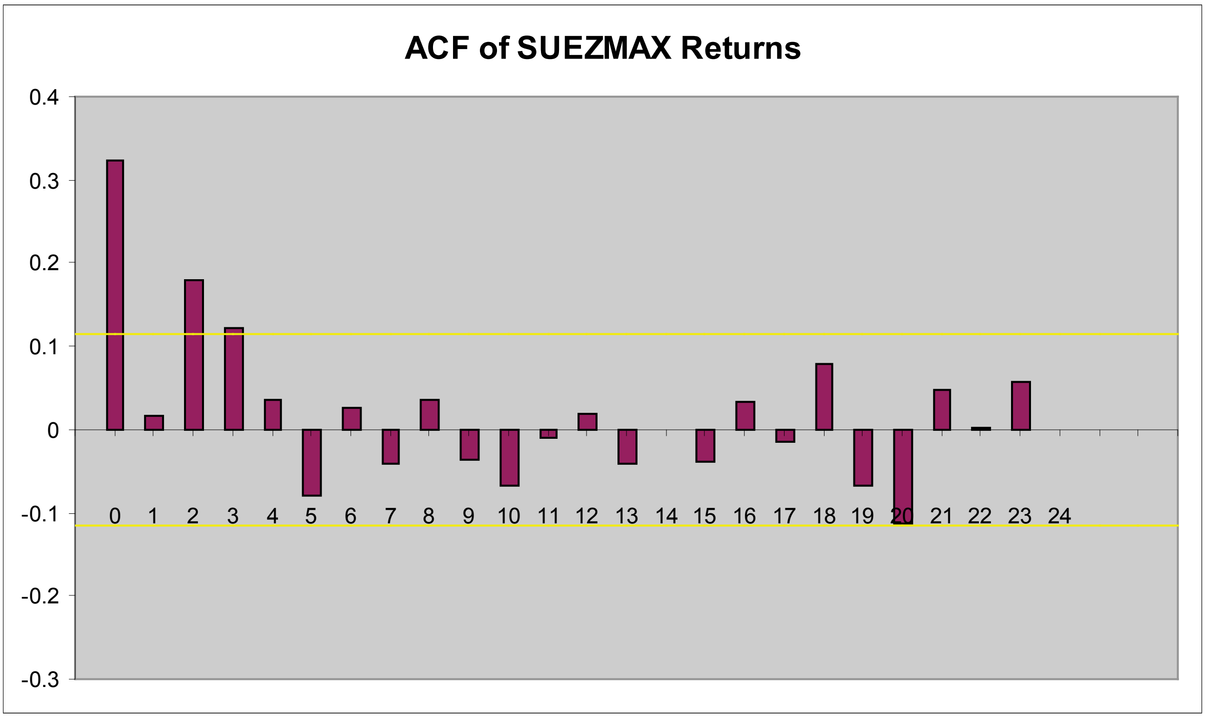

Figure 54 : ACF of SUEZMAX RETURNS

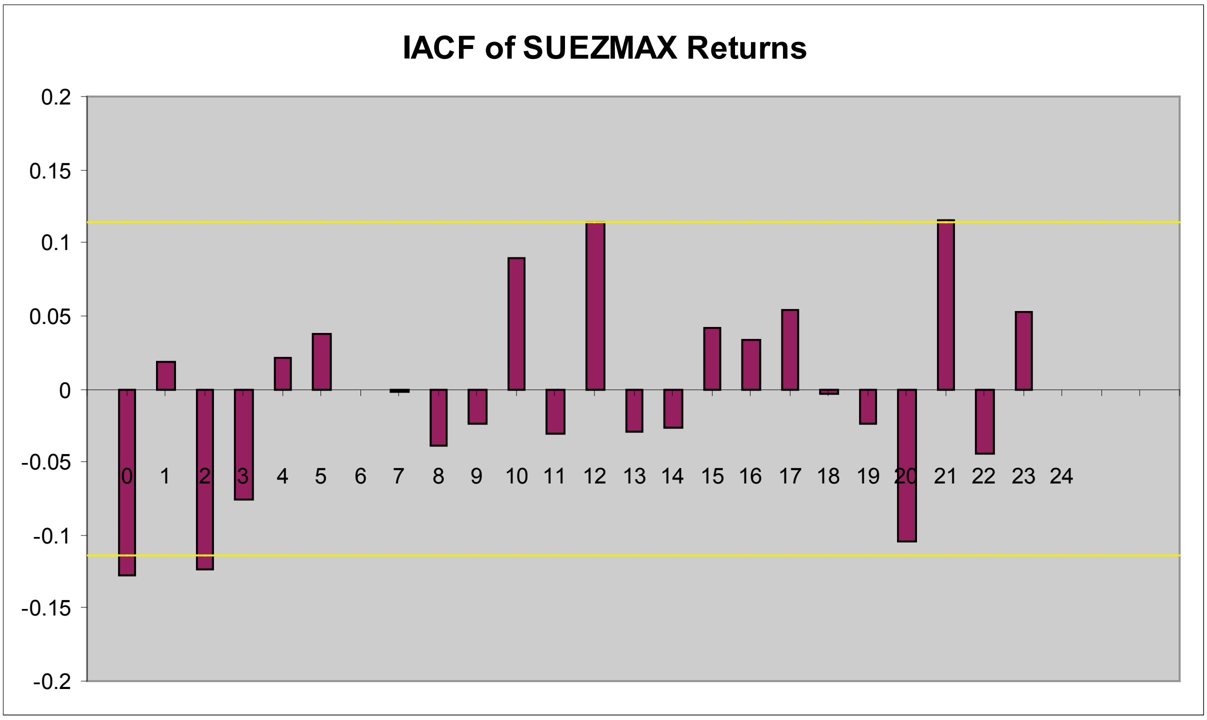

Figure 55 : IACF of SUEZMAX RETURNS

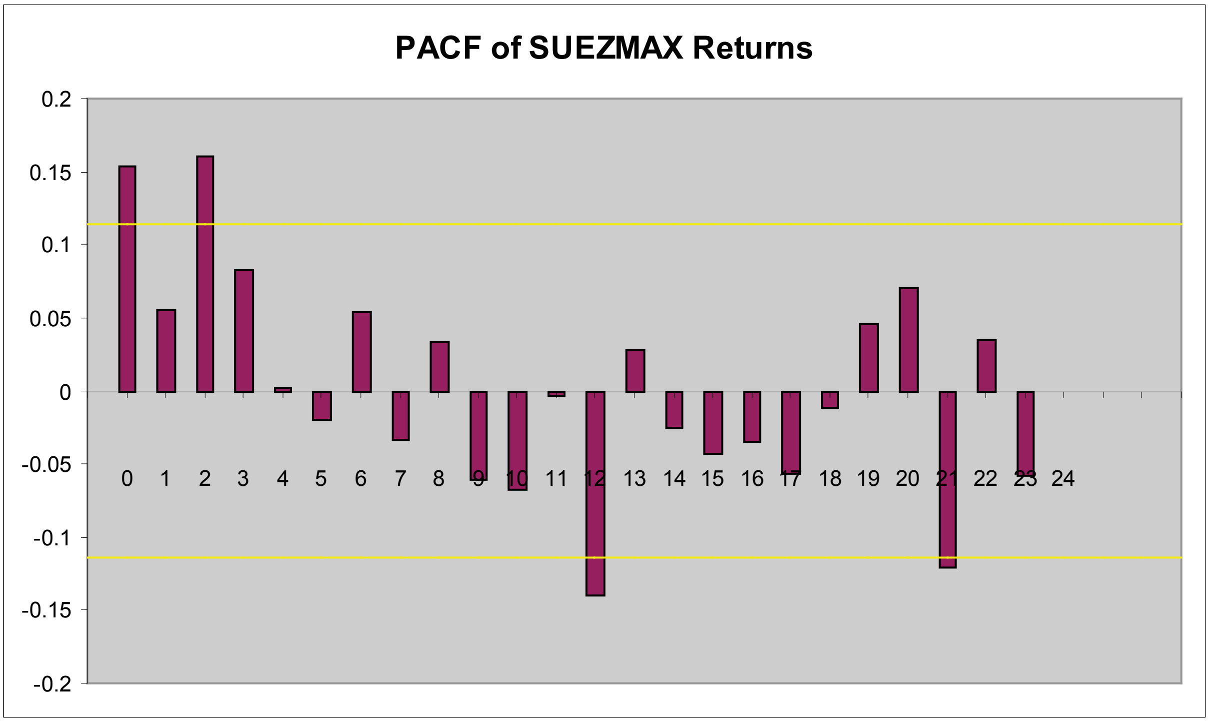

Figure 56 : PACF of SUEZMAX Returns

The Suezmax's time charter rate's partial autocorrelation does not die out quickly indicating that a simple autoregressive process will not be sufficient to model the process. The Autocorrelation and Unit Root Tests follow:

To Lag	Chi-Square	DF	Pr > ChiSq	Autocorrelations					
6	36.35	6	<.0001	0.247	0.124	0.189	0.102	0.044	0.036
12	43.21	12	<.0001	0.072	0.027	0.023	-0.042	-0.085	-0.087
18	64.13	18	<.0001	-0.142	-0.087	-0.072	-0.092	-0.076	-0.148
24	71.79	24	<.0001	-0.108	-0.009	0.027	-0.107	-0.017	-0.023

Augmented Dickey-Fuller Unit Root Tests

Type	Lags	Rho	Pr < Rho	Tau	Pr < Tau	F	Pr > F
Zero Mean	0	-213.293	0.0001	-13.02	<.0001		
	1	-185.530	0.0001	-9.58	<.0001		
	2	-124.074	0.0001	-7.07	<.0001		
Single Mean	0	-213.565	0.0001	-13.01	<.0001	84.60	0.0010
	1	-186.034	0.0001	-9.57	<.0001	45.79	0.0010
	2	-124.559	0.0001	-7.07	<.0001	25.01	0.0010
Trend	0	-214.095	0.0001	-13.00	<.0001	84.56	0.0010
	1	-187.016	0.0001	-9.57	<.0001	45.81	0.0010
	2	-125.442	0.0001	-7.07	<.0001	25.02	0.0010

Alternate Data Transformation

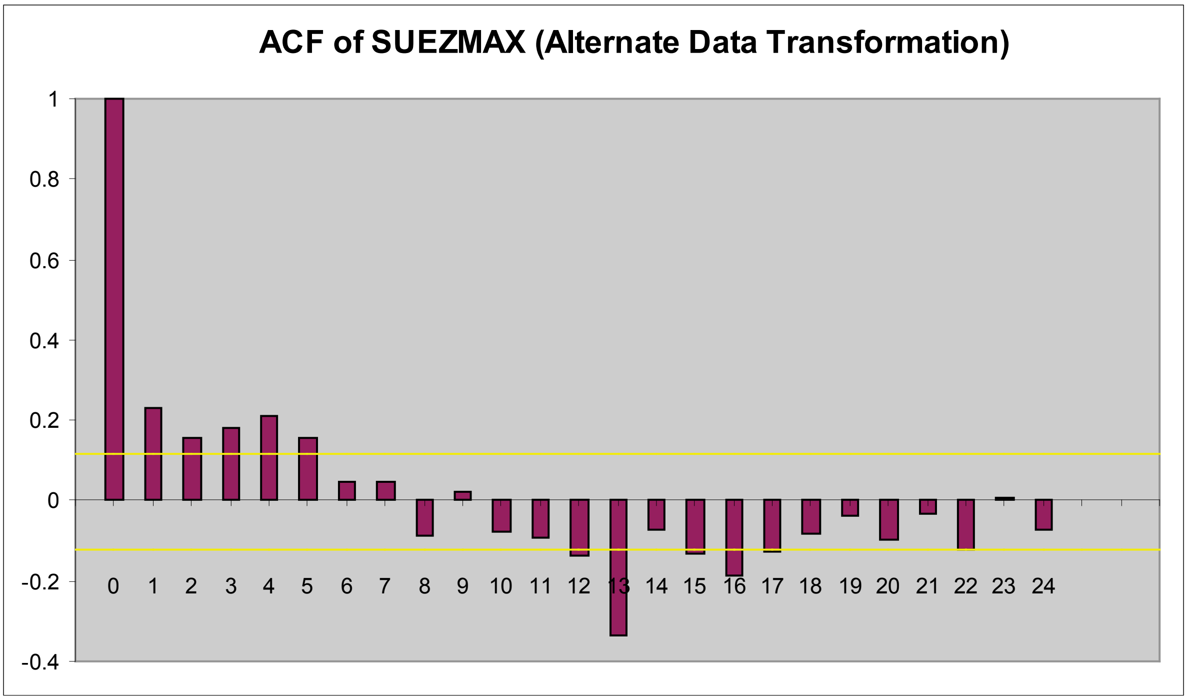

Figure 57 : Autocorrelation Function of SUEZMAX (Alternate Data Transformation)

We can safely assume the original transformation yields a more stationary series for this data. We have shown that in that the original autocorrelation plot for each time charter rate decays slowly, indicating a non-stationary series. By converting the series to returns (a standard practice) we can create a stationary series. This can be proved either by examining the autocorrelation plot which decays very quickly or by running either the Unit Root test or the check for white noise. It is worth mentioning that some series become more or less stationary and may require further or even different transformation in order to achieve strict stationarity requirements.

4. ARIMA MODELS

General

It has been shown that in order to make a non-stationary time series stationary, it must undergo certain transformations that render it stationary. Its also worth noting that, according to the results of the Partial Autocorrelation Functions for both time charter rates and their returns, a transformed series, while achieving stationarity, requires a more complex autoregressive model. One such transformation, differencing, will be used in order to show the theoretical process behind Autoregressive Integrated Moving Average Models (ARIMA models). This linear regression model means that the dependent variable z_t is explained by and regressed on its previous values z_{t-i}. We can see why this process is called autoregressive. The order of the process corresponds to the number of lagged z's that are included in the model.

We are trying to construct a model and the only information we have is the z_t components. To identify a model which governs the behavior of the series, we must filter out the principal components of the series and then use these components to deduce an appropriate model.

The equation (1) can be broadened to include more lagged variables. For example, if events two periods ago had an effect on what is happening today, we can extend (1) to include z_{t-2} , thus expressing z_t as

$$z_t = \phi_1 z_{t-1} + \phi_2 z_{t-2} + a_t$$

Where ϕ_t and ϕ_2 are autoregressive parameters to be estimated and a_t the error term. A model with a p^{th} order autoregressive parameter is thus

$$z_t = \phi_1 z_{t-1} + \phi_2 z_{t-2} + \ldots + \phi_p z_{t-p} + a_t$$

A simple extension of the AR model would be to include past errors to see if they can improve the time series representation of the data. We modify an AR(1) model and obtain

$$z_t = \phi_1 z_{t-1} + a_t - \theta_1 a_{t-1}$$

Where a_t-1 is the error at period t-1 and θ_1 is the moving average parameter.

The moving average parameter can also be extended to include additional lagged residual terms. A qth order MA processs can be expressed as

$$z_t = a_t - \theta_1 a_{t-1} - \theta_2 a_{t-2} - \ldots - \theta_q a_{t-q}$$

Combining both Autoregressive and Moving Average processes, we can obtain the final general ARMA(p,q) equation:

$$z_t = \phi_1 z_{t-1} + \phi_2 z_{t-2} + \ldots + \phi_p z_{t-p} + a_t - \theta_1 a_{t-1} - \theta_2 a_{t-2} - \ldots - \theta_q a_{t-q}$$

Now let us suppose that in order to make a series stationary, we must difference the series. The first difference of a series z(t) is $z_t - z_{t-1}$ will constitute a new stationary series w_t. The ARIMA(p,d,q) process of the differences series is

$$w_t = \phi_1 w_{t-1} + \phi_2 w_{t-2} + \ldots + \phi_p w_{t-p} + a_t - \theta_1 a_{t-1} - \theta_2 a_{t-2} - \ldots - \theta_q a_{t-q} \quad (1)$$

Where a_t a random shock (also known as errors or white noise) and $\phi_t \ldots \phi_q$ the autoregressive parameter which describes the effect of a unit change in z_{t-1} on z_t (or in this case, of w_{t-1} on w_t).

In forecasting, the simplest model is considered the best and most easily accepted. This means that autoregressive parameters of a model must be checked so as to prevent overfitting. On the other hand, the diagnostics of the time charter rates show that a simple model would not be sufficient to capture their memory. As a result, more complex models with large autoregressive parameters were introduced. Overfitting has not been an issue since the significance of the autoregressive parameters remains high for the most part. In many cases, redundant autoregressive lags are followed by important ones.

Models are developed in an attempt to model *data*. While data for the VLCC time charter rate may be modeled by a certain ARIMA model, this model is particular to the time frame leading up to current data availability. In attempting to forecast the same data with different time frame, the model many times failed to converge. This means that a model that has been developed for a time frame from 0 to t may not be the best choice for a model that is developed for a time frame from 0 to t +/- i where i is a considerable time frame. In this case, ARIMA models that were tested with data ending a year before current data ended (i = 12) did not converge. When attempting to find out which models are best for modeling a series, different models must be tested for the same time periods.

Methodology

SAS gives you the ability to automatically fit ARIMA models to your data set using the time Series Forecasting System. This is done using the following steps.

1. Load the Time Series Forecasting System from the Solutions -> Analysis menu.
2. The Time Series Forecasting window will load. Select a data set by clicking on "Browse". Keep in mind, the data set should have been already created by running the "proc data" procedure in the command window. In this case, the name of the created dataset is "Dimitris" which will be loaded from the "Work" library.
3. After the data set has been loaded, click on the "Fit Models Automatically" button.
4. At this point, you can see the ARIMA models that are already built into SAS, or (as in most cases) you need to create one or more models to test. To add a model to the model list
    a. Right click anywhere on the Automatic Model Fitting Window and choose "Options" and "Model Selection List".
    b. The Model List is shown and by clicking on "Actions" and "Add ARIMA model" you can specify the AR, MA and seasonal regressesors you wish to use. The model is added to the Automatic Model Selection List.
    c. Return to the "Automatic Model Fitting" window and right click again. Select "Options" and "Automatic Fit…". Here you can choose the type of models to fit

and the models to keep. Select “All Autofit Models In Selection List” so that SAS includes the model you have created in its selection process. Also, depending on the forecasting needs, select which models to keep.

5. By default, all the variables in the data set (VLCC, AFRAMAX, HANDYSIZE, SUEZMAX) are selected to be forecasted. If you need to choose only one, click on the “Select” button to the right of the “Series to Process” window and you can select any combination of series in your data set.
6. Click on the “Select” button next to the Model Selection Criterion to select the type of criterion (R-Squared, Akaike etc.) you will use. (Click “Select All” to see all the model selection criteria).
7. Finally, returning once again to the “Automatic Model Fitting” window, click on “Run”.
8. The resulting window gives you a choice of various outputs concerning the model that has been created.

For each time charter rate we have run a list of all the possible combinations of ARIMA models up to five autoregressive lags and five moving average lags, and have picked the likeliest model candidate by using the criterion of the minimization of the Root Mean Squared Error. This process will be repeated or each time charter rate. Note that the ARIMA modeling procedure does not require the data to be converted to returns beforehand.

4.1 Prediction for AFRAMAX

Below is the SAS plot of the AFRAMAX time charter rate:

Figure 58 : Model Predictions from AFRAMAX

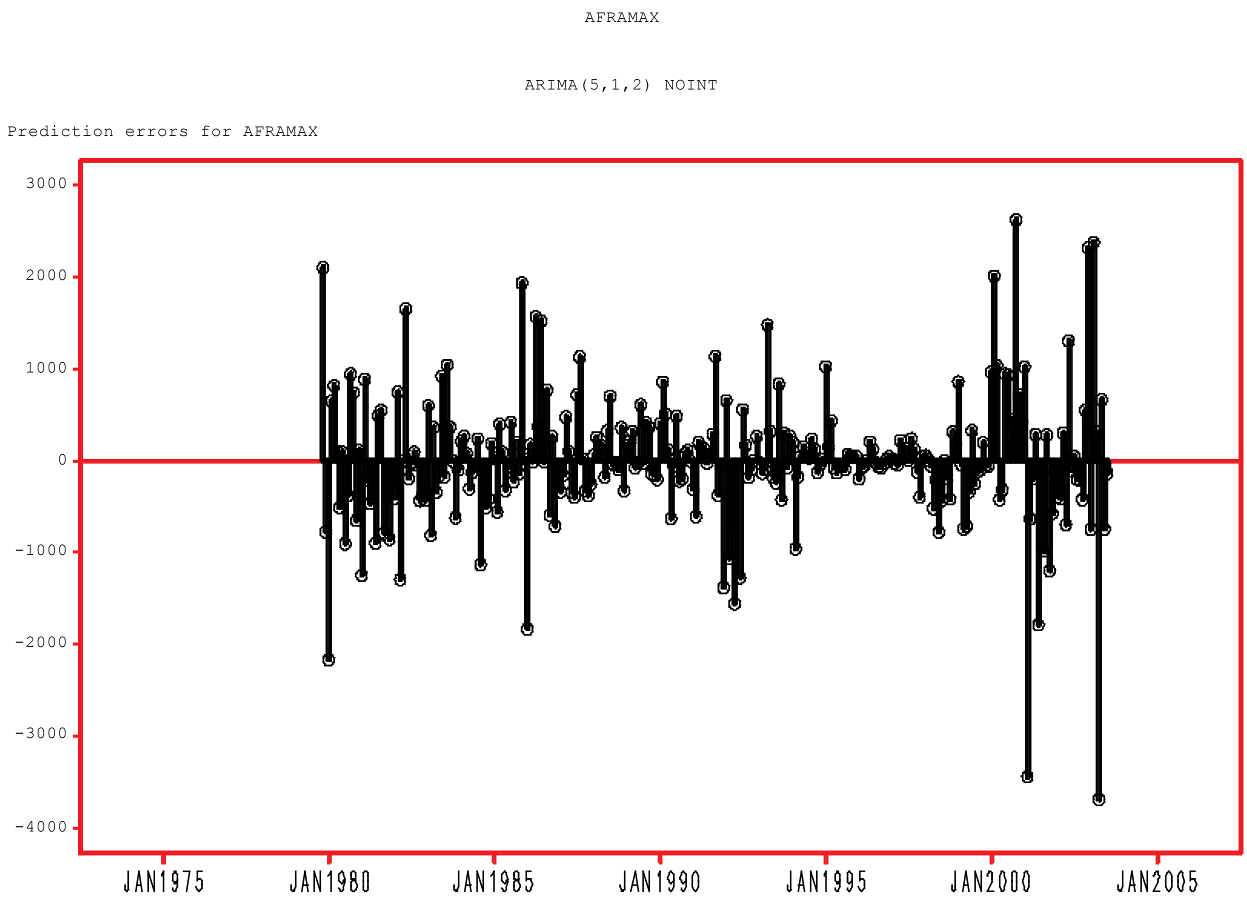

Figure 59 : Prediction Errors for AFRAMAX

Normally, the error term should be constant and as close to zero as possible. Its obvious here that the error term is not constant and that there are several instances of the model not being able to capture the time charter rate's volatility. This is expected for an ARIMA model since there are other models which are better suited to forecasting time charter rates that exhibit volatility. One of these models, GARCH, will be explored in the second part of this thesis.

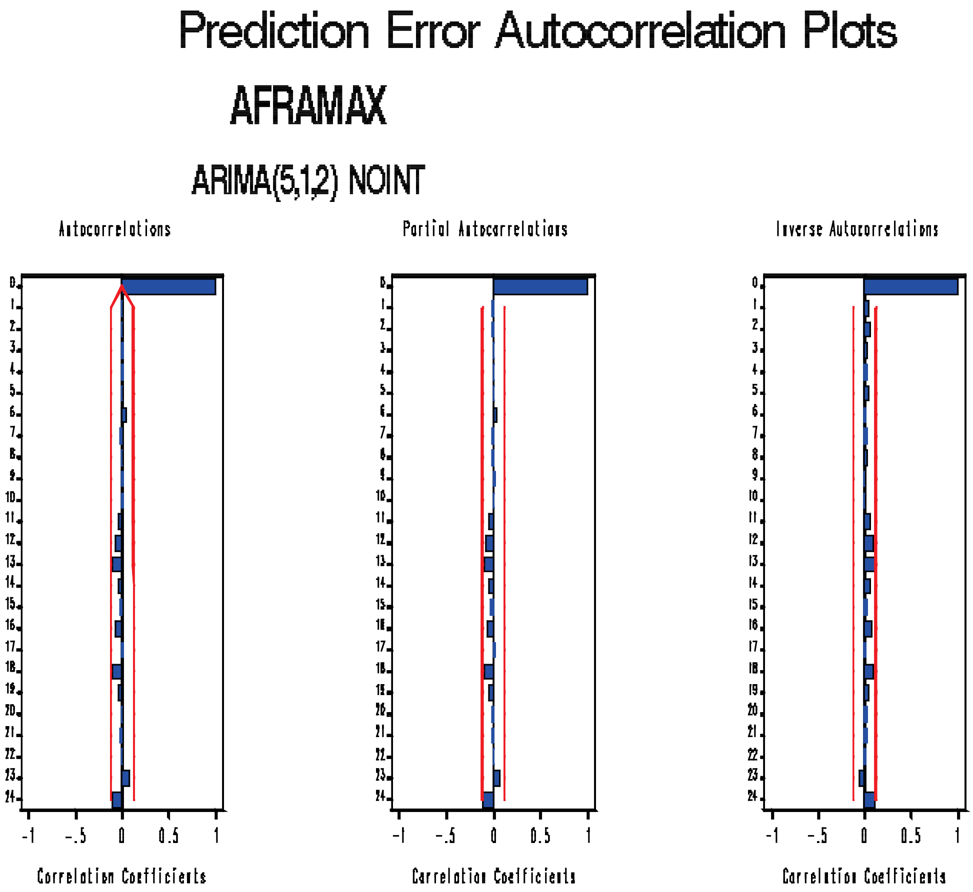

Figure 60 : Prediction Error Autocerrelation Plots for AFRAMAX

The ACF, PACF and IACF plots are within the standard errors, thus indicating that the model is correctly specified.

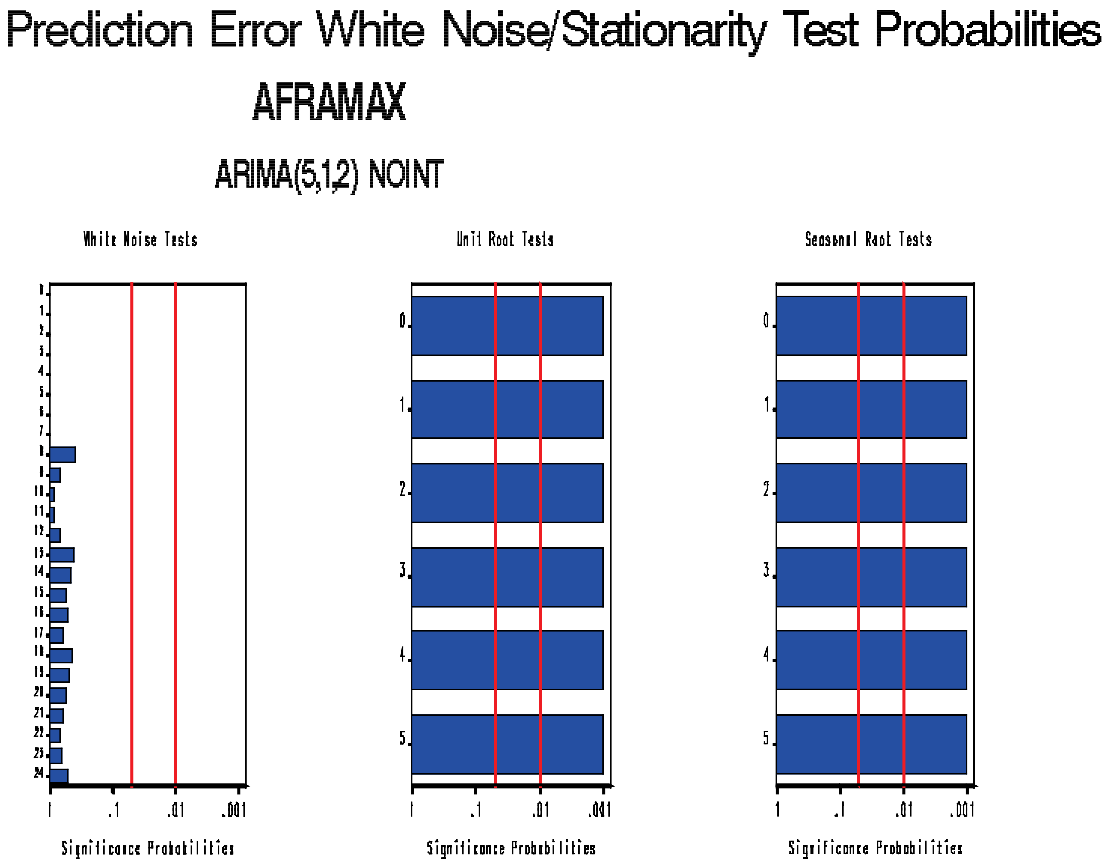

Figure 61 : Prediction Error White Noise/Stationarity Test Probabilities for AFRAMAX

The White Noise, Unit root and Seasonal Root tests indicate proper model fit for all lags. We can check the significance of each autoregressive parameter and seasonality by taking a look at the output table below.

PARM	VALUE	STDERR	T	P
Moving Average, Lag 1	0.98	0.04	23.30	0.00
Moving Average, Lag 2	-0.92	0.04	-22.03	0.00
Autoregressive, Lag 1	1.30	0.07	18.31	0.00
Autoregressive, Lag 2	-1.18	0.11	-10.87	0.00
Autoregressive, Lag 3	0.34	0.12	2.80	0.01
Autoregressive, Lag 4	-0.08	0.10	-0.75	0.46
Autoregressive, Lag 5	0.15	0.06	2.36	0.02
Model Variance (sigma squared)	487881.74			

The output shows that the 1st, 2nd 3rd and 5th order Autoregressive (AR) parameters are significant, as well as the 1st and 2nd order Moving Average (MA) parameters are significant.

Next, we can see the forecast for the AFRAMAX time series. The actual values for the next months are also shown, with upper and lower confidence limits and a standard error displayed next to the predicted value.

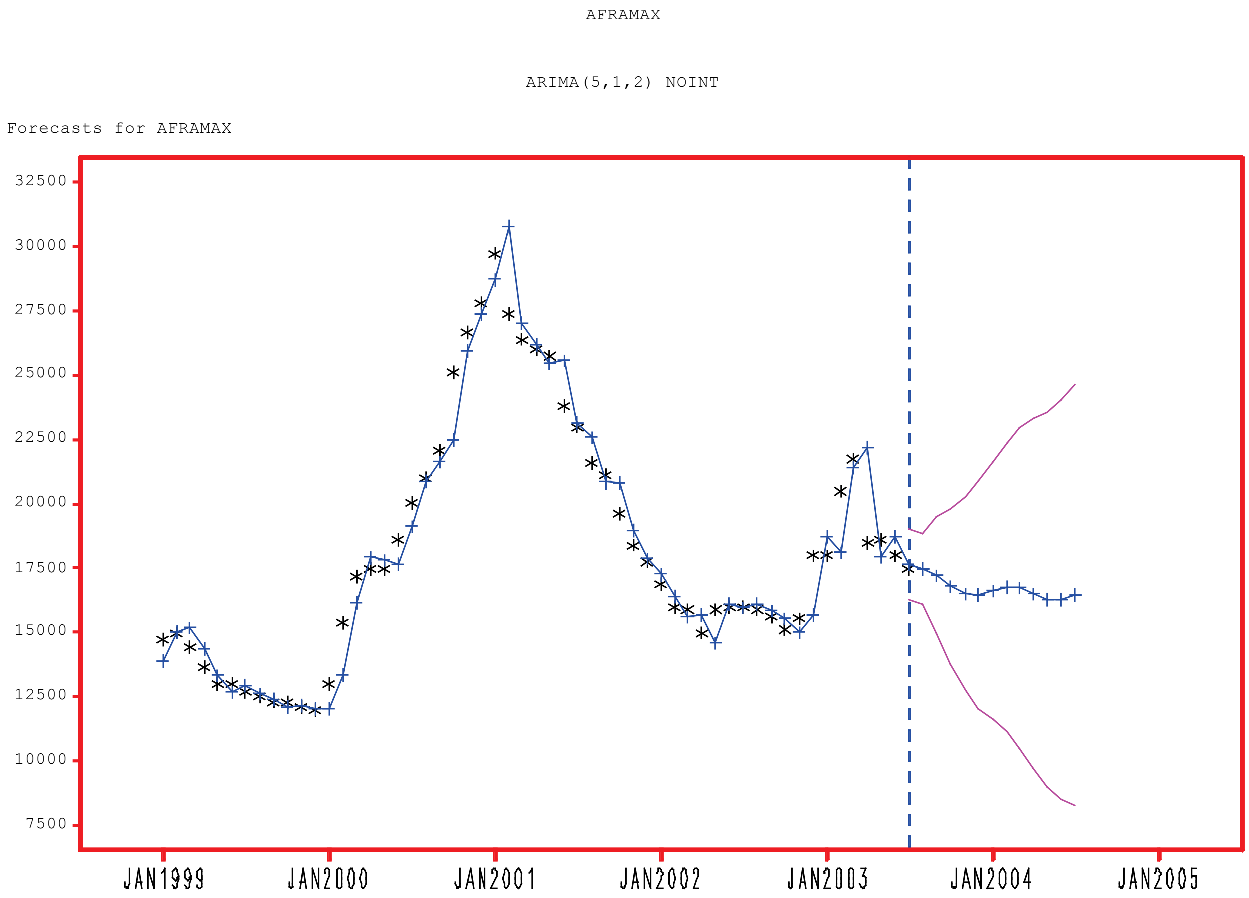

Figure 62 : Forecast for AFRAMAX

The forecasted values are shown, along with the upper and lower confidence limits (in pink). Below is the forecasted values along with the upper and lower confidence limits and the standard error.

FORECAST

DATE	PREDICT	UPPER	LOWER	STD
8/1/2003	17478.92	18847.93	16109.92	698.4853
9/1/2003	17218.79	19479.67	14957.91	1153.531
10/1/2003	16799.03	19820.67	13777.4	1541.679
11/1/2003	16502.03	20253.95	12750.11	1914.28
12/1/2003	16450.53	20854.88	12046.18	2247.159
1/1/2004	16609.37	21624.16	11594.57	2558.618
2/1/2004	16768.99	22391	11146.97	2868.428
3/1/2004	16730.83	22951.31	10510.35	3173.772
4/1/2004	16505.79	23289.76	9721.833	3461.268
5/1/2004	16292.78	23585.63	8999.93	3720.91
6/1/2004	16280.75	24032.49	8529.018	3955.04
7/1/2004	16467.54	24652.46	8282.629	4176.054

4.2 Prediction for HANDYSIZE

Below is the SAS plot of the HANDYSIZE time charter rate:

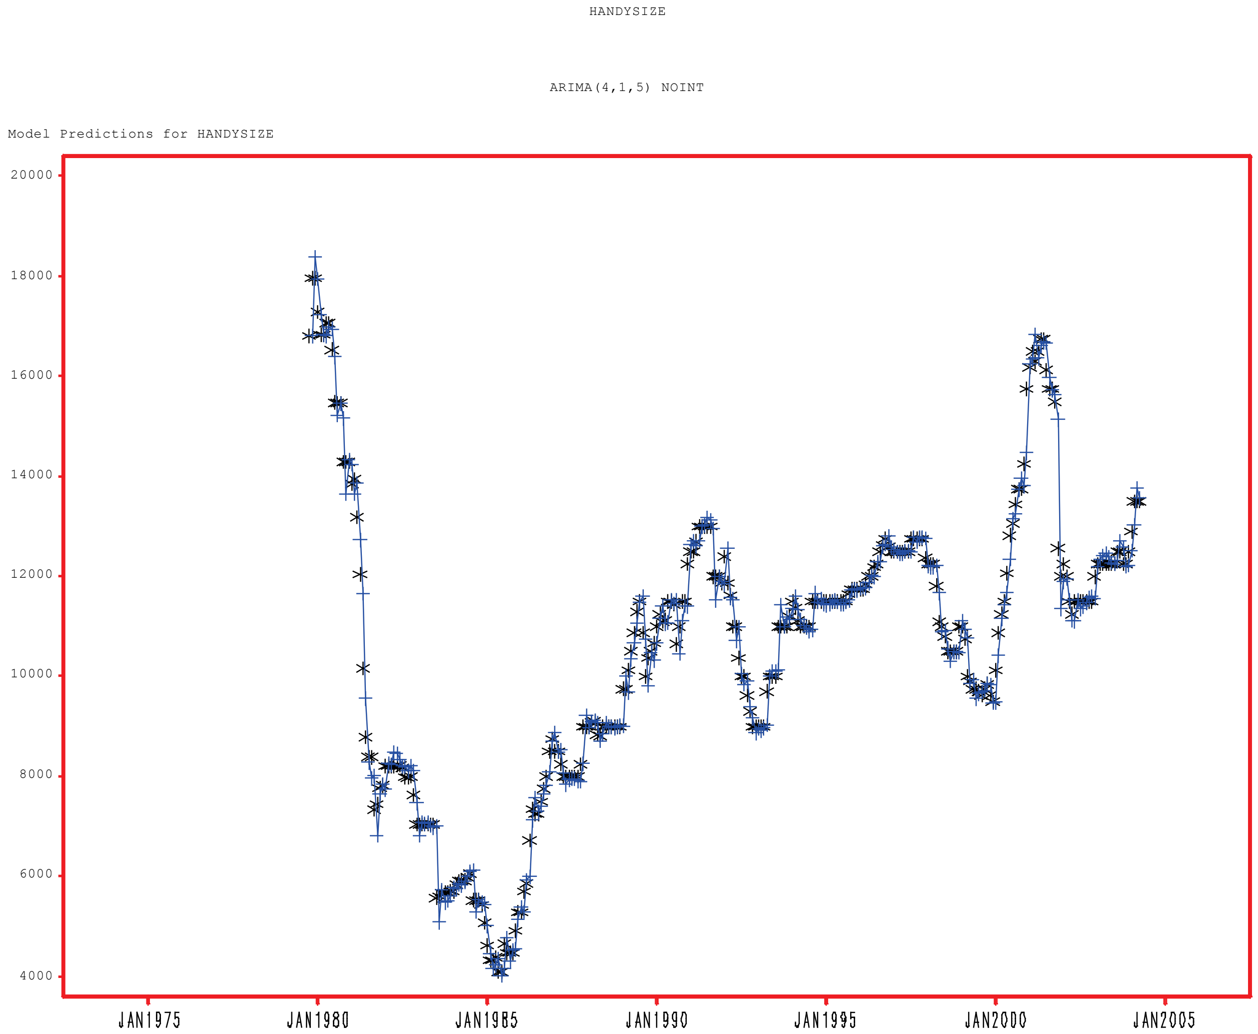

Figure 63 : Model Predictions for HANDYSIZE

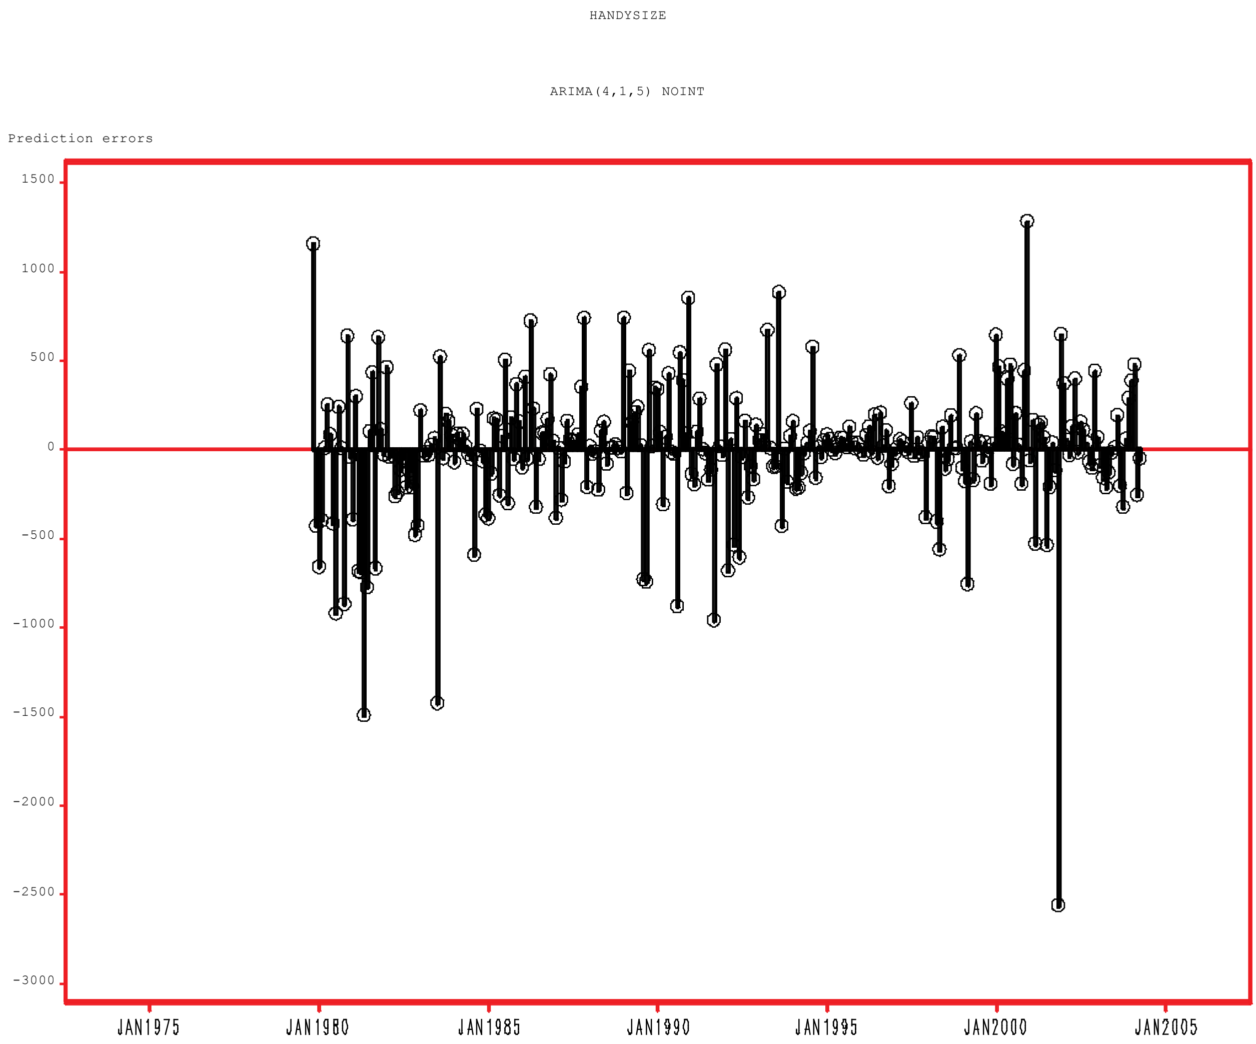

Figure 64 : Prediction Errors for HANDYSIZE

As with the previous time charter rate, the error term is not constant and there are several instances of the model not being able to capture the time charter rate's volatility. This is expected for an ARIMA model since there are other models which are better suited to forecasting time charter rates that exhibit volatility. These models will be explored later in the thesis.

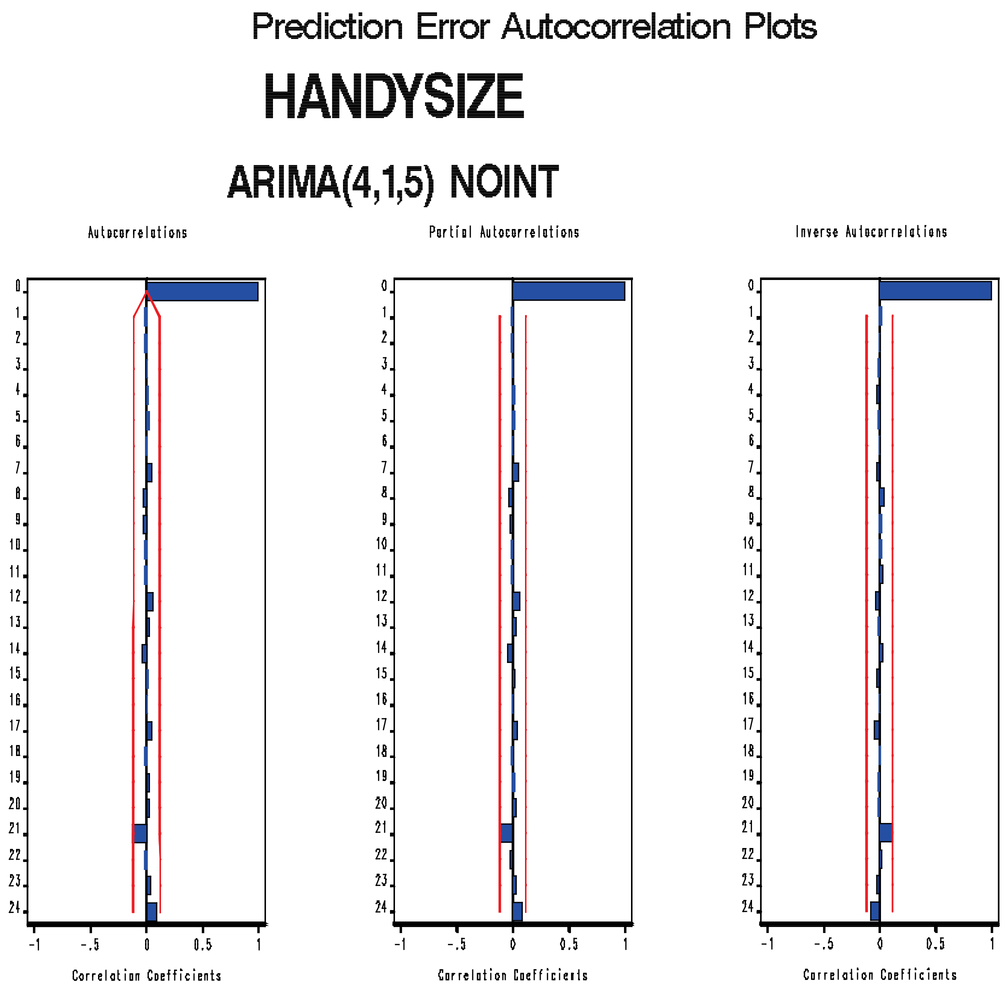

Figure 65 : Prediction Error Autocorrelation Plots for HANDYSIZE

The ACF, PACF and IACF plots are within the standard errors, thus indicating that the model is correctly specified.

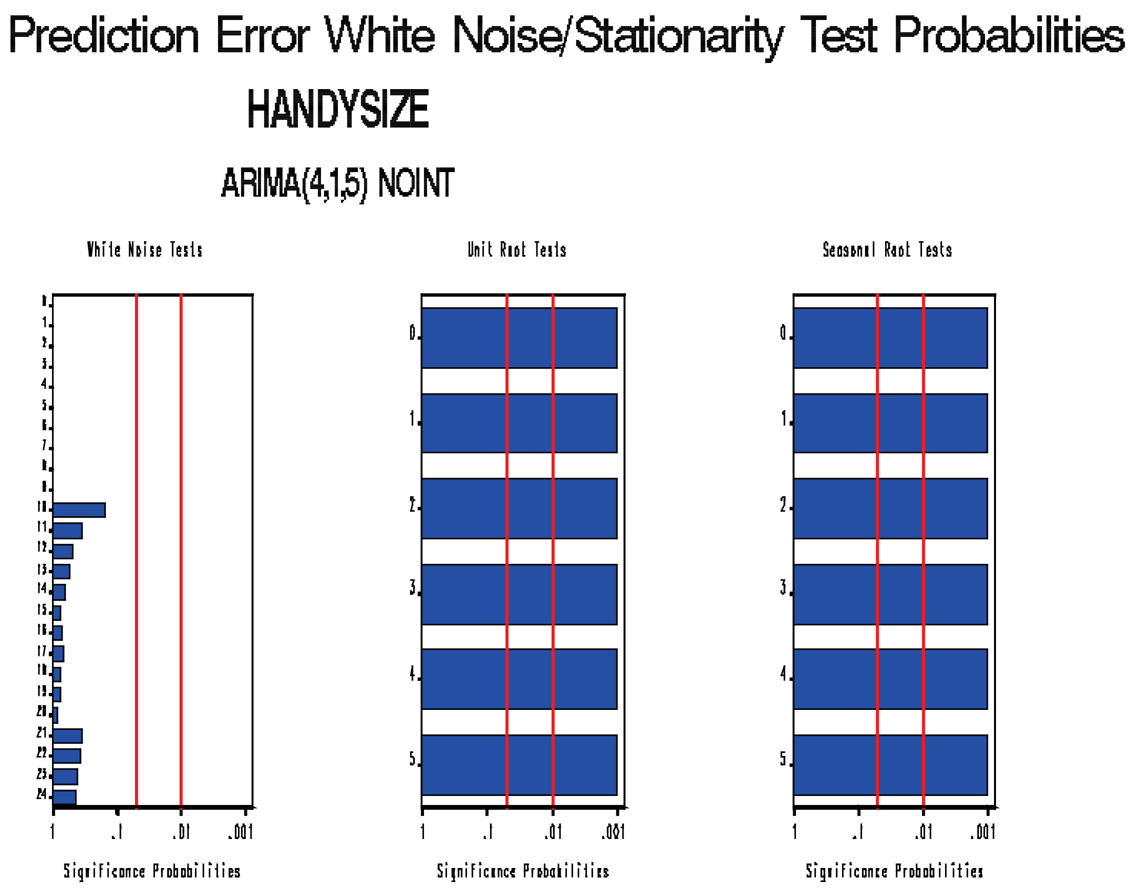

Figure 66 : Prediction Error White Noise/Stationarity Test Probabilities for HANDYSIZE

The White Noise, Unit root and Seasonal Root tests indicate proper model fit for all lags. We can check the significance of each autoregressive parameter by taking a look at the output table below.

PARM	VALUE	STDERR	T	P
Moving Average, Lag 1	0.14	0.15	0.98	0.33
Moving Average, Lag 2	-0.02	0.13	-0.13	0.90
Moving Average, Lag 3	**-0.81**	**0.06**	**-12.92**	**0.00**
Moving Average, Lag 4	**0.38**	**0.14**	**2.61**	**0.01**
Moving Average, Lag 5	0.10	0.10	1.00	0.32
Autoregressive, Lag 1	**0.52**	**0.13**	**4.00**	**0.00**
Autoregressive, Lag 2	-0.10	0.09	-1.04	0.30
Autoregressive, Lag 3	**-0.66**	**0.09**	**-7.60**	**0.00**
Autoregressive, Lag 4	**0.72**	**0.12**	**6.04**	**0.00**
Model Variance (sigma squared)	131141.37			

In this case many parameters seem to be significant in explaining the time charter rate. Next, we can see the forecast for the HANDYSIZE time series.

HANDYSIZE

ARIMA(4,1,5) NOINT

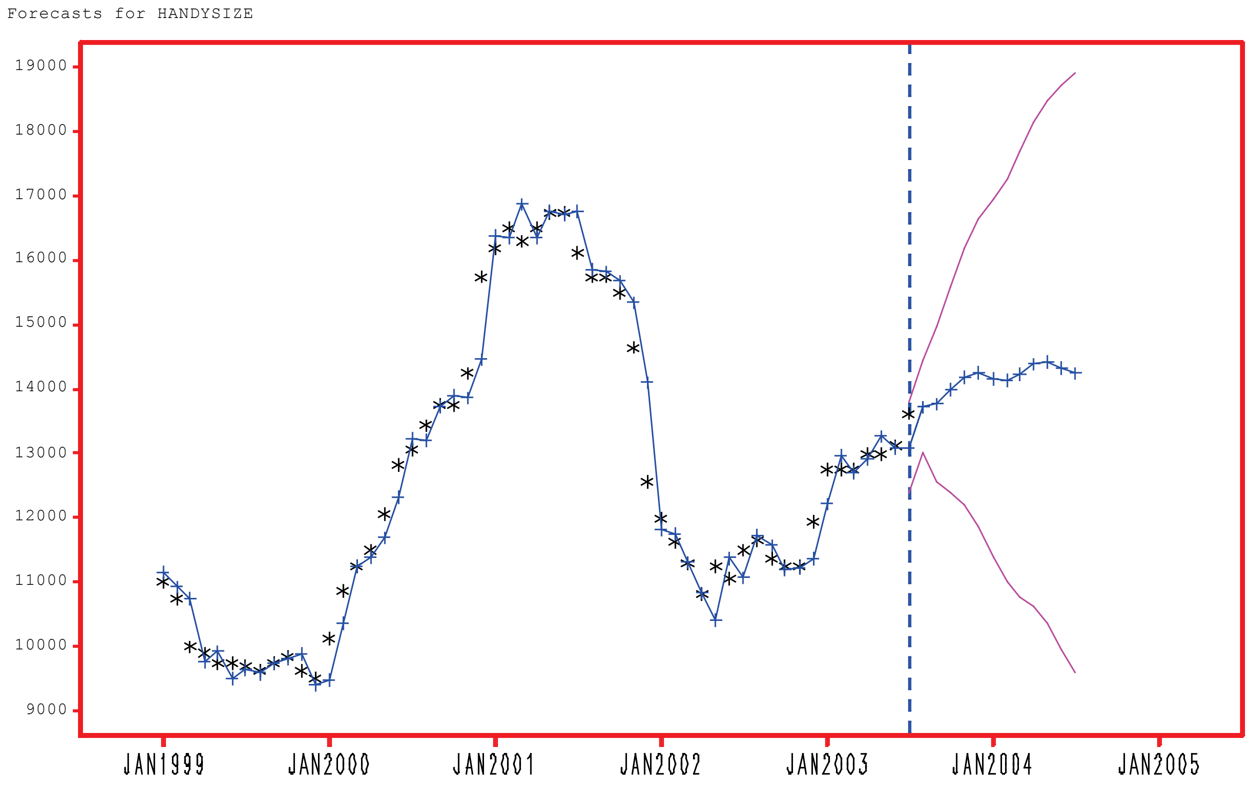

Figure 67 : Forecasts for HANDYSIZE

The forecasted values are shown, along with the upper and lower confidence limits (in pink). Below is the forecasted values along with the upper and lower confidence limits and the standard error.

FORECAST

DATE	PREDICT	UPPER	LOWER	STD
8/1/2003	13722.33	14432.1	13012.56	362.1345
9/1/2003	13771.09	14979.64	12562.53	616.6227
10/1/2003	13992.12	15600.18	12384.07	820.4508
11/1/2003	14192.44	16189.36	12195.52	1018.855
12/1/2003	14261.71	16647.23	11876.18	1217.127
1/1/2004	14166.69	16945.99	11387.38	1418.039
2/1/2004	14136.22	17273.33	10999.1	1600.597
3/1/2004	14227.37	17690.3	10764.45	1766.831
4/1/2004	14390.59	18154.98	10626.2	1920.642
5/1/2004	14418.77	18488.94	10348.61	2076.654
6/1/2004	14335.28	18707.65	9962.904	2230.844
7/1/2004	14246.17	18908.3	9584.04	2378.682

4.3 Prediction for SUEZMAX

The SUEZMAX time charter rate could not be forecasted without the help of a more complex model. This is because autocorrelation check of the residuals for a simple univariate model showed significant correlation. The addition of explanatory variables and the addition of seasonal dummy variables did not improve the model's performance. In the end, two models were created, both with low autoregressive values in order to avoid overfitting:

- ARIMA(1,1,1) with no intercept
- ARIMA(1,1,1) (2,0,0) (seasonal regression) + Crude Oil Purchase Price (Lagged 9 Periods) + Newbuilding Prices (Lagged 9 Periods) .

Finally, a third model was created as a combination of the previous two. By fitting regression weights the autocorrelation of the residuals was reduced to white noise.

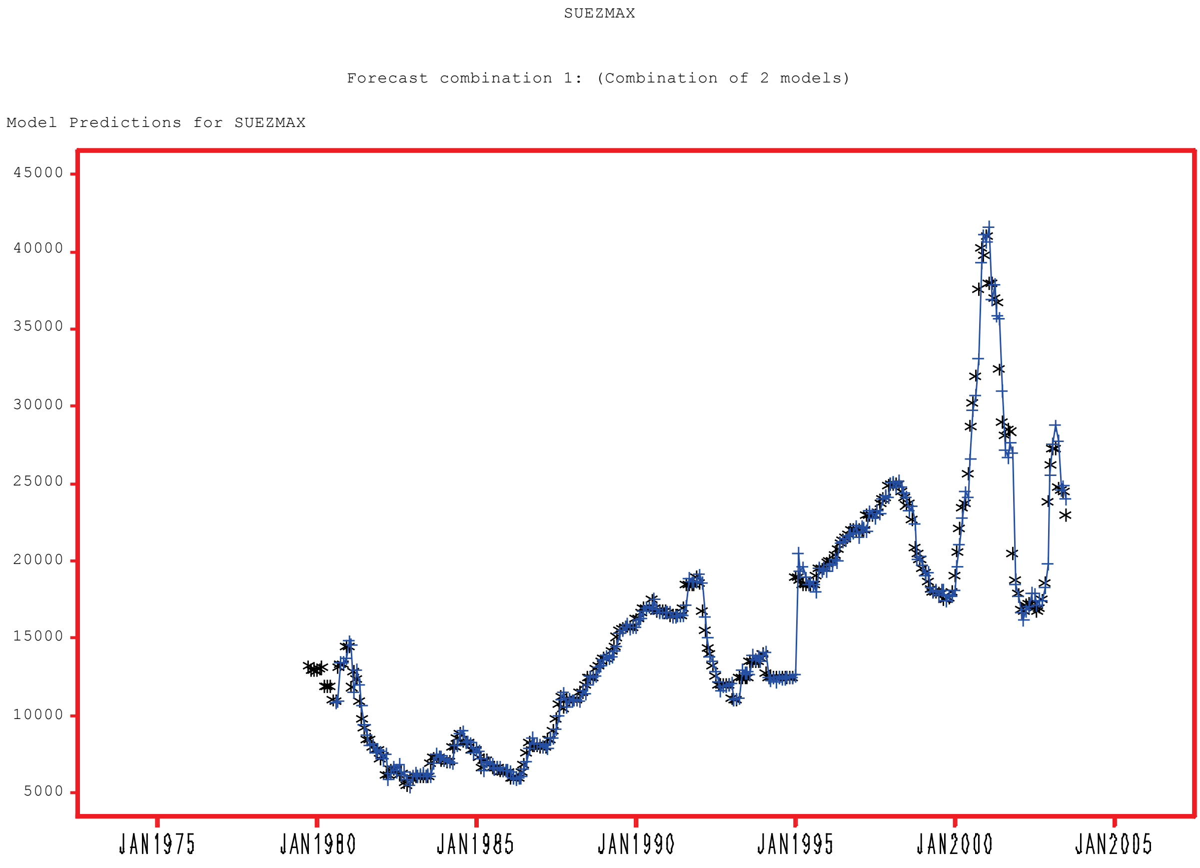

Figure 68 : Model Predictions for SUEZMAX

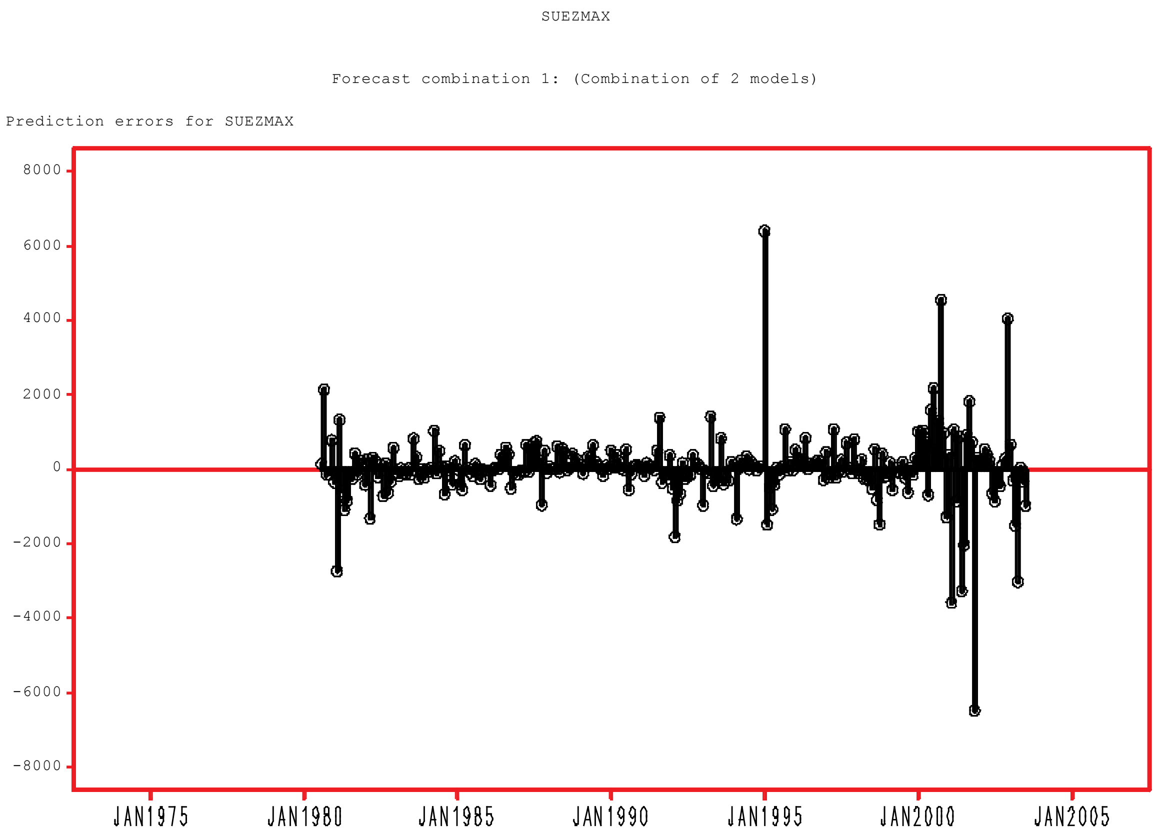

Figure 69 : Prediction Errors for SUEZMAX

Normally, the error term should be constant and as close to zero as possible. Its obvious here that the error term is not constant and that there are several instances of the model not being able to capture the time charter rate's volatility. This is expected for an ARIMA model since there are other models which are better suited to forecasting time charter rates that exhibit volatility. One of these models, GARCH, will be explored in the second part of this thesis.

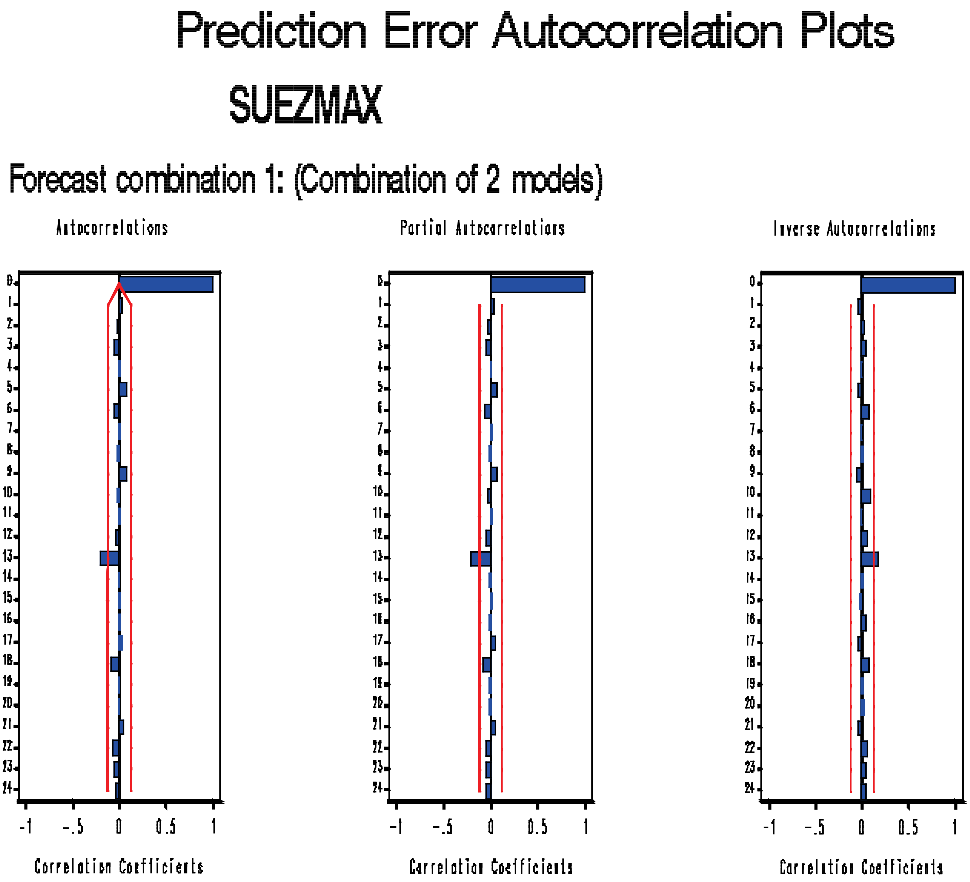

Figure 70 : Prediction Error Autocorrelation Plots for SUEZMAX

The ACF, PACF and IACF plots are within the standard errors, thus indicating that the model is correctly specified.

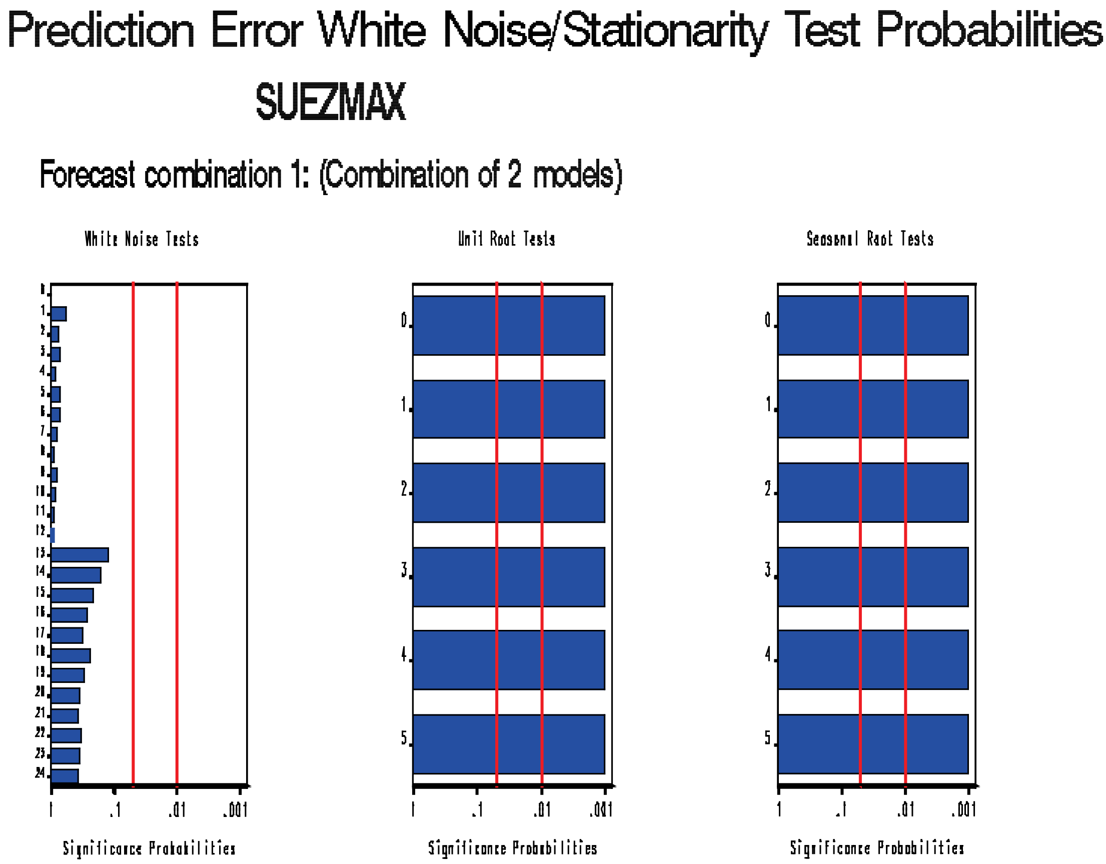

Figure 71 : Prediction Error White Noise/Stationarity Test Probabilities

The White Noise, Unit root and Seasonal Root tests indicate proper model fit for all lags. Each model contributes a standard value (shown in the table below) to the total combined model.

PARM	VALUE
copplag9 + nblag9 + ARIMA(4,1,4)	0.335830937
ARIMA(4,1,4) NOINT	0.662589288
Combined Model Variance	910254.8101

SUEZMAX

Forecast combination 1: (Combination of 2 models)

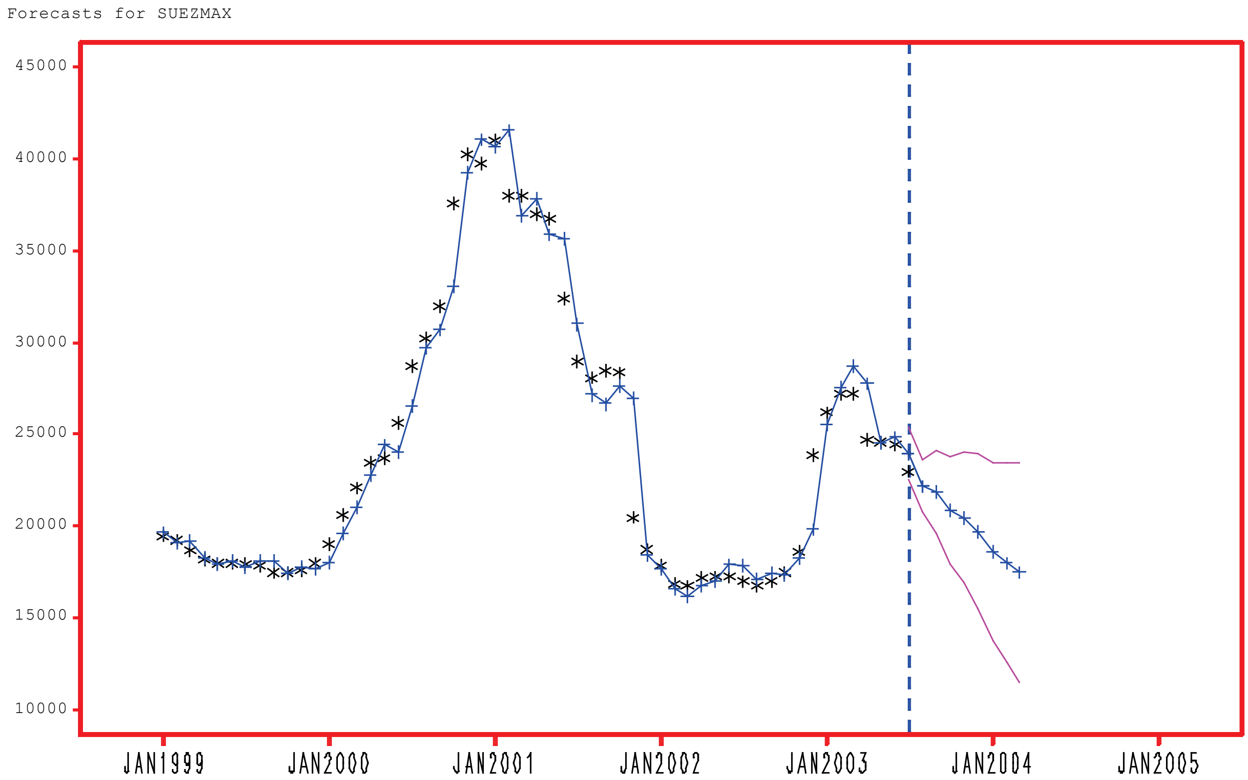

Figure 72 : Forecasts for SUEZMAX

Next, we can see the forecast for the SUEZMAX time series. The actual values for the next months are also shown, with upper and lower confidence limits and a standard error displayed next to the predicted value.

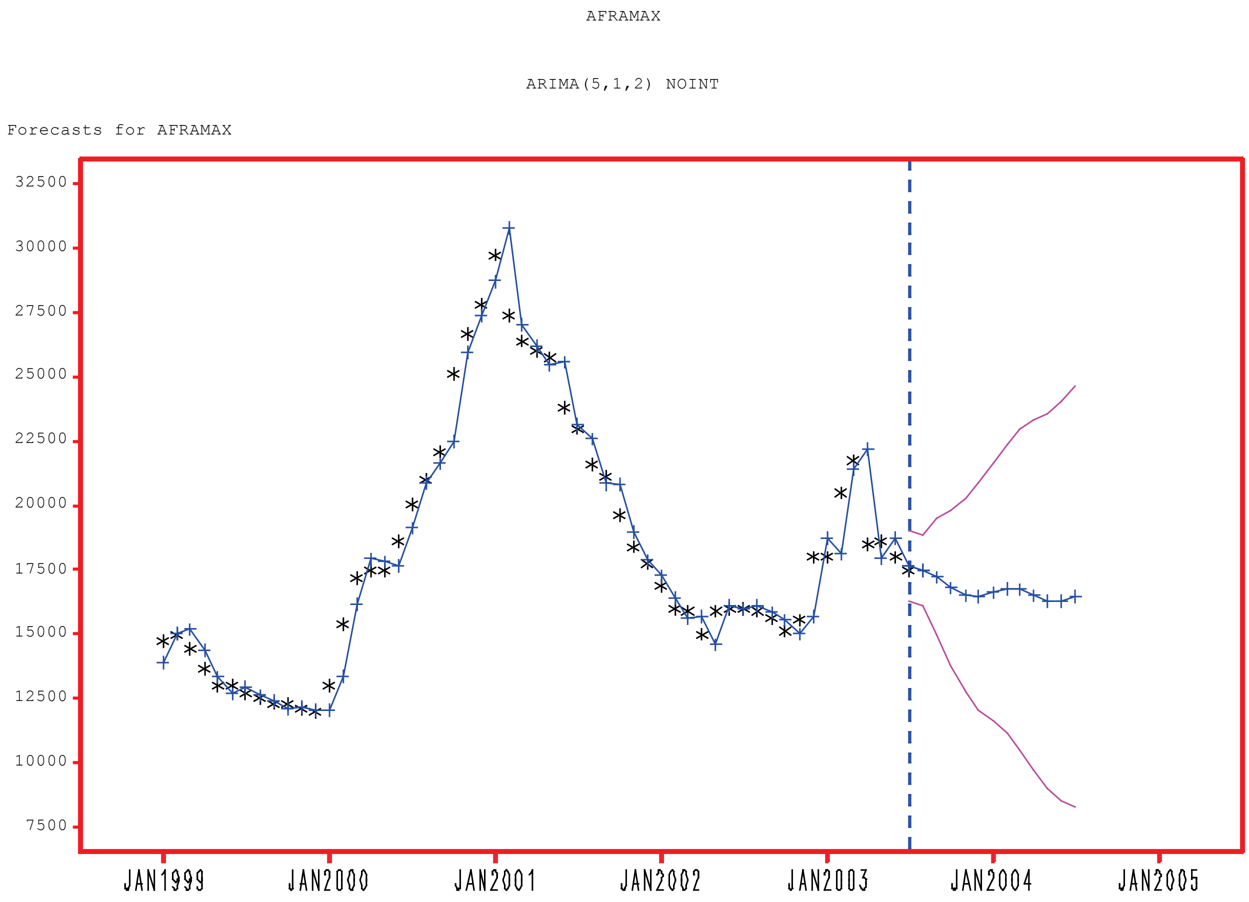

The forecasted values are shown, along with the upper and lower confidence limits (in pink). Below is the forecasted values along with the upper and lower confidence limits and the standard error.

FORECAST

DATE	PREDICT	UPPER	LOWER	STD
8/1/2003	22191.37	23594.61	20788.13	715.9538
9/1/2003	21865.08	24100.39	19629.77	1140.486
10/1/2003	20883.93	23795.29	17972.57	1485.416
11/1/2003	20466.72	24039.14	16894.31	1822.693
12/1/2003	19722.33	23965.67	15478.98	2165.012
1/1/2004	18628.96	23481.42	13776.49	2475.794
2/1/2004	17996.03	23434.35	12557.7	2774.706
3/1/2004	17487.01	23492.42	11481.61	3064.039

4.4 Prediction for VLCC

Below is the SAS plot of the VLCC time charter rate:

Figure 73 : Model Predictions for VLCC

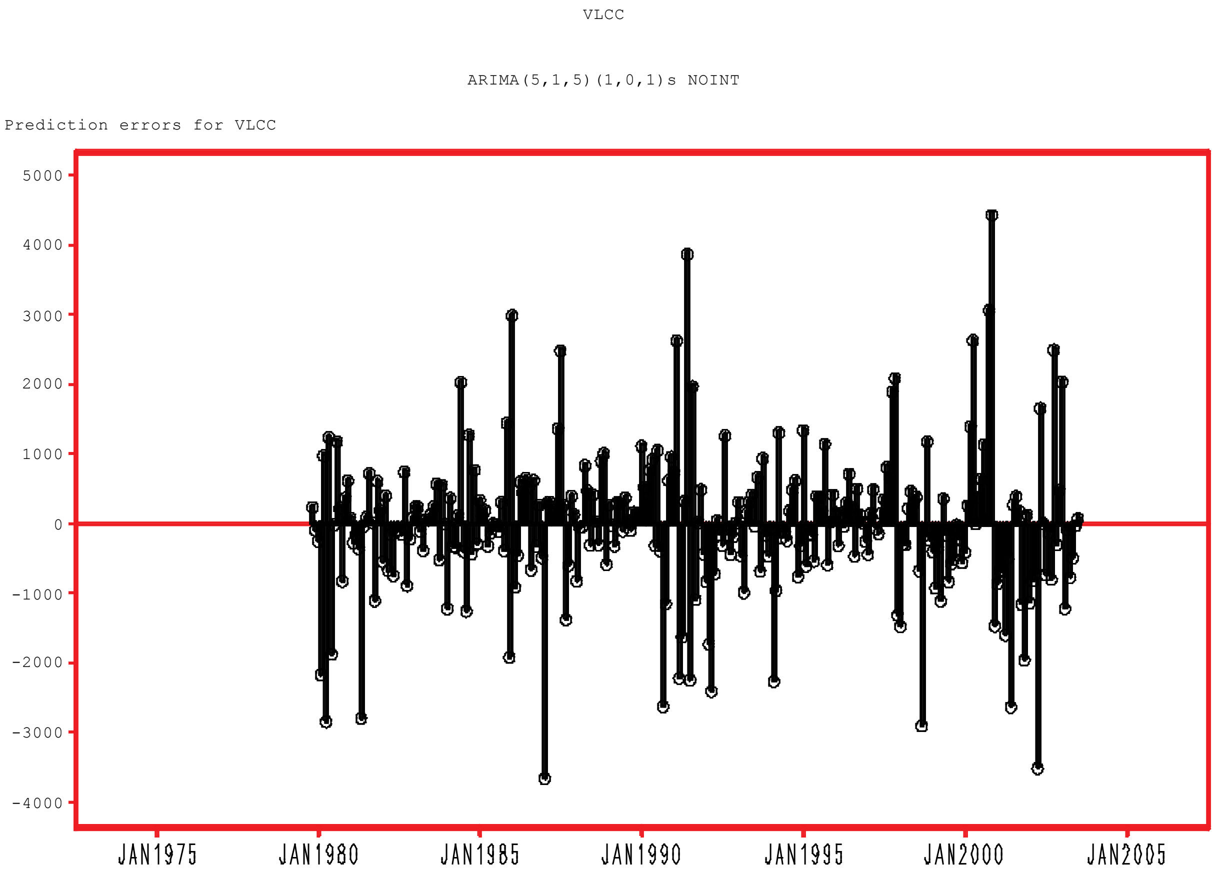

Figure 74 : Prediction Errors for VLCC

Normally, the error term should be constant and as close to zero as possible. Its obvious here that the error term is not constant and that there are several instances of the model not being able to capture the time charter rate's volatility. This is expected for an ARIMA model since there are other models which are better suited to forecasting time charter rates that exhibit volatility. One of these models, GARCH, will be explored in the second part of this thesis

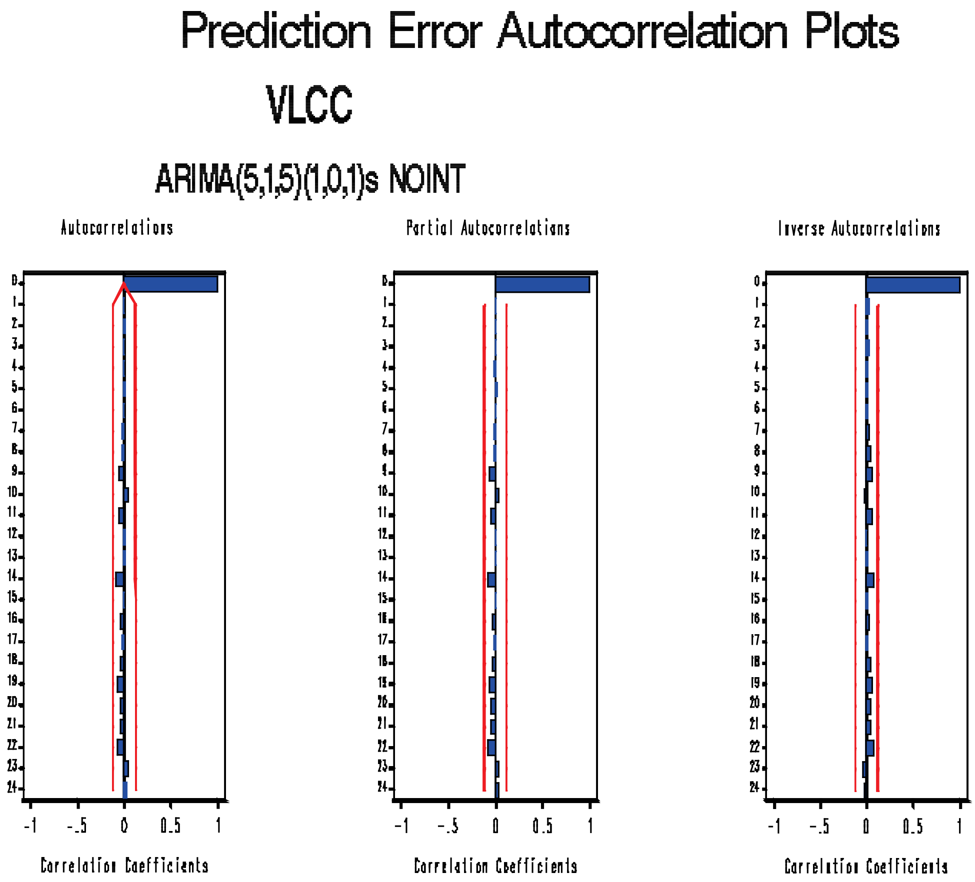

Figure 75 : Prediction Error Autocorrelation Plots for VLCC

The ACF, PACF and IACF plots are within the standard errors, thus indicating that the model is correctly specified.

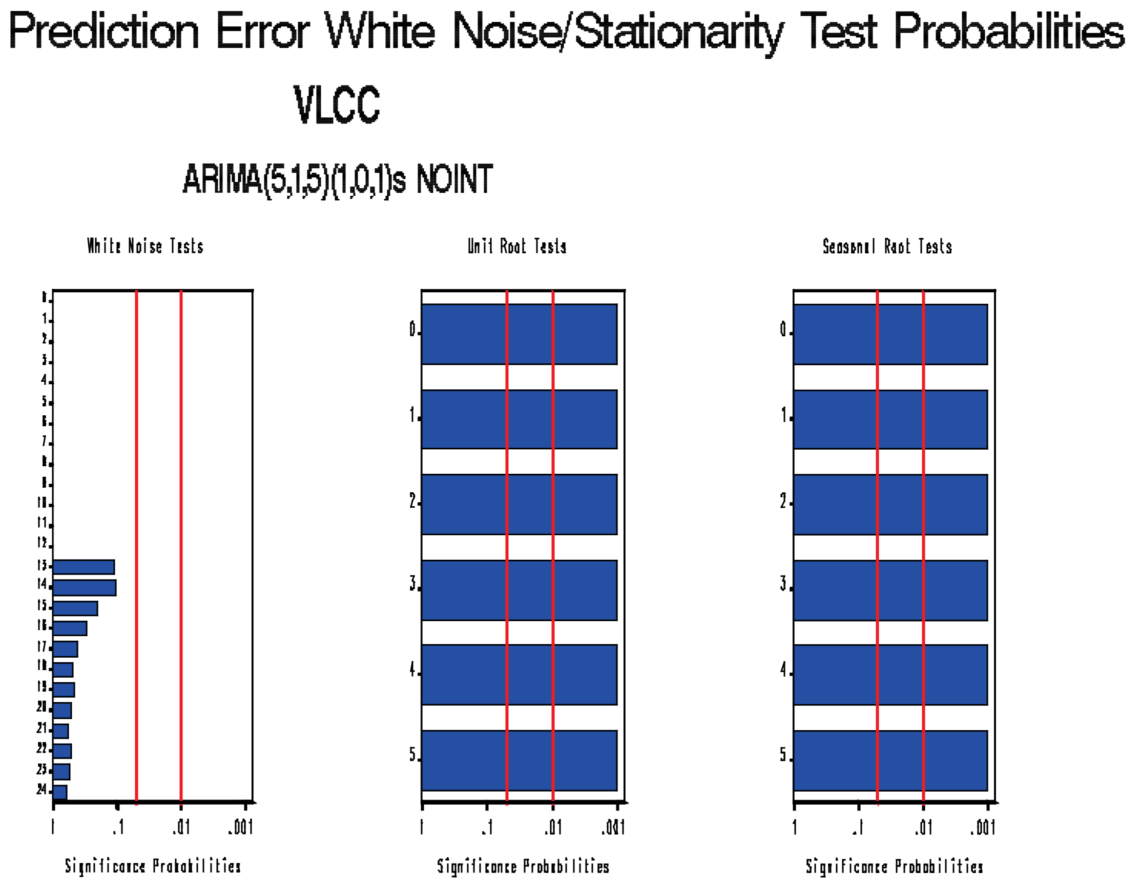

Figure 76 : Prediction Error White Noise/Stationarity Test Probabilities for VLCC

The White Noise, Unit root and Seasonal Root tests indicate proper model fit for all lags. We can check the significance of each autoregressive parameter and seasonality by taking a look at the output table below.

PARM	VALUE	STDERR	T	P
Moving Average, Lag 1	0.09	0.35	0.25	0.80
Moving Average, Lag 2	**-1.14**	**0.19**	**-6.18**	**0.00**
Moving Average, Lag 3	0.13	0.42	0.31	0.76
Moving Average, Lag 4	**-0.67**	**0.21**	**-3.24**	**0.00**
Moving Average, Lag 5	0.33	0.27	1.23	0.22
Seasonal Moving Average, Lag 12	-0.10	0.31	-0.31	0.76
Autoregressive, Lag 1	0.44	0.33	1.30	0.19
Autoregressive, Lag 2	**-1.19**	**0.15**	**-7.82**	**0.00**
Autoregressive, Lag 3	0.49	0.39	1.28	0.20
Autoregressive, Lag 4	**-0.66**	**0.17**	**-3.84**	**0.00**
Autoregressive, Lag 5	**0.45**	**0.21**	**2.12**	**0.04**
Seasonal Autoregressive, Lag 12	-0.32	0.30	-1.06	0.29
Model Variance (sigma squared)	1082147.49			

Next, we can see the forecast for the VLCC time series. The actual values for the next months are also shown, with upper and lower confidence limits and a standard error displayed next to the predicted value

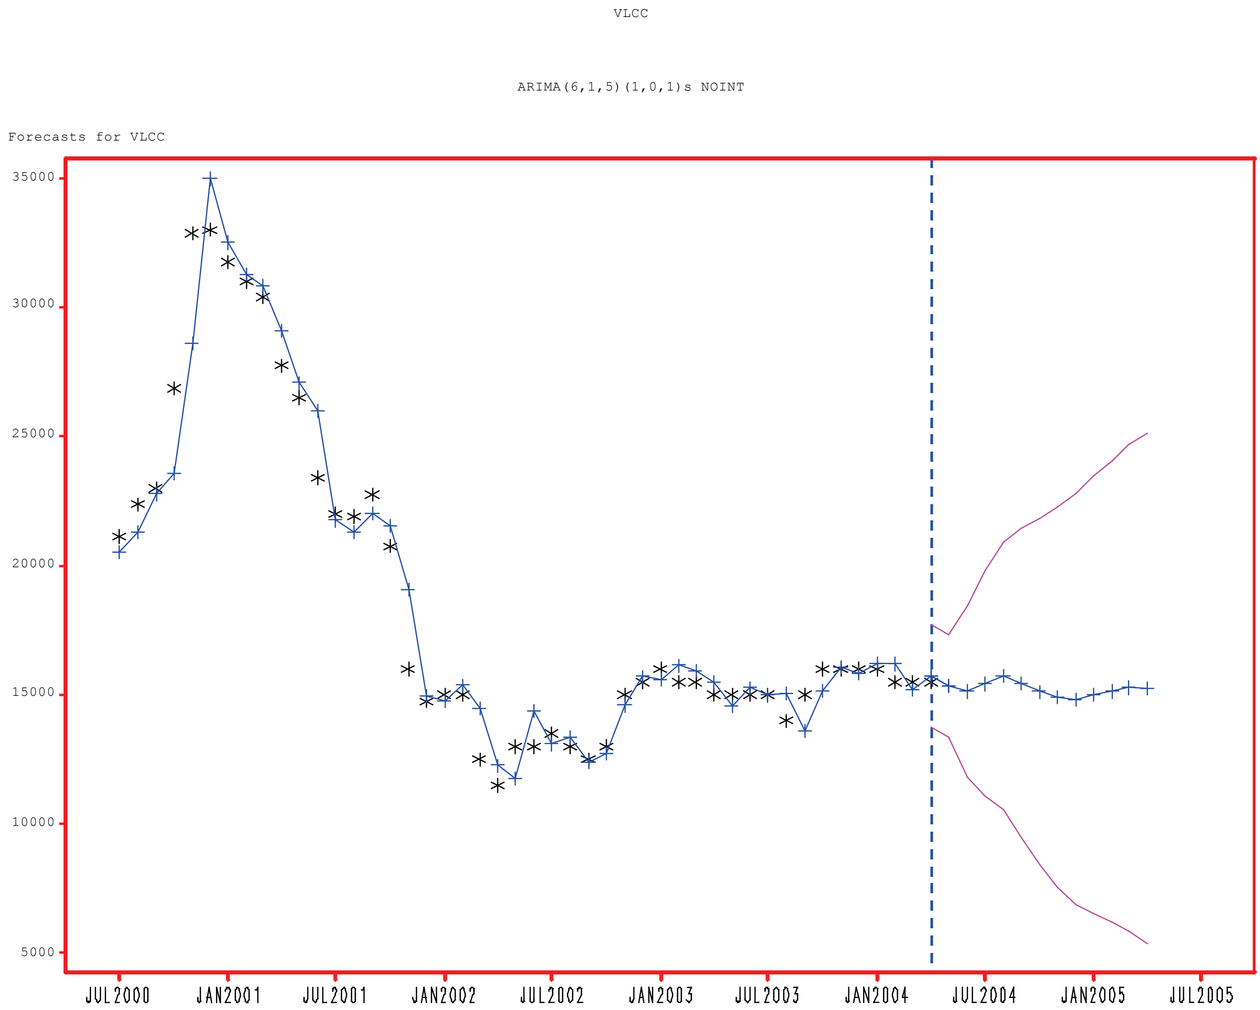

Figure 77 : Forecasts for VLCC

The forecasted values are shown, along with the upper and lower confidence limits (in pink). Below is the forecasted values along with the upper and lower confidence limits and the standard error.

FORECAST

DATE	PREDICT	UPPER	LOWER	STD
8/1/2003	15573.83	17612.71	13534.95	1040.263
9/1/2003	15293.91	18715.76	11872.05	1745.878
10/1/2003	15058.71	19589.22	10528.2	2311.527
11/1/2003	15411.4	20828.16	9994.633	2763.706
12/1/2003	15035.4	21256.34	8814.468	3174.005
1/1/2004	13826.63	20723.61	6929.653	3518.932
2/1/2004	13630.71	21155.65	6105.764	3839.326
3/1/2004	13960.51	22162.88	5758.134	4184.96
4/1/2004	14454.39	23272.59	5636.186	4499.166
5/1/2004	14682.37	24007.85	5356.88	4757.989
6/1/2004	14276.32	24086.34	4466.296	5005.204
7/1/2004	13867.13	24162.46	3571.791	5252.818

5. ARIMAX Model

As shown from the autocorrelation plots, two time series can be correlated with each other at different time periods. For instance, lets see again the returns for the VLCC time charter rate when compared to the Crude Oil Purchase Price, Newbuilding Price and HANDYSIZE TCR Returns.

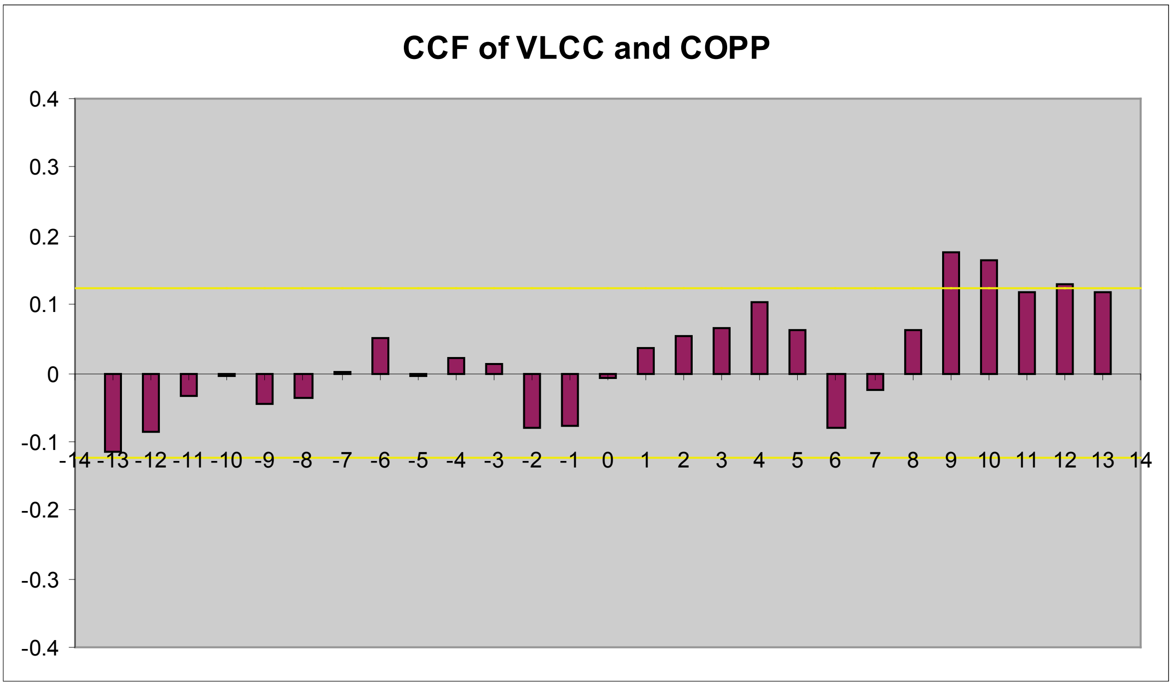

Figure 78 : CCF of VLCC and COPP

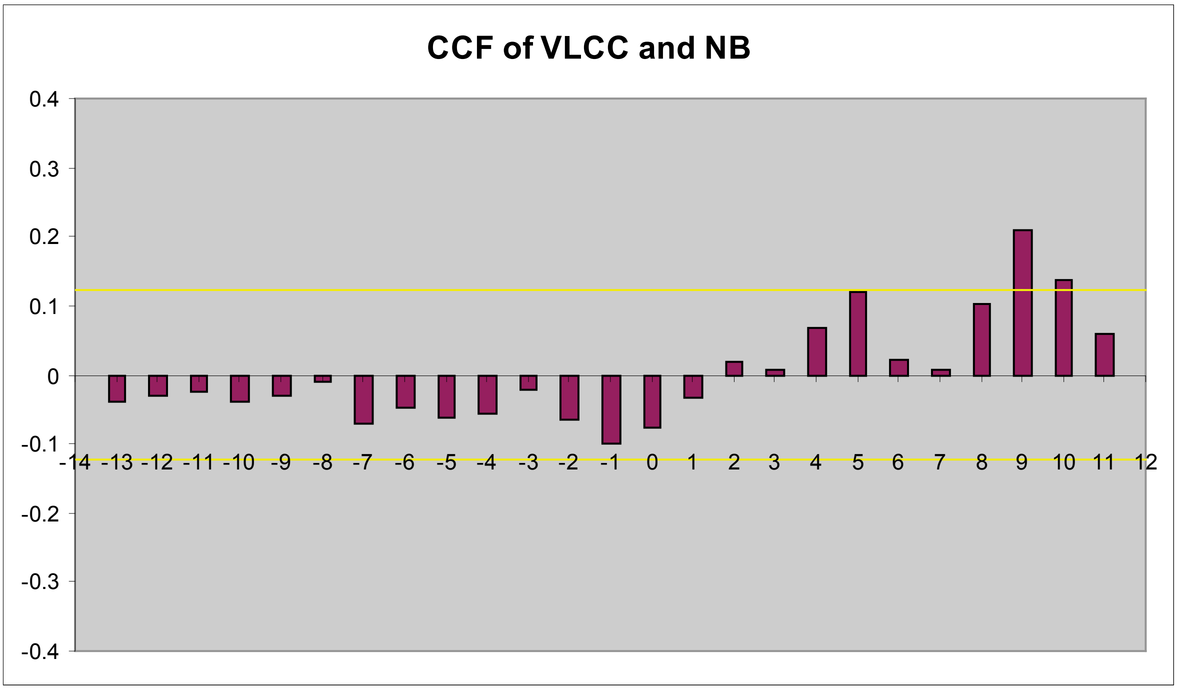

Figure 79 : CCF of VLCC Newbuildnig Prices

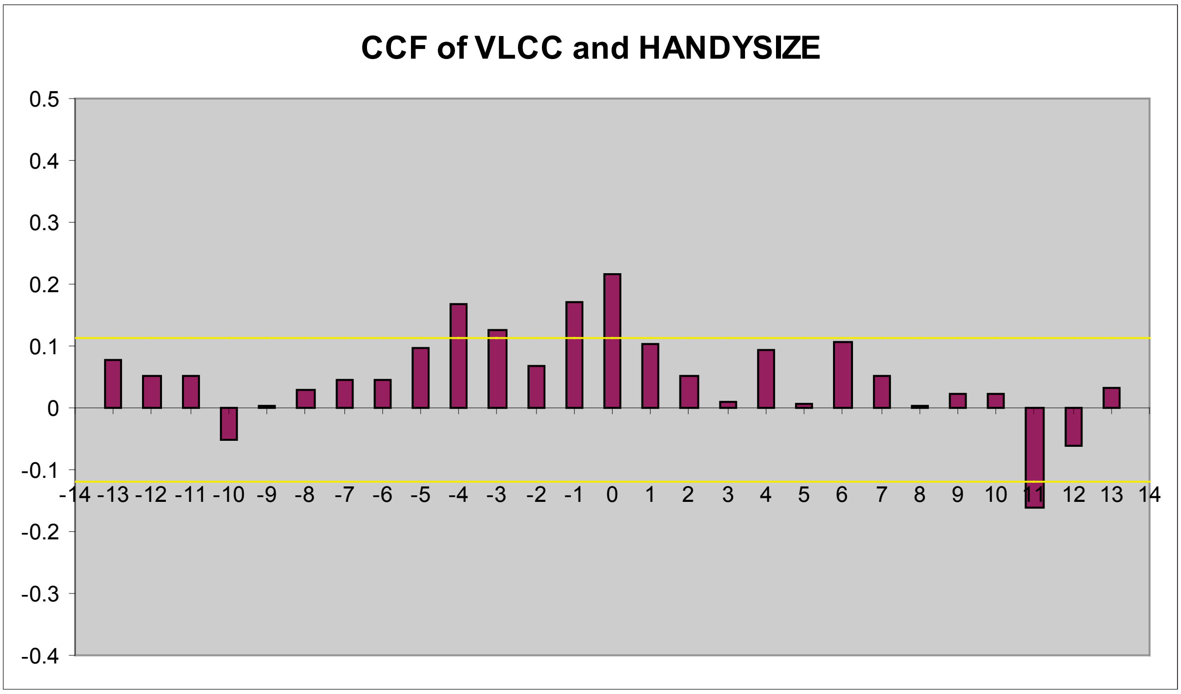

Figure 80 : CCF of VLCC and HANDYSIZE Time Charter Rates

One task in economic modeling is that of describing the possible impacts over time of a change in one or more explanatory variables on a dependent variable. One important feature of the ARMAX model is its ability to model any possible lag distribution shape. Depending on the explanatory variables used, the Cross Correlation function may show a maximum correlation at a lag other than zero. In our case, all three time charter rates have a maximum CCF at lag 0. For long horizon forecasts, a case could be made for the HANDYSIZE TCR which shows a spike in its CCF at lag t= + 11. Of the alternate explanatory variables used, only the NewBuilding and VLCC Scrap Price have a long enough forecast horizon with a good CCF lag at t = + 9.

In order to create a large horizon forecast without forecasting the explanatory variables, the following three explanatory variables will be used in an ARIMAX model:

- HANDYSIZE Time Charter Rate (t=+11)
- Crude Oil Purchase Price (t=+9)
- NewBuilding Prices (t=+9)

It would be logical to conclude that the ARIMAX model would be specified with explanatory variables with a lag equal to the lag specified by the cross correlation plot. The basic ARMAX model is:

$$\begin{aligned} z_t = {} & \phi_1 z_{t-1} + \phi_2 z_{t-2} + \ldots + \phi_p z_{t-p} + \\ & + \beta_0 x_{1t} + \beta_1 x_{1t-1} + \ldots + \beta_v x_{1t-p} + \\ & + \gamma_0 x_{2t} + \gamma_1 x_{2t-1} + \ldots + \gamma_2 x_{2t-p} + \\ & .... \\ & + a_t - \theta_1 a_{t-1} - \theta_2 a_{t-2} - \ldots - \theta_q a_{t-q} \end{aligned}$$

Where $\phi_1 z_{t-1} + \phi_2 z_{t-2} + \ldots + \phi_p z_{t-p}$ is the autoregressive parameter, $\beta_0 x_{1t} + \beta_1 x_{1t-1} + \ldots + \beta_v x_{1t-p} + \gamma_0 x_{2t} + \gamma_1 x_{2t-1} + \ldots + \gamma_2 x_{2t-p} + ...$ are the explanatory variables $x_1, x_2 \ldots$ that are regressed up to AR(p) and finally, $a_t - \theta_1 a_{t-1} - \theta_2 a_{t-2} - \ldots - \theta_q a_{t-q}$ is the moving average term.

The ARIMAX model developed below uses explanatory variables as well as dummy seasonal variables to adjust for the seasonality that appears in the 12^{th} order lag of the VLCC autocorrelations. The fact that the 12^{th} order lag doesn't also appear at lags 24, 36 etc means that the seasonality is stationary or because of the effect of an outlier variable.

5.1 Seasonal Dummy Variables

We have assumed that all independent variables (AFRAMAX, SUEZMAX, HANDYSIZE) are quantitative and measured in a well defined scale (time). Frequently however, variables are qualitative and have many distinct levels. For instance, consider the effect of sex on the starting salary of Engineering School graduates. Or consider predicting the speed of adoption of an innovation in terms of the size of the firm and type of ownership (public, private). In these examples we observe a qualitative variable at several different levels. In order to model its effect, we have to introduce additional variables.

The effect of a qualitative variable that is observed at k different levels (for instance from k different months) on the response variable z(t) can be represented by k-1 indicator variables. These indicator or *dummy variables* are defined as IND_{ti} = 1 if the observation comes from level i (for $1 \le i \le k-1$), and 0 otherwise.

In our case, there are k = 12 indicator variables (seasonal dummy variables) and they are defined as SD_{ti} = 1 when the data comes from i=1 (January) and 0 otherwise. The same process is followed for i = 2 (February), up to i = k-1 = 11 (November).

$$\sum_{i=1}^{k-1} \delta_i SD_{ti} \xrightarrow{k=12} \delta_1 SD_{t1} + \delta_2 SD_{t2} + \ldots + \delta_{11} SD_{t11}$$

As a result, the basic ARMAX model with seasonal dummy variables (monthly) is

$$
\begin{aligned}
z_t = {} & \phi_1 z_{t-1} + \phi_2 z_{t-2} + \ldots + \phi_p z_{t-p} + \\
& + \beta_0 x_{1t} + \beta_1 x_{1t-1} + \ldots + \beta_v x_{1t-p} + \\
& + \gamma_0 x_{2t} + \gamma_1 x_{2t-1} + \ldots + \gamma_2 x_{2t-p} + \\
& \ldots. \\
& + \delta_1 SD_{t1} + \delta_2 SD_{t2} + \ldots + \delta_{11} SD_{t11} + \\
& + a_t - \theta_1 a_{t-1} - \theta_2 a_{t-2} - \ldots - \theta_q a_{t-q}
\end{aligned}
$$

5.2 ARIMAX Model

The ARIMAX model that has been selected includes the lagged explanatory variables. Seasonal dummies were initially included.

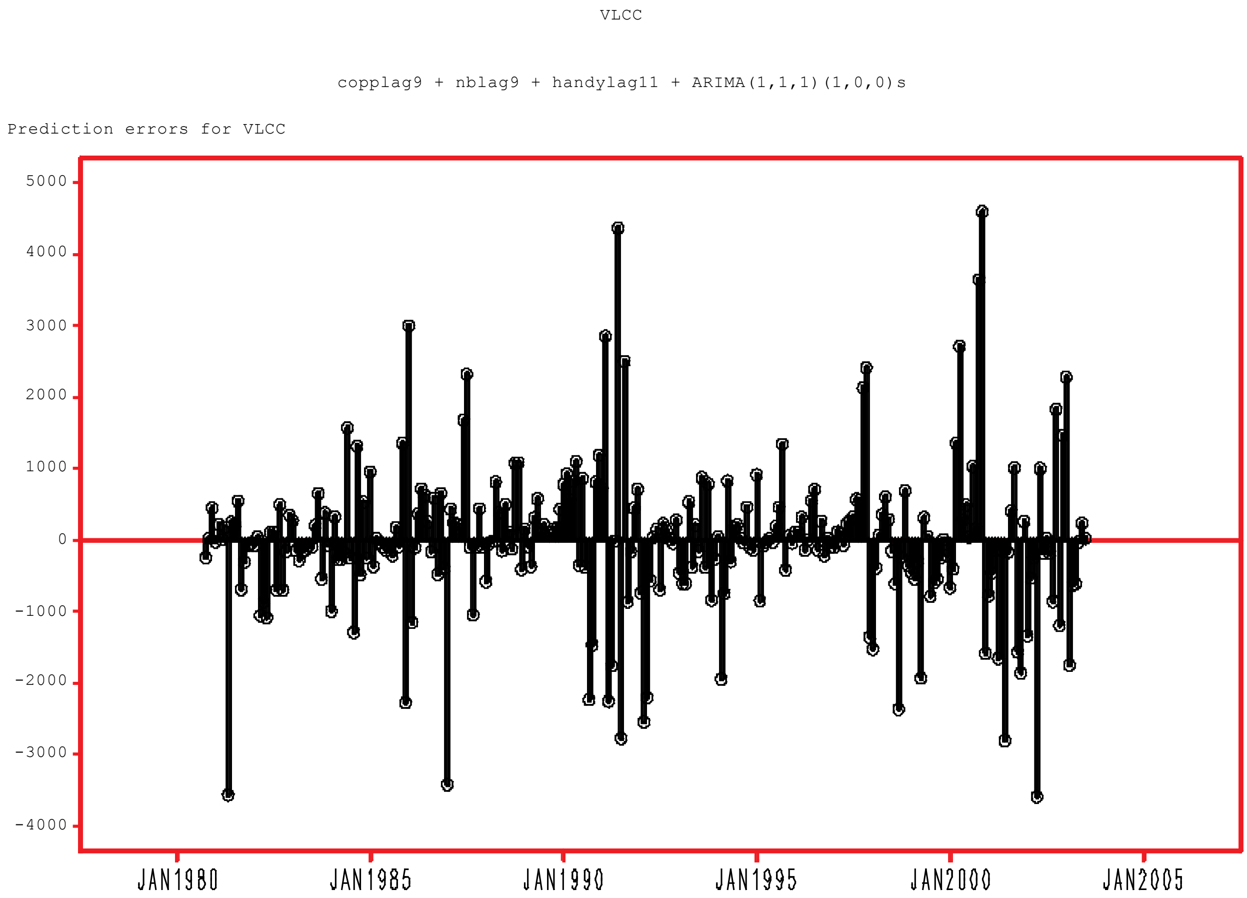

Figure 81 : Prediction Errors for VLCC (ARMAX)

Normally, the error term should be constant and as close to zero as possible. Its obvious here that the error term is not constant and that there are several instances of the model not being able to capture the time charter rate's volatility. This is expected for an ARIMA model since there are other models which are better suited to forecasting time charter rates that exhibit volatility. One of these models, GARCH, will be explored in the second part of this thesis.

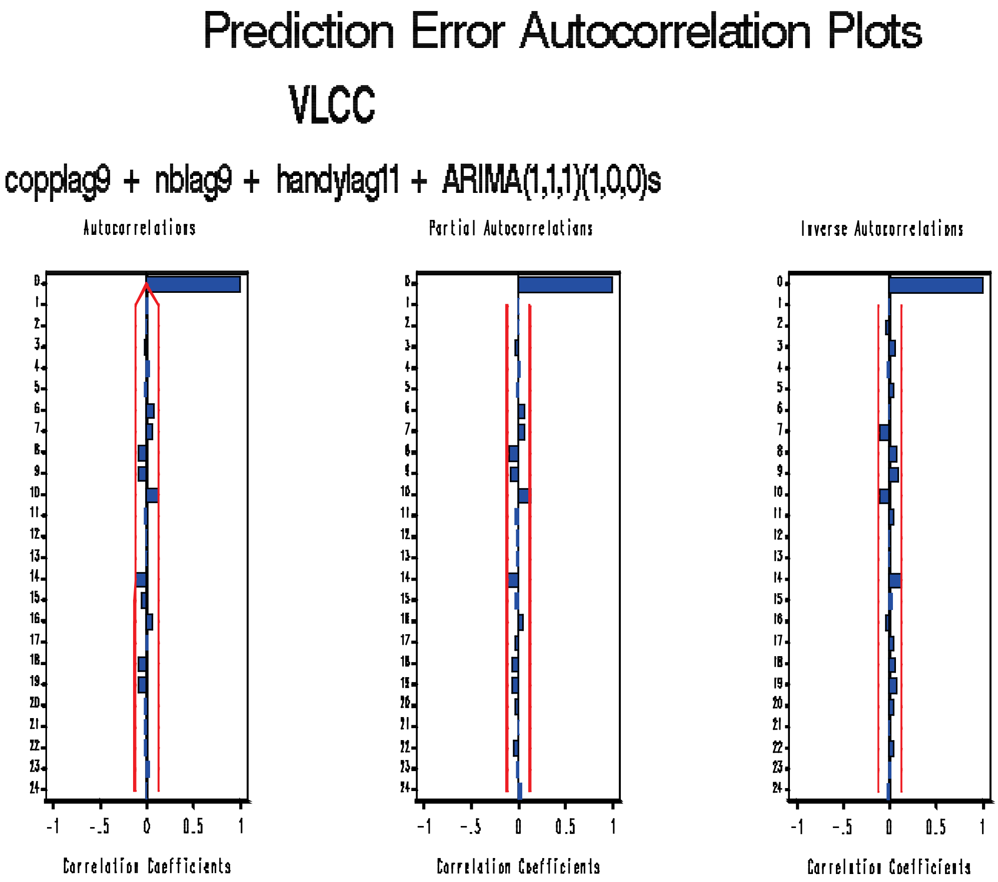

Figure 82 : Prediction Error Autocorrelation Plots for VLCC (ARMAX)

We see that the ACE, PACF and IACF indicate a stationary series.

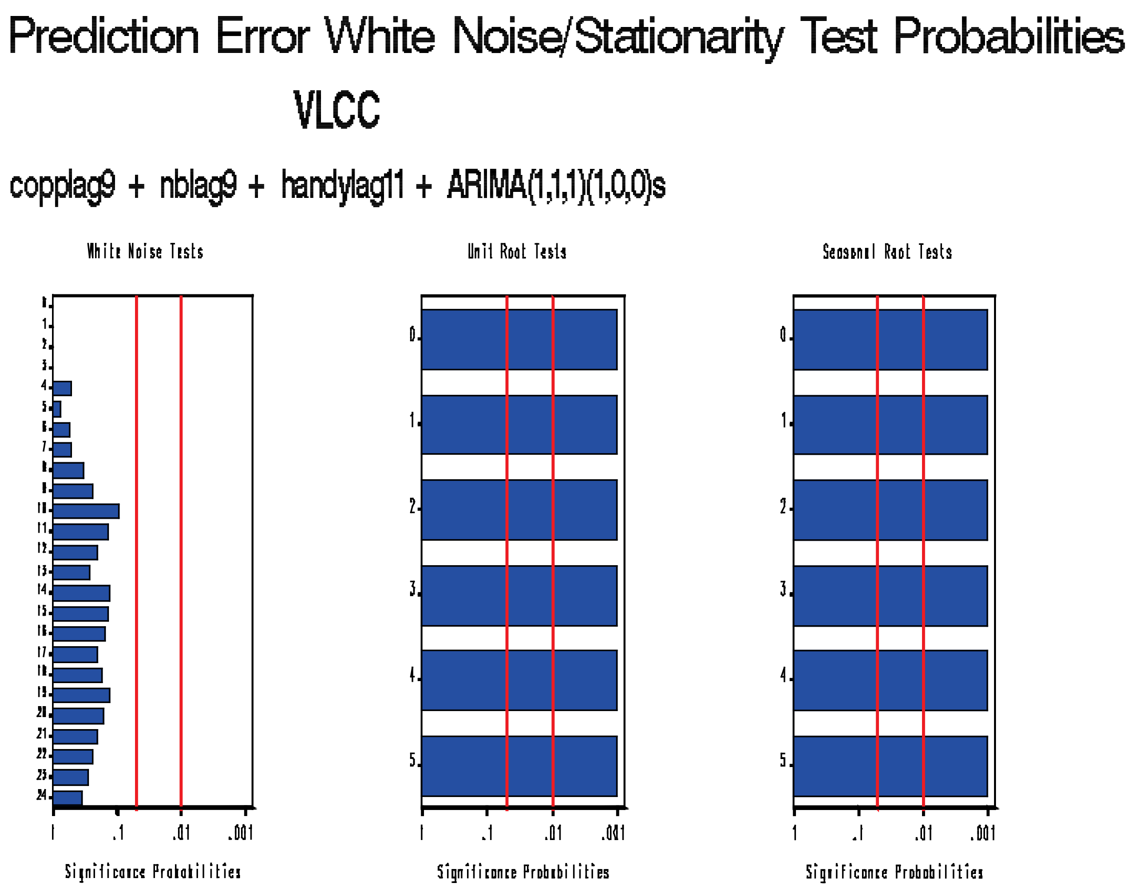

Figure 83 : Prediction Error White Noise/Stationarity Test Probabilities for VLCC (ARIMAX)

The White Noise, Unit root and Seasonal Root tests indicate proper model fit for all lags. We can check the significance of each autoregressive parameter and seasonality by taking a look at the output table below.

VLCC (ARIMAX) Model Forecast

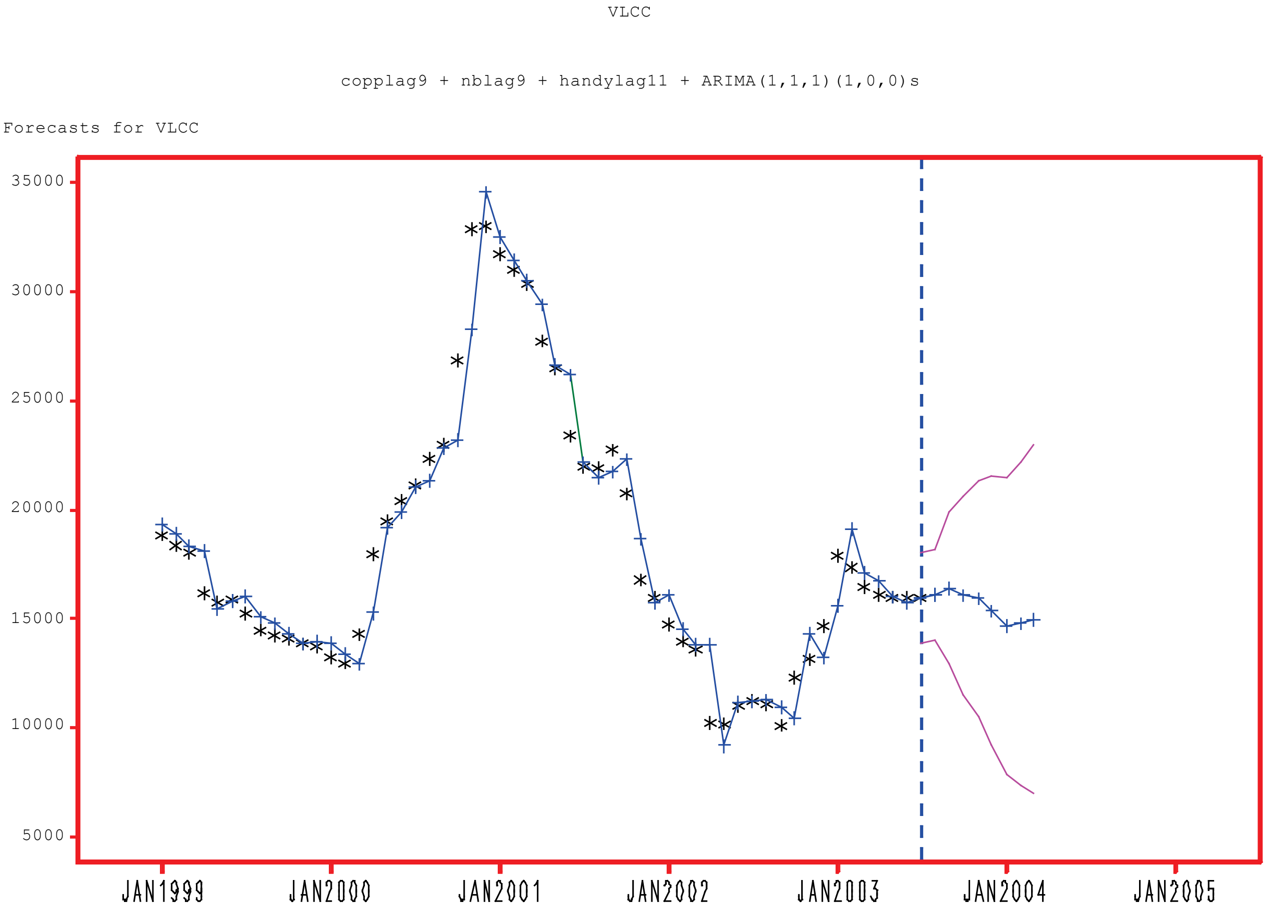

Figure 84 : Forecasts for VLCC (ARIMAX)

The significant model parameters are highlighted in bold. Both the Crude Oil Purchase Price and the Newbuilding variables were insignificant in defining the model.

PARM	VALUE	STDERR	T	P
Intercept	15.68	76.56	0.20	0.84
Moving Average, Lag 1	-0.21	0.17	-1.19	0.24
Autoregressive, Lag 1	0.15	0.18	0.87	0.39
Seasonal Autoregressive, Lag 12	**-0.19**	**0.06**	**-2.97**	**0.00**
copplag9	0.32	0.50	0.65	0.52
nblag9	-5.86	6.30	-0.93	0.35
handylag11	**-0.33**	**0.18**	**-1.87**	**0.06**
Model Variance (sigma squared)	1102289			

If there is any one statistic that normally takes precedence over the others in a model comparison, it is the mean squared error within the estimation period, or equivalently its square root, the root mean squared error (RMSE). This is the statistic whose value is minimized during the parameter estimation process, and it is the statistic that determines the width of the confidence intervals for predictions. When comparing the two VLCC models that have been created, (Univariate and ARMAX) we look to the model that minimizes the RMSE.

MODEL	RMSE
UNIVARIATE	991.69
ARIMAX	1037.3

Surprisingly, the ARIMAX does not show a significantly lower RMSE. The ARIMAX's forecast is:

DATE	FORECAST	ACTUAL
8/1/2003	16097	16000
9/1/2003	16390	16625
10/1/2003	16085	16200
11/1/2003	15945	18000
12/1/2003	15377	19000
1/1/2004	14686	18200
2/1/2004	14800	17375
3/1/2004	14989	17000

6 GARCH Models

General

Some financial data appears to have variance that changes locally. This change of variance, otherwise known as heteroscedasticity was studied and modeled using ARCH and GARCH models by Engle (1982) and Bollerslev (1986). In these models, the innovations variance h(t) at time t is assumed to follow an autoregressive moving average model, with squared residuals where the uncorrelated shocks usually go. The variance model is

$$h(t) = \omega + \sum_{i=1}^{q} \alpha_i e^2(t-i) + \sum_{j=1}^{p} \gamma_j h(t-j)$$

where e(t) is the residual at time t. The model can even be fit with unit roots in the "autoregressive" part in which case the models are called GARCH or EGARCH models.

Time Charter rates fall under this category and it seems logical for GARCH models to be tested on the Time Charter rate returns in order to find whether GARCH models perform better than their ARIMA and ARMAX counterparts.

6.1 Regression with Autocorrelated Errors

Ordinary regression analysis is based on several statistical assumptions. One key assumption is that the errors are independent of each other. However, with time series data, the ordinary regression residuals usually are correlated over time. It is not desirable to use ordinary regression analysis for time series data since the assumptions on which the classical linear regression model is based on will usually be violated.

Violation of the independent errors assumption has three important consequences for ordinary regression. First, statistical tests of the significance of the parameters and the confidence limits for the predicted values are not correct. Second, the estimates of the regression coefficients are not as efficient as they would be if the autocorrelation were taken into account. Third, since the ordinary regression residuals are not independent, they contain information that can be used to improve the prediction of future values. By augmenting the regression model with an autoregressive model for the random error, we can account for the autocorrelation of the errors. Instead of the usual regression model, the following autoregressive error model is used:

$$yt = x_t^{'}\beta + v_t$$
$$v_t = -\phi_1 v_{t-1} - \phi_2 v_{t-2} - ... - \phi_m v_{t-m} + \varepsilon_t$$
$$\varepsilon_t \sim IN(0, \sigma^2)$$

The notation $\varepsilon_t \sim IN(0, \sigma^2)$ indicates that each σ^2 is normally and independently distributed with mean 0 and variance σ^2. By simultaneously estimating the regression coefficients β and the autoregressive error model parameters φ_i, we correct the regression estimates for autocorrelation. Thus, this kind of regression analysis is often called *autoregressive error correction* or *serial correlation correction*.

6.2 Parameter Estimates for Generalized Linear Models

The **Parameter Estimates** table for generalized linear models includes the following:

Variable : names the variable associated with the estimated parameter. The name **INTERCEPT** represents the estimate of the intercept parameter.

DF : is the degrees of freedom associated with each parameter estimate. There is one degree of freedom unless the model is not full rank. In this case, any parameter that is confounded with previous parameters in the model has its degrees of freedom set to 0.

Estimate : is the parameter estimate.

Std Error : is the estimated standard deviation of the parameter estimate.

ChiSq : is the chi-squared test statistic for testing that the parameter is 0. This is computed as the square of the ratio of the parameter estimate divided by the standard error.

Pr > ChiSq : is the probability of obtaining an chi-squared statistic greater than that observed given that the true parameter is 0. A small *p*-value is evidence for concluding that the parameter is not 0.

FIT STATISTICS

ADJRSQ	computes adjusted R^2
AIC	computes Akaike's information criterion
MSE	computes MSE for each model
RMSE	displays root MSE for each model
SBC	computes the SBC statistic
DW	Computes Durbin Watson statistic
DFE	specifies the degrees of freedom associated with the root mean square error
SSE	computes error sum of squares for each model

6.3 Ordinary Least Squares Analysis

In the first In the later GARCH procedures, the Autoreg procedure will always compute an OLS Regression. An example of an OLS Regression of the VLCC time charter rate is shown below.

6.3.1 VLCC OLS

```
                     Dependent Variable    vlcc


                   Ordinary Least Squares Estimates

   SSE                  5103111710    DFE                     284
   MSE                    17968703    Root MSE               4239
   SBC                  5598.32244    AIC              5591.01045
   Regress R-Square         0.4613    Total R-Square       0.4613
   Durbin-Watson            0.0741


                                       Standard                Approx
Variable        DF     Estimate           Error    t Value    Pr > |t|

Intercept        1        -3729            1179      -3.16      0.0017
date             1       1.5556          0.0997      15.60      <.0001
                                                                          8.32
```

Here we see the statistics for the model residuals. The model root mean square error (Root MSE) is 4239, and the model R^2 is .4613. Notice that two R^2 statistics are shown, one for the regression model (Reg Rsq) and one for the full model (Total Rsq) that includes the autoregressive error process, if any. In this case, an autoregressive error model is not used, so the two R^2 statistics are the same. Other statistics shown are the sum of square errors (SSE), mean square error (MSE), error degrees of freedom (DFE, the number of observations minus the number of parameters), the information criteria SBC and AIC, and the Durbin-Watson statistic. A table of regression coefficients, with standard errors and *t*-tests is also shown. In this case, the estimated model is

$$y_t = -3729 + 1.1556t$$
$$Est.Var(\varepsilon_t) = 17968703$$

The same process can be used in order to create an OLS model for the rest of the time charter rates. The volatility of the time charter rates, and the fact that we must use return in order to achieve independence and stationarity forces us to use autoregressive ARIMA, GARCH processes and an OLS Regression is of no real importance.

6.4 Autoregressive Error Model

6.4.1 Testing for Autocorrelation

When time series data are used in regression analysis, often the error term is not independent through time. If the error term is autocorrelated, the efficiency of ordinary least-squares (OLS) parameter estimates is adversely affected and standard error estimates are biased. The Durbin-Watson d statistic can be used to test for the presence of first-order autocorrelation in OLS residuals. When autocorrelation is detected, using regression with correction for autocorrelation gives you several alternate estimation methods that produce better estimates. In many cases, the parameter estimates produced are similar to the OLS estimates. However, the standard errors can be quite different, affecting the tests of significance.

6.4.2 Durbin Watson & ARCH Tests

Using the Durbin-Watson test you can decide if autocorrelation correction is needed. However, generalized Durbin-Watson tests should not be used to decide on the autoregressive order. The higher-order tests assume the absence of lower-order autocorrelation. If the ordinary Durbin-Watson test indicates no first-order autocorrelation, you can use the second-order test to check for second-order autocorrelation. Once autocorrelation is detected, further tests at higher orders are not appropriate. If first-order Durbin-Watson tests are significant, the orders 2 through 12 can be ignored. One can use the DW= option to request higher-order Durbin-Watson statistics. Since the ordinary Durbin-Watson statistic only tests for first-order autocorrelation, the Durbin-Watson statistics for higher-order autocorrelation are called *generalized Durbin-Watson statistics*.

When using Durbin-Watson tests to check for autocorrelation, you should specify an order at least as large as the order of any potential seasonality, since seasonality produces autocorrelation at the seasonal lag. For example, for quarterly data use DW=4, and for monthly data use DW=12.

6.4.3 Q and LM tests for ARCH Disturbances

The Q and LM tests for each time charter return shows that:

VLCC: Normally, a strong case for a heteroscedastic series would be that both the Q and LM tests would have a Pr > Q and Pr > LM that is < 0.0001. This isn't exactly the case, although we can't rule it out either.

AFRAMAX: The aframax retu rns show a strong heteroscedastic behavior as indicated by the value of both the Q and LM tests.

HANDYSIZE : Heteroscedastic effects for the series begin strongly yet drop off as be look backward in time.

SUEZMAX : We can definitely rule out heteroscedastic effects because the Q and LM tests are strongly against the hypothesis (Pr > Q amd Pr > LM close to one).

On the other hand, the Q and LM tests show that series themselves (not the returns) have a very strong heteroscedastic behavior as indicated by Pr > Q and Pr > LM that are < 0.0001 at all lags. This means that by using transforming a series to returns, heteroscedasticity is removed and the series is able to be better modeled using non-garch methods. AFRAMAX still retains very strong heteroscedastic properties which will be compensated for by a GARCH model. Having created the best possible model for each time charter rate, we can obtain a prediction and analyze the residuals.

6.4.4 VLCC Returns (Logged Differences)

Dependent Variable vlcc_r

Ordinary Least Squares Estimates

SSE	1.90525963	DFE	283
MSE	0.00673	Root MSE	0.08205
SBC	-607.14326	AIC	-614.44824
Regress R-Square	0.0007	Total R-Square	0.0007

Durbin-Watson Statistics

Order	DW	Pr < DW	Pr > DW
1	1.6260	0.0006	0.9994
2	1.9639	0.3803	0.6197
3	1.9432	0.3373	0.6627
4	1.9311	0.3217	0.6783
5	2.0808	0.8057	0.1943
6	1.9855	0.5464	0.4536
7	1.8051	0.0884	0.9116
8	2.0889	0.8668	0.1332
9	2.1205	0.9249	0.0751
10	1.7299	0.0349	0.9651
11	1.8853	0.3333	0.6667
12	2.1424	0.9644	0.0356
13	1.9087	0.4552	0.5448

NOTE: Pr<DW is the p-value for testing positive autocorrelation, and Pr>DW is the p-value for testing negative autocorrelation.

Q and LM Tests for ARCH Disturbances

Order	Q	Pr > Q	LM	Pr > LM
1	0.6567	0.4177	0.6624	0.4157
2	1.0968	0.5779	1.0671	0.5865
3	1.7709	0.6213	1.7956	0.6159
4	2.6504	0.6179	2.5651	0.6330
5	7.2556	0.2023	7.8040	0.1674
6	9.9982	0.1247	9.8403	0.1315
7	10.0676	0.1848	9.8903	0.1949
8	10.2087	0.2507	10.0863	0.2590
9	10.2098	0.3338	10.2081	0.3339
10	10.5932	0.3901	10.7950	0.3737
11	10.5932	0.4779	11.0190	0.4417
12	13.8836	0.3082	13.8771	0.3086

Variable	DF	Estimate	Standard Error	t Value	Approx Pr > |t|
Intercept	1	-0.009507	0.0230	-0.41	0.6792
date	1	8.423E-7	1.9409E-6	0.43	0.6646

8.34

6.4.5 AFRAMAX RETURNS (Logged Differences)

Dependent Variable aframax_r

Ordinary Least Squares Estimates

SSE	1.6604546	DFE	283
MSE	0.00587	Root MSE	0.07660
SBC	-646.33842	AIC	-653.6434
Regress R-Square	0.0006	Total R-Square	0.0006

Durbin-Watson Statistics

Order	DW	Pr < DW	Pr > DW
1	1.4654	<.0001	1.0000
2	1.9579	0.3611	0.6389
3	2.0192	0.5879	0.4121
4	1.9343	0.3315	0.6685
5	1.6683	0.0041	0.9959
6	1.6367	0.0022	0.9978
7	1.7019	0.0128	0.9872
8	1.9945	0.6230	0.3770
9	2.1892	0.9784	0.0216
10	2.1140	0.9256	0.0744
11	1.8031	0.1295	0.8705
12	1.8855	0.3560	0.6440
13	1.8819	0.3670	0.6330

NOTE: Pr<DW is the p-value for testing positive autocorrelation, and Pr>DW is the p-value for testing negative autocorrelation.

Q and LM Tests for ARCH Disturbances

Order	Q	Pr > Q	LM	Pr > LM
1	1.7990	0.1798	1.7320	0.1882
2	27.5578	<.0001	26.1157	<.0001
3	29.8806	<.0001	26.8488	<.0001
4	31.6054	<.0001	26.9640	<.0001
5	51.9056	<.0001	42.0172	<.0001
6	55.0233	<.0001	43.2116	<.0001
7	60.6806	<.0001	43.2538	<.0001
8	60.7488	<.0001	45.6035	<.0001
9	70.1250	<.0001	51.8192	<.0001
10	70.1957	<.0001	53.0552	<.0001
11	70.9181	<.0001	54.9502	<.0001
12	70.9189	<.0001	55.0305	<.0001

Variable	DF	Estimate	Standard Error	t Value	Approx Pr > |t|
Intercept	1	-0.006855	0.0214	-0.32	0.7494
date	1	7.2495E-7	1.8119E-6	0.40	0.6894

8.34

6.4.6 SUEZMAX RETURNS (Logged Differences)

Dependent Variable suezmax_r

Ordinary Least Squares Estimates

SSE	0.95551675	DFE	283
MSE	0.00338	Root MSE	0.05811
SBC	-803.82782	AIC	-811.1328
Regress R-Square	0.0024	Total R-Square	0.0024

Durbin-Watson Statistics

Order	DW	Pr < DW	Pr > DW
1	1.5049	<.0001	1.0000
2	1.7516	0.0176	0.9824
3	1.6200	0.0007	0.9993
4	1.7860	0.0452	0.9548
5	1.9010	0.2555	0.7445
6	1.9059	0.2892	0.7108
7	1.8252	0.1189	0.8811
8	1.8500	0.1816	0.8184
9	1.8473	0.1914	0.8086
10	1.9772	0.6127	0.3873
11	2.0317	0.7902	0.2098
12	2.0341	0.8126	0.1874
13	2.1436	0.9696	0.0304

NOTE: Pr<DW is the p-value for testing positive autocorrelation, and Pr>DW is the p-value for testing negative autocorrelation.

Q and LM Tests for ARCH Disturbances

Order	Q	Pr > Q	LM	Pr > LM
1	0.0080	0.9287	0.0086	0.9263
2	0.0848	0.9585	0.0806	0.9605
3	0.1033	0.9914	0.0952	0.9924
4	0.1378	0.9977	0.1336	0.9979
5	0.4466	0.9939	0.4487	0.9939
6	0.7936	0.9922	0.7746	0.9927
7	0.9104	0.9961	0.8608	0.9968
8	0.9217	0.9987	0.8693	0.9989
9	0.9277	0.9996	0.8764	0.9997
10	0.9432	0.9999	0.8917	0.9999
11	0.9999	0.9999	0.9655	1.0000
12	1.0026	1.0000	0.9689	1.0000

Variable	DF	Estimate	Standard Error	t Value	Approx Pr > |t|
Intercept	1	-0.0111	0.0163	-0.68	0.4940
date	1	1.1315E-6	1.3745E-6	0.82	0.4111

8.35

6.4.7 HANDYSIZE RETURNS (Logged Differences)

Dependent Variable handysize_r

Ordinary Least Squares Estimates

SSE	0.45999142	DFE	283
MSE	0.00163	Root MSE	0.04032
SBC	-1012.1755	AIC	-1019.4805
Regress R-Square	0.0153	Total R-Square	0.0153

Durbin-Watson Statistics

Order	DW	Pr < DW	Pr > DW
1	1.3197	<.0001	1.0000
2	1.7005	0.0055	0.9945
3	1.5991	0.0004	0.9996
4	1.5848	0.0003	0.9997
5	1.7071	0.0105	0.9895
6	1.9186	0.3268	0.6732
7	1.8915	0.2682	0.7318
8	1.9479	0.4677	0.5323
9	1.8693	0.2462	0.7538
10	1.9820	0.6282	0.3718
11	2.0883	0.9008	0.0992
12	2.1241	0.9505	0.0495
13	2.0386	0.8379	0.1621

NOTE: Pr<DW is the p-value for testing positive autocorrelation, and Pr>DW is the p-value for testing negative autocorrelation.

Q and LM Tests for ARCH Disturbances

Order	Q	Pr > Q	LM	Pr > LM
1	4.2549	0.0391	4.1299	0.0421
2	4.2687	0.1183	4.1499	0.1256
3	4.2763	0.2331	4.1567	0.2450
4	6.2331	0.1824	6.0405	0.1961
5	6.6786	0.2457	6.1521	0.2917
6	6.7092	0.3486	6.2321	0.3977
7	7.6570	0.3638	7.3154	0.3968
8	7.8834	0.4449	7.9691	0.4365
9	7.8978	0.5445	8.0095	0.5332
10	8.0125	0.6276	8.1125	0.6178
11	8.1228	0.7023	8.1897	0.6962
12	8.1337	0.7746	8.2081	0.7687

Variable	DF	Estimate	Standard Error	t Value	Approx Pr > |t|
Intercept	1	-0.0238	0.0113	-2.11	0.0356
date	1	1.9968E-6	9.5368E-7	2.09	0.0372

8.36

As we can see from the Q and LM tests, the log-differenced data exhibit the following characteristics

- In all the cases, the first-order Durbin-Watson test is highly significant, with *p* close to 0, rejecting the hypothesis of no first-order autocorrelation. Thus, autocorrelation correction is needed. The higher-order tests assume the absence of lower-order autocorrelation. If the ordinary Durbin-Watson test indicates no first-order autocorrelation, you can use the second-order test to check for second-order autocorrelation. Once autocorrelation is detected, further tests at higher orders are not appropriate. Since the first-order Durbin-Watson test is significant, the higher tests in higher orders can be ignored.
- The Q and LM tests for ARCH disturbances show that the ARCH effects have been removed from the time series and the GARCH models in all the time charter rates except for AFRAMAX.

For the GARCH models to work correctly there must exist ARCH disturbances (meaning that the Pr > Q should be close to 0). In order to maintain non-constant variance throughout the time series, the alternate transformation tested for ARIMA models will undergo the same analysis.

6.4.8 VLCC Returns (Alternate Data Transformation)

Dependent Variable vlcc_r

Ordinary Least Squares Estimates

SSE	363389261	DFE	283
MSE	1284061	Root MSE	1133
SBC	4826.7713	AIC	4819.46632
Regress R-Square	0.0002	Total R-Square	0.0002

Durbin-Watson Statistics

Order	DW	Pr < DW	Pr > DW
1	1.3487	<.0001	1.0000
2	1.8679	0.1319	0.8681
3	1.9354	0.3135	0.6865
4	1.8641	0.1516	0.8484
5	1.9417	0.3772	0.6228
6	1.7944	0.0665	0.9335
7	1.7936	0.0737	0.9263
8	2.1169	0.9113	0.0887
9	2.1250	0.9302	0.0698
10	1.8166	0.1415	0.8585
11	2.0662	0.8641	0.1359
12	2.4153	1.0000	<.0001
13	2.1789	0.9852	0.0148

NOTE: Pr<DW is the p-value for testing positive autocorrelation, and Pr>DW is the p-value for testing
negative autocorrelation.

Q and LM Tests for ARCH Disturbances

Order	Q	Pr > Q	LM	Pr > LM
1	9.5411	0.0020	9.4929	0.0021
2	10.0239	0.0067	9.5145	0.0086
3	10.4489	0.0151	10.1377	0.0174
4	10.6003	0.0314	10.5643	0.0319
5	15.5080	0.0084	14.9659	0.0105
6	20.1126	0.0026	16.6599	0.0106
7	37.9457	<.0001	29.5765	0.0001
8	41.1710	<.0001	29.9589	0.0002
9	41.7702	<.0001	30.0396	0.0004
10	41.7706	<.0001	30.0479	0.0008
11	42.1437	<.0001	30.0876	0.0015
12	59.4324	<.0001	39.7636	<.0001

Variable	DF	Estimate	Standard Error	t Value	Approx Pr > |t|
Intercept	1	-67.4533	317.2075	-0.21	0.8318
date	1	0.006203	0.0268	0.23	0.8172

6.4.9 AFRAMAX Returns (Alternate Data Transformation)

Dependent Variable aframax_r

Ordinary Least Squares Estimates

SSE	161393793	DFE	283
MSE	570296	Root MSE	755.17951
SBC	4595.45752	AIC	4588.15254
Regress R-Square	0.0003	Total R-Square	0.0003

Durbin-Watson Statistics

Order	DW	Pr < DW	Pr > DW
1	1.2756	<.0001	1.0000
2	1.6244	0.0007	0.9993
3	1.6448	0.0015	0.9985
4	1.7252	0.0135	0.9865
5	1.6763	0.0050	0.9950
6	1.5378	0.0001	0.9999
7	1.7062	0.0140	0.9860
8	1.8305	0.1411	0.8589
9	1.9003	0.3360	0.6640
10	1.9478	0.5151	0.4849
11	1.9448	0.5291	0.4709
12	1.9400	0.5367	0.4633
13	2.0295	0.8184	0.1816

NOTE: Pr<DW is the p-value for testing positive autocorrelation, and Pr>DW is the p-value for testing negative autocorrelation.

Q and LM Tests for ARCH Disturbances

Order	Q	Pr > Q	LM	Pr > LM
1	7.4275	0.0064	6.9832	0.0082
2	39.7799	<.0001	34.4305	<.0001
3	43.4397	<.0001	34.5562	<.0001
4	68.7746	<.0001	46.5569	<.0001
5	70.7686	<.0001	46.6464	<.0001
6	71.5515	<.0001	49.8140	<.0001
7	74.2809	<.0001	51.0100	<.0001
8	83.4476	<.0001	57.1173	<.0001
9	84.9488	<.0001	57.1784	<.0001
10	85.8059	<.0001	57.3342	<.0001
11	88.6389	<.0001	57.9300	<.0001
12	93.1527	<.0001	58.1598	<.0001

Variable	DF	Estimate	Standard Error	t Value	Approx Pr > |t|
Intercept	1	-36.6567	211.3980	-0.17	0.8625
date	1	0.005036	0.0179	0.28	0.7782

6.4.10 HANDYSIZE Returns (Alternate Data Transformation)

Dependent Variable handysize_r

Ordinary Least Squares Estimates

SSE	46734223	DFE	283
MSE	165139	Root MSE	406.37249
SBC	4242.2369	AIC	4234.93192
Regress R-Square	0.0201	Total R-Square	0.0201

Durbin-Watson Statistics

Order	DW	Pr < DW	Pr > DW
1	1.2179	<.0001	1.0000
2	1.6951	0.0048	0.9952
3	1.6011	0.0004	0.9996
4	1.5561	0.0001	0.9999
5	1.6343	0.0017	0.9983
6	1.7611	0.0371	0.9629
7	1.8673	0.2053	0.7947
8	1.9358	0.4275	0.5725
9	1.8034	0.1064	0.8936
10	1.8967	0.3466	0.6534
11	1.9524	0.5544	0.4456
12	2.0652	0.8750	0.1250
13	2.0944	0.9276	0.0724

NOTE: Pr<DW is the p-value for testing positive autocorrelation, and Pr>DW is the p-value for testing negative autocorrelation.

Q and LM Tests for ARCH Disturbances

Order	Q	Pr > Q	LM	Pr > LM
1	15.9827	<.0001	15.2904	<.0001
2	16.5800	0.0003	15.3321	0.0005
3	16.6745	0.0008	15.3458	0.0015
4	17.2931	0.0017	15.7713	0.0033
5	17.3227	0.0039	15.8187	0.0074
6	17.3239	0.0082	15.8274	0.0147
7	19.5173	0.0067	18.0634	0.0117
8	19.6096	0.0119	18.2808	0.0192
9	19.6908	0.0199	18.4181	0.0306
10	19.9750	0.0295	18.8705	0.0419
11	21.6533	0.0272	19.5909	0.0513
12	28.4133	0.0048	23.3663	0.0248

Variable	DF	Estimate	Standard Error	t Value	Approx Pr > |t|
Intercept	1	-278.9706	113.7562	-2.45	0.0148
date	1	0.0232	0.009613	2.41	0.0166

6.4.11 SUEZMAX Returns (Alternate Data Transformation)

Dependent Variable suezmax_r

Ordinary Least Squares Estimates

SSE	314523800	DFE	283
MSE	1111392	Root MSE	1054
SBC	4785.61306	AIC	4778.30809
Regress R-Square	0.0010	Total R-Square	0.0010

Durbin-Watson Statistics

Order	DW	Pr < DW	Pr > DW
1	1.2764	<.0001	1.0000
2	1.5722	0.0001	0.9999
3	1.6014	0.0004	0.9996
4	1.6406	0.0016	0.9984
5	1.6883	0.0068	0.9932
6	1.8189	0.0979	0.9021
7	1.9033	0.3022	0.6978
8	1.9846	0.5907	0.4093
9	1.8586	0.2186	0.7814
10	2.0437	0.8019	0.1981
11	2.0916	0.9056	0.0944
12	2.3058	0.9993	0.0007
13	2.5395	1.0000	<.0001

NOTE: Pr<DW is the p-value for testing positive autocorrelation, and Pr>DW is the p-value for testing negative autocorrelation.

Q and LM Tests for ARCH Disturbances

Order	Q	Pr > Q	LM	Pr > LM
1	2.7813	0.0954	2.7619	0.0965
2	2.8045	0.2460	2.7620	0.2513
3	3.1343	0.3714	3.0809	0.3793
4	11.0227	0.0263	10.3668	0.0347
5	18.3213	0.0026	15.1680	0.0097
6	18.3305	0.0055	15.4572	0.0170
7	18.5045	0.0099	15.5760	0.0293
8	20.2951	0.0093	16.2858	0.0385
9	22.6780	0.0070	16.7663	0.0525
10	22.7044	0.0119	17.1906	0.0703
11	22.7641	0.0191	17.2700	0.1001
12	24.8019	0.0158	18.2630	0.1079

Variable	DF	Estimate	Standard Error	t Value	Approx Pr > |t|
Intercept	1	-119.5658	295.1102	-0.41	0.6857
date	1	0.0133	0.0249	0.54	0.5929

The Q and LM Tests indicate that the alternate data transformation retains the time serie's ARCH effects and is more suitable for all time charter rates. Thus, unlike the ARIMA models which used simple differencing, the alternate data transformation is suited for GARCH models and will be used in the analysis and models that follow.

6.5 Stepwise Autoregression

Once we determine that autocorrelation correction is needed, we must select the order of the autoregressive error model to use. One way to select the order of the autoregressive error model is *stepwise autoregression*. The stepwise autoregression method initially fits a high-order model with many autoregressive lags and then sequentially removes autoregressive parameters until all remaining autoregressive parameters have significant *t*-tests.

We use stepwise autoregression, specifying the BACKSTEP option a large order with the NLAG= option. The output is in the form of two tables labeled "Backward Elimination of Autoregressive Terms" and "Estimates of Autoregressive Parameters". The first table lists the eliminates autoregressive parameters while the second table lists the parameters with significant *t*-tests.

The following statements show the stepwise feature, using an initial order of 13 – note that the whole output is exactly the same as the previous OLS Regression command, but with the addition of the two tables described above. In the name of clarity, the first part of the output, the OLS Regression, is withheld.

6.5.1 VLCC RETURNS

Backward Elimination of Autoregressive Terms

Lag	Estimate	t Value	Pr > |t|
2	0.003671	0.06	0.9533
5	0.011384	0.18	0.8539
13	-0.020317	-0.34	0.7354
3	0.038901	0.67	0.5057
4	-0.026945	-0.49	0.6242
11	0.066837	1.10	0.2728
7	-0.062930	-1.04	0.2988
9	0.072999	1.21	0.2290
10	-0.097623	-1.79	0.0746

Estimates of Autoregressive Parameters

Lag	Coefficient	Standard Error	t Value
1	-0.313378	0.054541	-5.75
6	-0.110569	0.054517	-2.03
8	0.110551	0.054558	2.03
12	0.231671	0.054686	4.24

8.37

6.5.2 AFRAMAX RETURNS

Backward Elimination of Autoregressive Terms

Lag	Estimate	t Value	Pr > |t|
7	-0.007744	-0.12	0.9036
4	0.023195	0.37	0.7141
12	0.022609	0.36	0.7206
2	-0.027745	-0.44	0.6601
11	0.032739	0.54	0.5887
5	-0.031513	-0.53	0.5994
8	0.032644	0.55	0.5854
10	0.046914	0.79	0.4306
9	0.074275	1.32	0.1891
3	-0.070384	-1.24	0.2156

Estimates of Autoregressive Parameters

Lag	Coefficient	Standard Error	t Value
1	-0.327748	0.055709	-5.88
6	-0.121316	0.055602	-2.18
13	0.111569	0.055547	2.01

8.38

6.5.3 HANDYSIZE RETURNS

Backward Elimination of Autoregressive Terms

Lag	Estimate	t Value	Pr > |t|
7	0.002887	0.04	0.9644
6	-0.006980	-0.11	0.9092
11	-0.026268	-0.41	0.6809
10	0.037874	0.63	0.5262
5	-0.051297	-0.85	0.3983
2	0.053974	0.86	0.3912
8	0.054693	0.91	0.3649
9	-0.059888	-1.06	0.2907
13	0.062201	1.05	0.2939
3	-0.099176	-1.68	0.0935
12	0.091590	1.68	0.0946

Estimates of Autoregressive Parameters

Lag	Coefficient	Standard Error	t Value
1	-0.349154	0.055642	-6.28
4	-0.136185	0.055642	-2.45

8.39

6.5.4 SUEZMAX RETURNS

Backward Elimination of Autoregressive Terms

Lag	Estimate	t Value	Pr > |t|
11	-0.002060	-0.03	0.9730
6	0.002797	0.05	0.9631
7	-0.013645	-0.23	0.8158
2	-0.030344	-0.51	0.6137
5	-0.052063	-0.90	0.3706
10	0.053160	0.92	0.3568
8	0.066728	1.16	0.2469
4	-0.065988	-1.14	0.2538
9	-0.073785	-1.36	0.1744

Estimates of Autoregressive Parameters

Lag	Coefficient	Standard Error	t Value
1	-0.266964	0.055283	-4.83
3	-0.118421	0.054050	-2.19
12	0.112004	0.056591	1.98
13	0.232628	0.057723	4.03

8.40

The estimates of the autocorrelations are shown for 13 lags. The backward elimination of autoregressive terms report shows which autoregressive parameters at lags were insignificant and eliminated, resulting in the *Estimates of Autoregressive Parameters*. The stepwise autoregressive process is performed using the Yule-Walker method. The maximum likelihood estimates are produced after the order of the model is determined from the significance tests of the preliminary Yule-Walker estimates. When using stepwise autoregression, it is a good idea to specify an NLAG= option value larger than the order of any potential seasonality, since seasonality produces autocorrelation at the seasonal lag. In this case, we have monthly data which uses NLAG=13. We see that the BACKSTEP option in the variables has dropped for example lags 4 or 5 and 10 or 11. This means that a parameter at a longer lag (for instance lag=12) was kept while some smaller lags are dropped. This is called a *subset model*, since the number of estimated autoregressive parameters is smaller than the order of the model. Subset models are common for seasonal data and often correspond to *factored* autoregressive models. A factored model is the product of simpler autoregressive models. For example, the best model for seasonal monthly data may be the combination of a first-order model for recent effects with a higher order subset model for the seasonality (i.e. a single parameter at lag 12). As is the case in the SUEZMAX time charter rate, this results in an order 13 subset model with nonzero parameters at lags 1, 3, and 13.

For each variable, the subset model is

t-value (lag)													
RATE	1	2	3	4	5	6	7	8	9	10	11	12	13
VLCC	-5.75	-	-	-	-	-2.03	-	2.03	-	-	-	4.24	-
AFRAMAX	-5.88	-	-	-	-	-2.18	-	-	2.48	-	-	-	2.01
HANDYSIZE	-6.28	-		-2.45	-	-	-	-	-	-	-	-	-
SUEZMAX	-4.83	-	-2.19	-	-	-	-	-	-	-	-	1.98	4.03

The information we can gain from this table on our time charter rates can give us an idea of what time lag drives its prices. Alternately, the author assumes the characteristic drivers are more or less the same in the market but act on a different time-periods on the series. VLCC and AFRAMAX seem to be affected by events occurring at longer time periods than SUEZMAX and HANDYSIZE ships. In the GARCH modeling process that follows, the changing of GARCH parameters p and q have, as a result, a change in the importance of the t-values of the lags. The following GARCH models have been created with the purpose of creating a model whose GARCH and Autoregressive parameters are significant (as shown by the t-value and probability tests).

6.6 Heteroscedasticity

Modeling non-constant variance, or heteroscedasticity, improves the efficiency of estimates of the parameters associated with the mean of a series and provides insight into the volatility of a series. One of the key assumptions of regression analysis is that the variance of the errors is constant across observations. This assumption is often violated when modeling time series or panel data, resulting in inefficient parameter estimates and inaccurate forecast error variance. If the errors for a model are heteroscedastic and the functional form of the variance is known, the model for the variance can be estimated along with the regression function. Our variables use monthly data created from an average of weekly figures from October 1979 to July 2003.

6.6.1 Testing for Heteroscedasticity

One of the key assumptions of the ordinary regression model is that the errors have the same variance throughout the sample. This is also called the *homoscedasticity* model. If the error variance is not constant, the data are said to be *heteroscedastic*. Since ordinary least-squares regression assumes constant error variance, heteroscedasticity causes the OLS estimates to be inefficient. Models that take into account the changing variance can make more efficient use of the data. Also, heteroscedasticity can make the OLS forecast error variance inaccurate since the predicted forecast variance is based on the average variance instead of the variability at the end of the series.

To test for heteroscedasticity with PROC AUTOREG, specify the ARCHTEST option. We regress VLCC and use the ARCHTEST option to test for heteroscedastic OLS residuals. The DWPROB option is also used to test

The Q statistics test for changes in variance across time using lag windows ranging from 1 through 12. The *p*-values for the test statistics are given and strongly indicate heteroscedasticity, with $p < 0.0001$ for all lag windows.

The Lagrange multiplier (LM) tests also indicate heteroscedasticity. These tests can also help determine the order of the ARCH model appropriate for modeling the heteroscedasticity, assuming that the changing variance follows an autoregressive conditional heteroscedasticity model.

6.6.2 Heteroscedasticity and GARCH Models

There are several approaches to dealing with heteroscedasticity. If the error variance at different times is known, weighted regression is a good method. If, as is usually the case, the error variance is unknown and must be estimated from the data, you can model the changing error variance.

The *generalized autoregressive conditional heteroscedasticity* (GARCH) model is one approach to modeling time series with heteroscedastic errors. The GARCH regression model with autoregressive errors is

$$
\begin{aligned}
y_t &= x_t^{'}\beta + v_t \\
v_t &= \varepsilon_t - \varphi_1 y_{t-1} - \dots - \varphi_m v_{t-m} \\
\varepsilon_t &= \sqrt{h_t e_t} \\
h_t &= \omega + \sum_{i=1}^{q} a_i \varepsilon_{t-i}^2 + \sum_{j=1}^{p} \gamma_j h_{t-j} \\
e_t &\sim IN(0,1)
\end{aligned}
$$

This model combines the *m*th-order autoregressive error model with the GARCH(p,q) variance model. It is denoted as the AR(m)- GARCH (p,q) regression model. The Lagrange multiplier (LM) tests shown in Figure 1 can help determine the order of the ARCH model appropriate for the data. If the tests are significant ($p<.0001$) through order 12, they would indicates that a very high-order ARCH model is needed to model the heteroscedasticity. The basic ARCH(q) model (p=0) is a *short memory* process in that only the most recent q squared residuals are used to estimate the changing variance. The GARCH model (p>0) allows *long memory* processes, which use all the past squared residuals to estimate the current variance.

The GARCH (p,q) model is specified with the GARCH =(P=p,Q=q) option in the MODEL statement. The basic ARCH(q) model is the same as the GARCH (0,q) model and is specified with the GARCH =(Q=q) option. Using the statements below we create a GARCH model for each of the returns. The autoregressive parameters are the ones calculated by the *Estimates of Autoregressive Parameters* table in the previous output. In this first part we will not opt for a prediction or analyze the residuals, but will tinker with the GARCH parameters in order to try and create a model where each AR,Q and P are significant. Notice that the heteroscedastic properties of the returns (as indicated by the Q and LM tests) are different than in the series themselves.

6.6.3 GARCH Model of VLCC Returns

GARCH Estimates

SSE	298373395	Observations	285
MSE	1046924	Uncond Var	1277677.24
Log Likelihood	-2372.3872	Total R-Square	0.1791
SBC	4789.99432	AIC	4760.7744
Normality Test	147.4617	Pr > ChiSq	<.0001

Variable	DF	Estimate	Standard Error	t Value	Approx Pr > |t|
Intercept	1	28.3607	416.7211	0.07	0.9457
date	1	-0.003192	0.0347	-0.09	0.9268
AR1	1	-0.3735	0.0821	-4.55	<.0001
AR6	1	-0.0918	0.0498	-1.84	0.0652
AR8	1	0.0982	0.0573	1.71	0.0865
AR12	1	0.2080	0.0577	3.60	0.0003
ARCH0	1	1059293	0.3744	2829426	<.0001
ARCH1	1	0.1709	0.0807	2.12	0.0342
GARCH1	1	-4.65E-24	1.6021E-9	-0.00	1.0000

8.41

6.6.4 GARCH MODEL OF AFRAMAX RETURNS

GARCH Estimates

SSE	137700917	Observations	285
MSE	483161	Uncond Var	632269.043
Log Likelihood	-2255.7717	Total R-Square	0.1470
SBC	4562.41584	AIC	4529.54344
Normality Test	156.8327	Pr > ChiSq	<.0001

Variable	DF	Estimate	Standard Error	t Value	Approx Pr > |t|
Intercept	1	12.8673	277.0704	0.05	0.9630
date	1	0.000873	0.0222	0.04	0.9687
AR1	1	-0.3025	0.0675	-4.48	<.0001
AR6	1	-0.0986	0.0818	-1.21	0.2279
AR9	1	0.1246	0.0675	1.85	0.0650
AR13	1	0.0800	0.0730	1.10	0.2730
ARCH0	1	486228	0.1037	4688605	<.0001
ARCH1	1	0.002469	0.0449	0.06	0.9561
ARCH2	1	0.2285	0.0814	2.81	0.0050
GARCH1	1	-1.26E-10	3.9829E-9	-0.03	0.9747

8.42

6.6.5 GARCH MODEL OF HANDYSIZE RETURNS

GARCH Estimates

SSE	39390470.9	Observations	285
MSE	138212	Uncond Var	167277.34
Log Likelihood	-2086.0987	Total R-Square	0.1741
SBC	4206.11236	AIC	4184.19742
Normality Test	118.8557	Pr > ChiSq	<.0001

Variable	DF	Estimate	Standard Error	t Value	Approx Pr > |t|
Intercept	1	-263.8236	185.6564	-1.42	0.1553
date	1	0.0240	0.0168	1.43	0.1538
AR1	1	-0.3378	0.0964	-3.51	0.0005
AR4	1	-0.1220	0.0590	-2.07	0.0387
ARCH0	1	139389	0.7563	184299	<.0001
ARCH1	1	0.1667	0.0563	2.96	0.0031
GARCH1	1	5.666E-23	8.326E-8	0.00	1.0000

8.43

6.6.6 GARCH MODEL OF SUEZMAX RETURNS

GARCH Estimates

SSE	243286284	Observations	285
MSE	853636	Uncond Var	1322127.75
Log Likelihood	-2340.7482	Total R-Square	0.2273
SBC	4726.71633	AIC	4697.49641
Normality Test	2819.8213	Pr > ChiSq	<.0001

Variable	DF	Estimate	Standard Error	t Value	Approx Pr > |t|
Intercept	1	-66.0773	383.2728	-0.17	0.8631
date	1	0.006669	0.0293	0.23	0.8199
AR1	1	-0.2592	0.0940	-2.76	0.0058
AR3	1	-0.1594	0.0453	-3.52	0.0004
AR12	1	0.2497	0.0501	4.98	<.0001
AR13	1	0.2310	0.0550	4.20	<.0001
ARCH0	1	859100	0.3997	2149404	<.0001
ARCH1	1	0.3502	0.1251	2.80	0.0051
GARCH1	1	7.321E-23	7.6442E-9	0.00	1.0000

8.44

6.7 FORECASTING

6.7.1 FORECASTING VLCC RETURNS WITH GARCH MODELS

Having decided on the model, we can now forecast the GARCH models for each of the variables (VLCC_R, AFRAMAX_R, HANDYSIZE_R, SUEZMAX_R) in PROC AUTOREG. The output is the same as in the previous commands and will not be shown. The only difference is found in the *output* command where the residuals (r = <name>_r_resid), predictions (p = <name>_r_pred), upper (ucl) and lower (lcl) confidence limits are specified and added to the output table.

The first step in predicting and evaluating the series is to create the output file *out1* from which the predicted values and confidence limits can be exported. In this case the output in the routines is excluded since it is the same as in the routines specified previously (without the 'output' line).

Then, having run the specified GARCH model, we can print the time charter rate returns and the fitted model to get an idea about how well the model has captured volatile regions in the series.

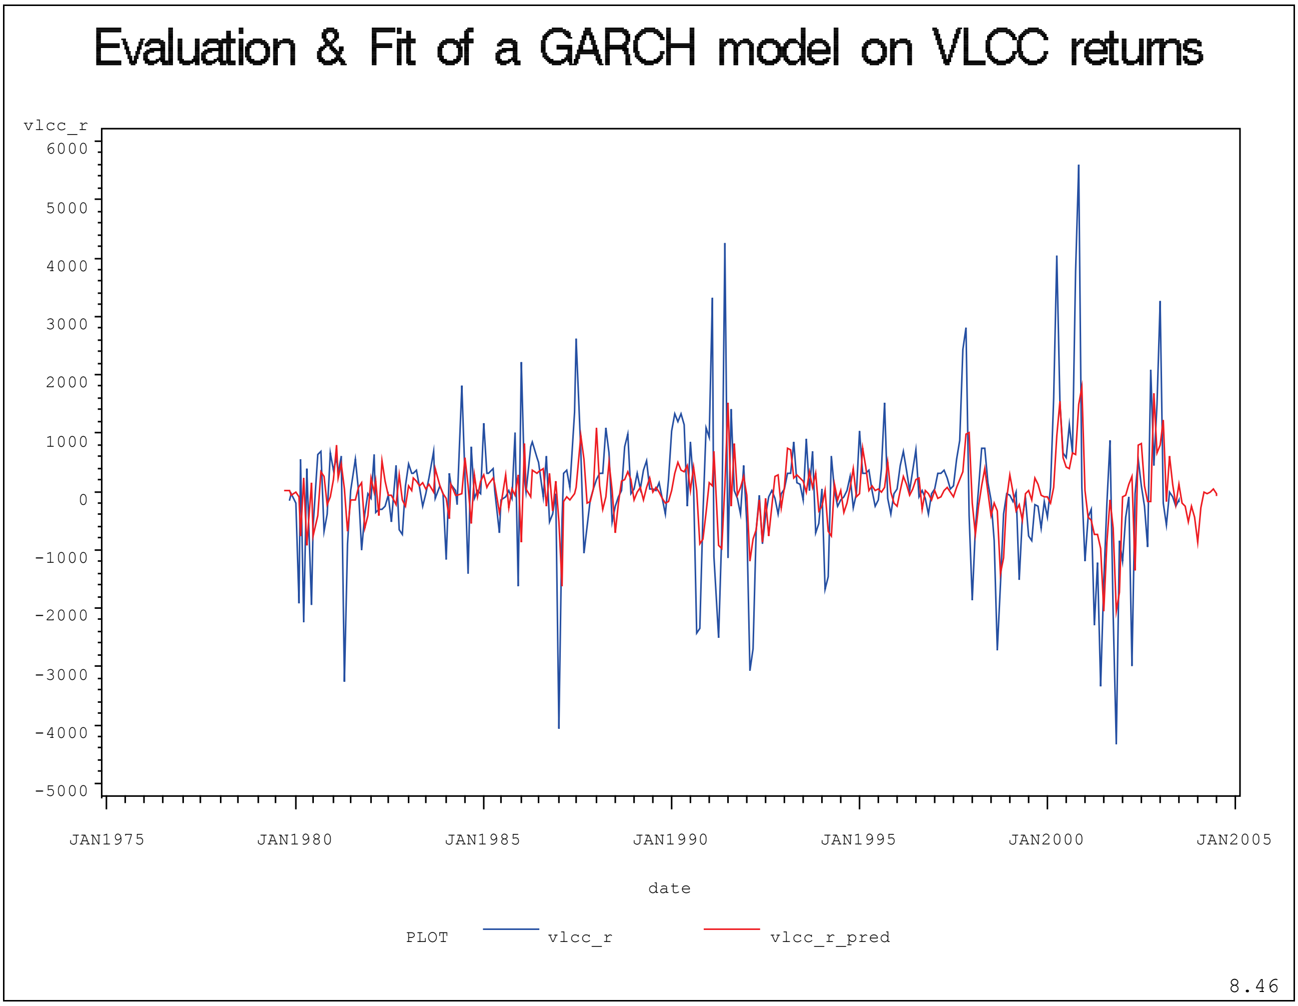

Figure 85 : Evaluation and Fit of a GARCH model on VLCC Returns

We then zoom into the end of the series, creating a reference line when our data ends. To the right of the reference line is the prediction of the time charter rate for the next 12 months.

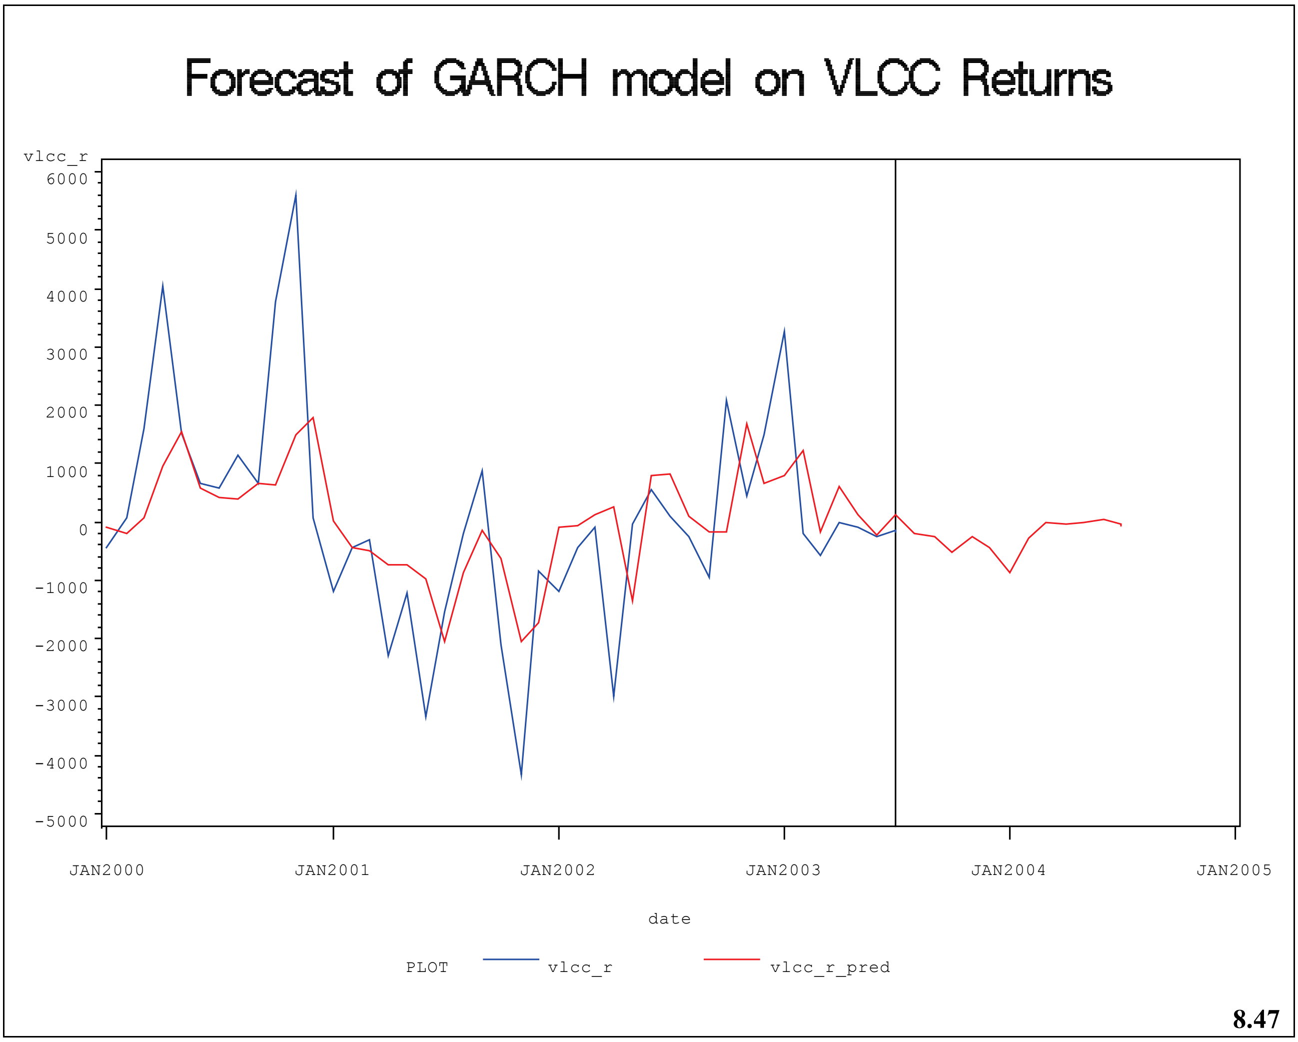

Figure 86 : Forecast of GARCH model on VLCC Returns

SAS can export the result file into an Excel worksheet for further data analysis. The table below shows the forecasted returns for the forecast period from July 2003 to July 2004. The series w_t SAS log transformation is:

	VLCC FORECAST	
DATE	FORECAST	CONVERTED
8/1/2003	-188.636655	15914
9/1/2003	-264.00673	15628
10/1/2003	-531.882387	15204
11/1/2003	-261.695589	15341
12/1/2003	-447.796825	14938
1/1/2004	-868.045736	14023
2/1/2004	-292.800269	13413
3/1/2004	-17.8070703	13097
4/1/2004	-51.2339905	12695
5/1/2004	-14.9322172	12659
6/1/2004	37.9795113	12947
7/1/2004	-27.6374115	13072

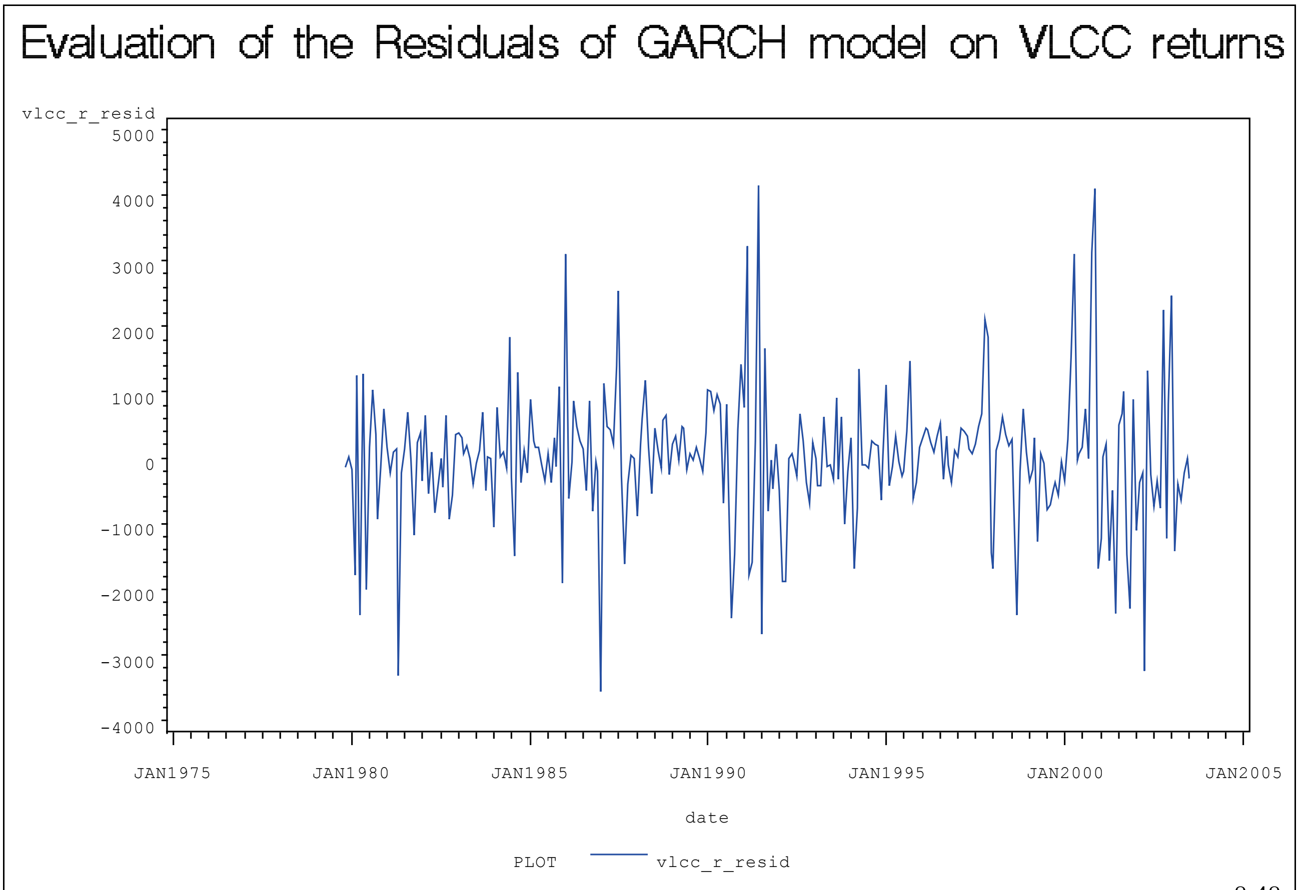

8.48

Figure 87 : Evaluation of Residuals of GARCH Model on VLCC Returns

We can use the identify statement to check for autocorrelation of the residuals. If the residuals are truly white noise then the autocorrelation function should have no spikes. In this case, the residuals are all within the standard error. Using (8.48) we obtain

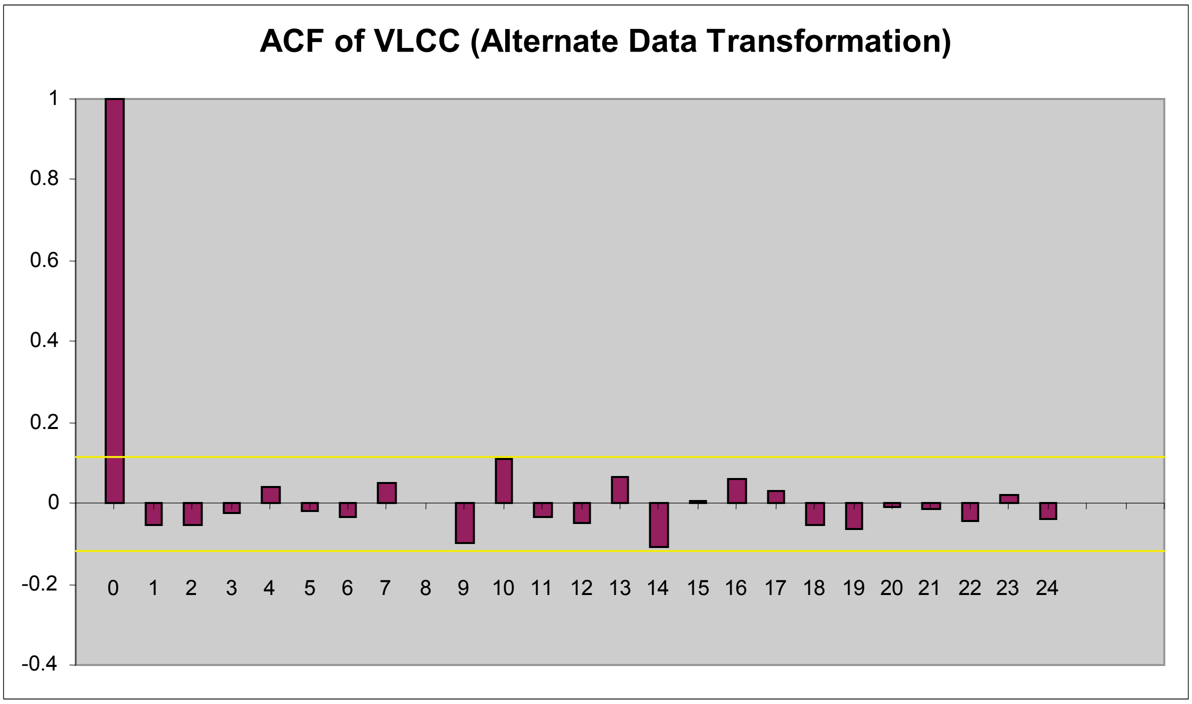

Autocorrelation Check for White Noise

To Lag	Chi-Square	DF	Pr > ChiSq	--------------------Autocorrelations--------------------					
6	2.86	6	0.8260	-0.056	-0.055	-0.026	0.040	-0.017	-0.033
12	10.89	12	0.5387	0.049	-0.000	-0.098	0.107	-0.033	-0.049
18	17.71	18	0.4752	0.063	-0.106	0.003	0.059	0.032	-0.054
24	20.34	24	0.6771	-0.066	-0.007	-0.016	-0.044	0.022	-0.039

The identify statement shows that the residuals are uncorrelated white noise which shows that the model, in theory, has been correctly fitted to the data.

FORECASTING AFRAMAX RETURNS WITH GARCH MODELS

As with the VLCC time charter rate, the first step in predicting and evaluating the series is to create the output file from which the predicted values and confidence limits can be exported. In this case the ouput is excluded since it is the same as in the routines specified previously (without the 'output' line). Having run the specified GARCH model, we can print the time charter rate returns and the fitted model to get an idea about how well the model has captured volatile regions in the series.

The actual and fitted returns are shown superimposed over each other.

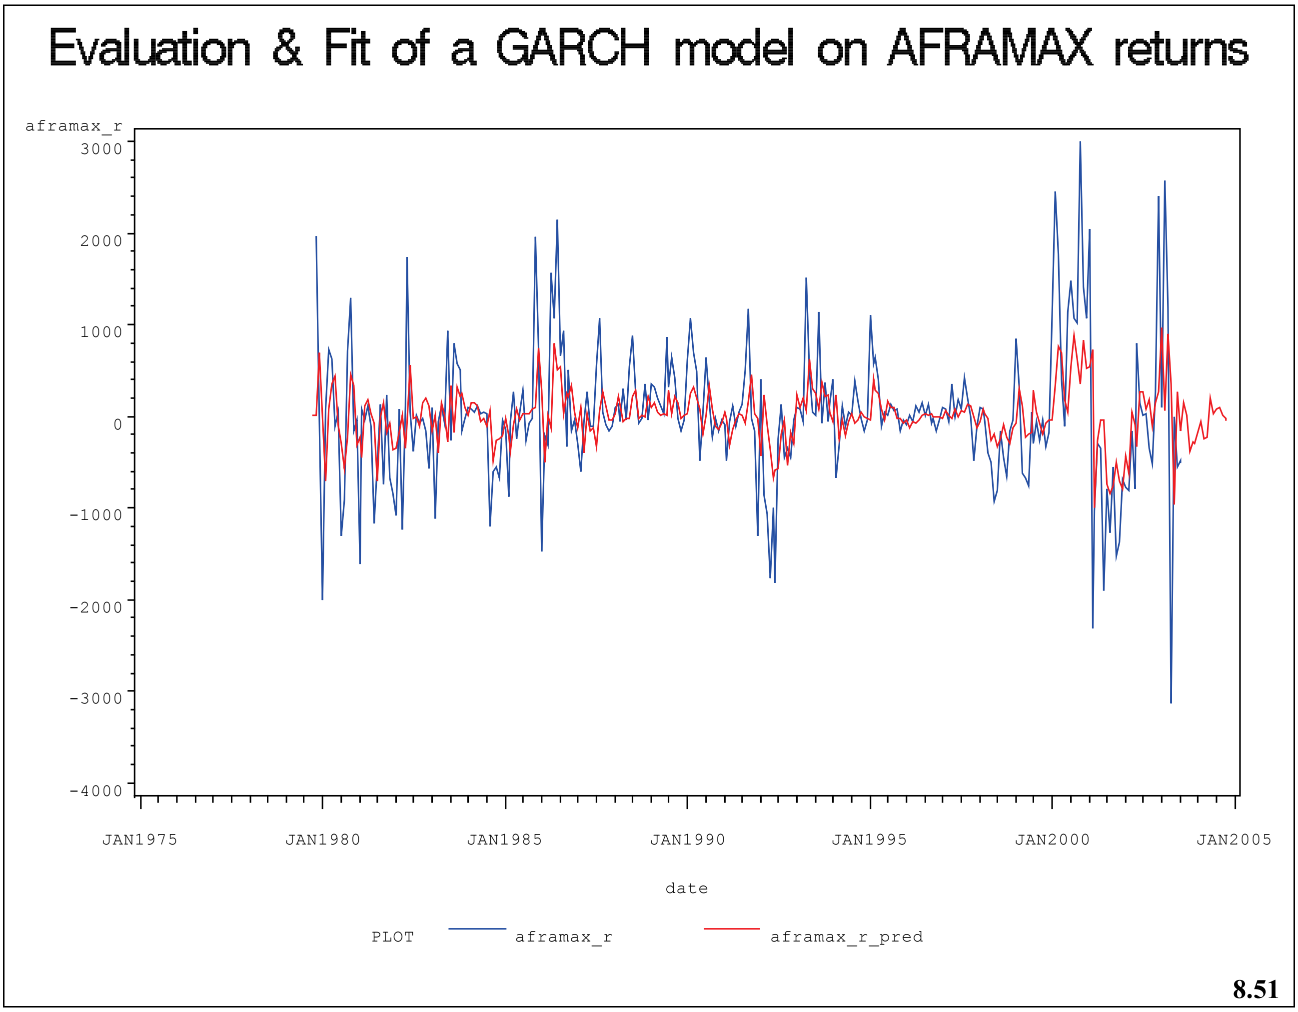

Figure 88 : Evaluation & Fit of a GARCH model on AFRAMAX Returns

We can then zoom to the end of the model data and see the model prediction (in red). The perpendicular line running along the Y axis indicates the end of the known data and the beginning of the forecasted period.

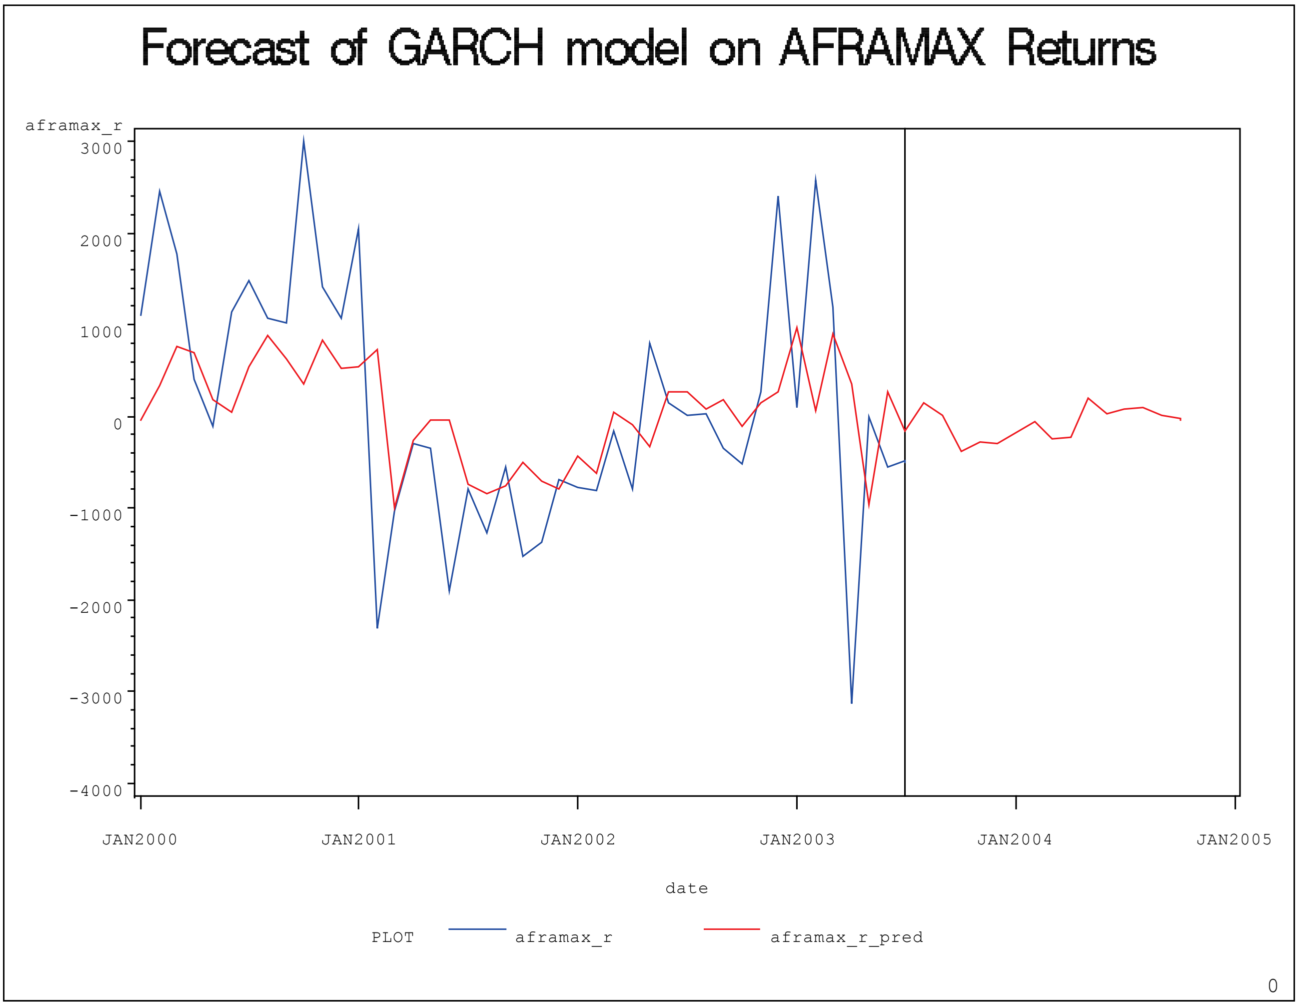

Figure 89 : Forecast of GARCH model on AFRAMAX Returns

The forecasted values for the AFRAMAX time charter rate are shown below:

DATE	FORECAST	FORECAST CONVERTED
8/1/2003	150.636067	17515
9/1/2003	19.1233408	17613
10/1/2003	-379.68342	17259
11/1/2003	-274.776765	17139
12/1/2003	-289.404329	16896
1/1/2004	-172.953913	16622
2/1/2004	-49.6336229	16502
3/1/2004	-241.757173	16317
4/1/2004	-220.660575	15988
5/1/2004	207.246267	16298
6/1/2004	207.246267	16290
7/1/2004	31.4603659	16356

Evaluation of AFRAMAX Residuals

We now print and evaluate the AFRAMAX residuals.

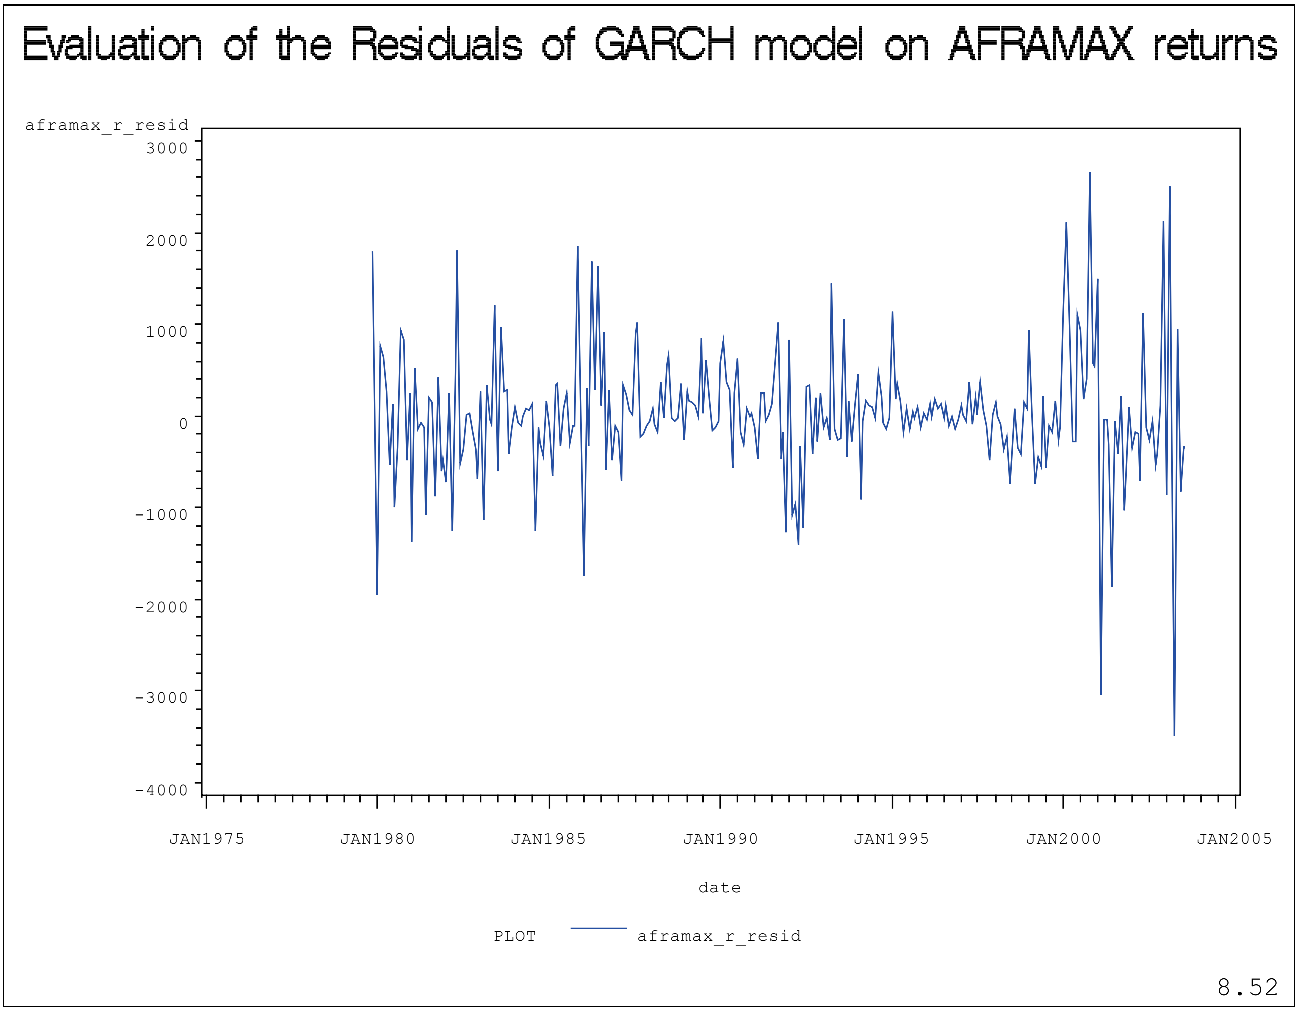

Figure 90 : Evaluation of the Residuals of GARCH model on AFRAMAX Returns

Using (8.54) we obtain

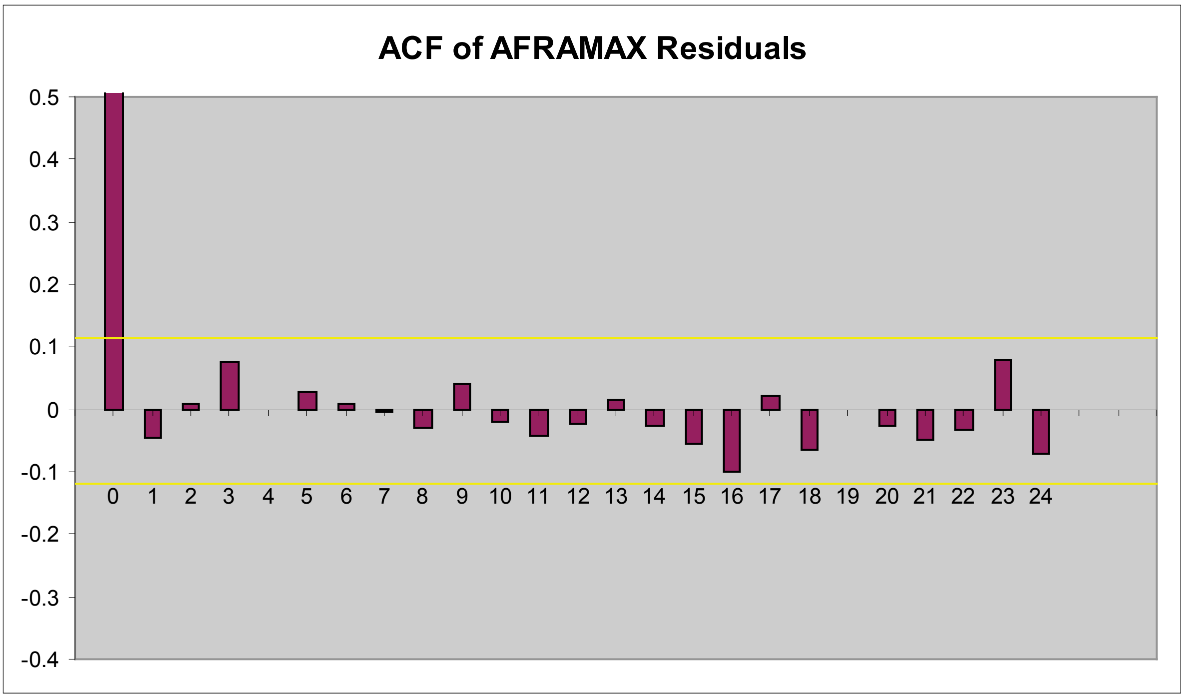

Figure 91 : ACE of AFRAMAX Residuals

Autocorrelation Check for White Noise

To Lag	Chi-Square	DF	Pr > ChiSq	--------------------Autocorrelations--------------------					
6	2.44	6	0.8747	-0.045	0.010	0.075	-0.001	0.027	0.008
12	4.01	12	0.9833	-0.005	-0.031	0.039	-0.019	-0.043	-0.023
18	9.68	18	0.9418	0.015	-0.026	-0.057	-0.099	0.022	-0.065
24	14.50	24	0.9345	-0.001	-0.025	-0.049	-0.032	0.080	-0.071

The identify statement shows that the residuals are uncorrelated white noise which shows that the model has been correctly fitted to the data.

6.7.2 FORECASTING HANDYSIZE RETURNS WITH GARCH MODELS

Using (8.55) we obtain the actual and fitted returns are shown superimposed over each other::

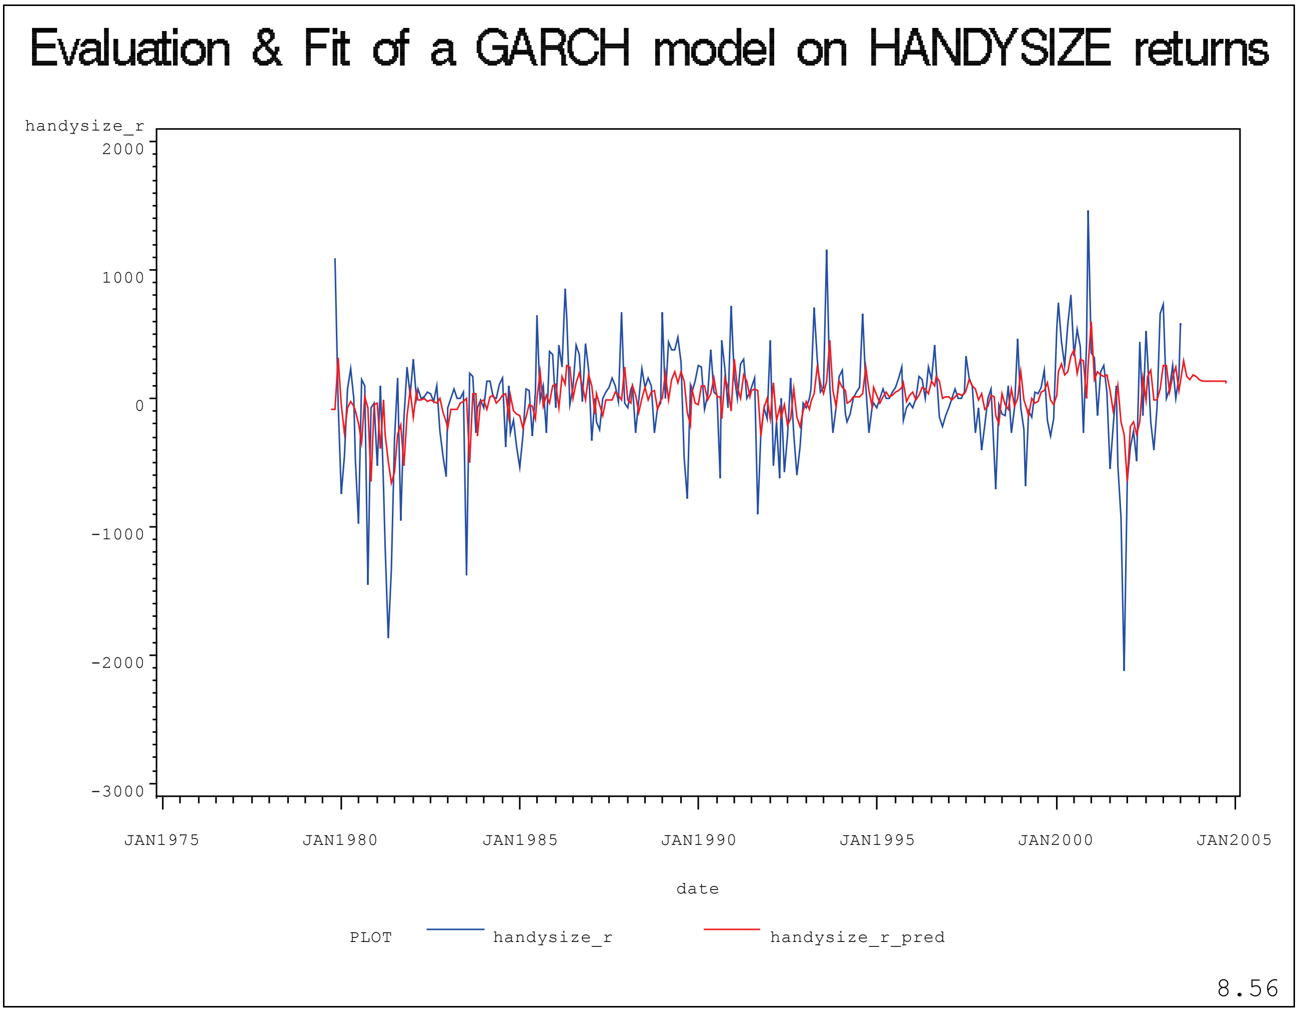

Figure 92 : Evaluation and Fit of a GARCH model on HANDYSIZE Returns

We can then zoom to the end of the model data and see the model prediction (in red). The perpendicular line running along the Y axis indicates the end of the known data and the beginning of the forecasted period.

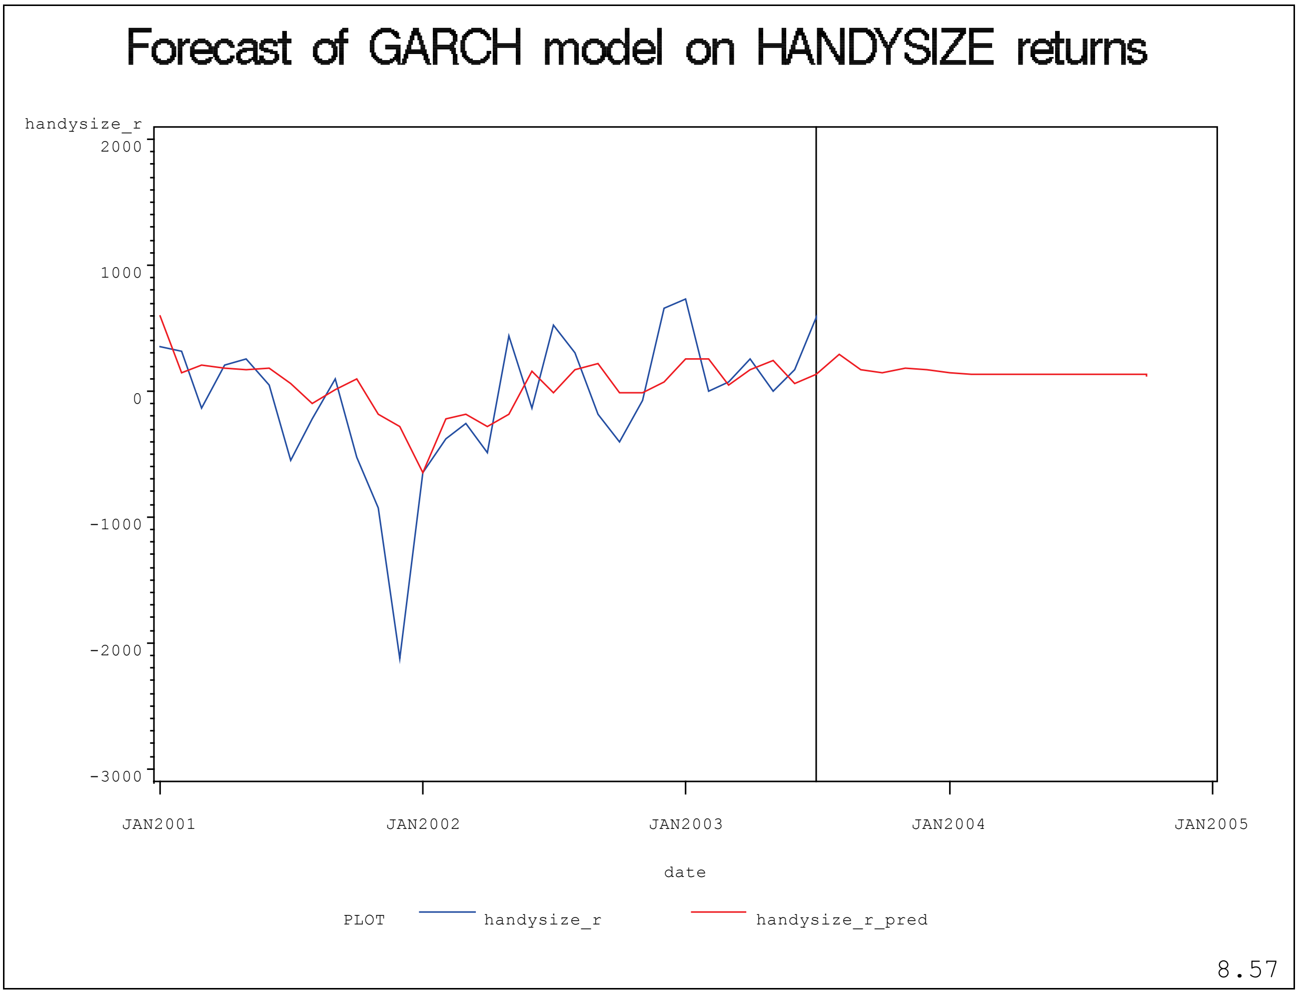

Figure 93 : Forecast of GARCH model on HANDYSIZE Returns

The forecasted values for the HANDYSIZE time charter rate are shown below

DATE	FORECAST	FORECAST CONVERTED
8/1/2003	291.491299	13756
9/1/2003	163.995149	13825
10/1/2003	141.955715	14240
11/1/2003	184.535873	14499
12/1/2003	163.984371	14695
1/1/2004	141.902647	14916
2/1/2004	132.16029	15049
3/1/2004	134.416325	15116
4/1/2004	133.093162	15248
5/1/2004	130.331136	15373

We now print and evaluate the HANDYSIZE residuals.

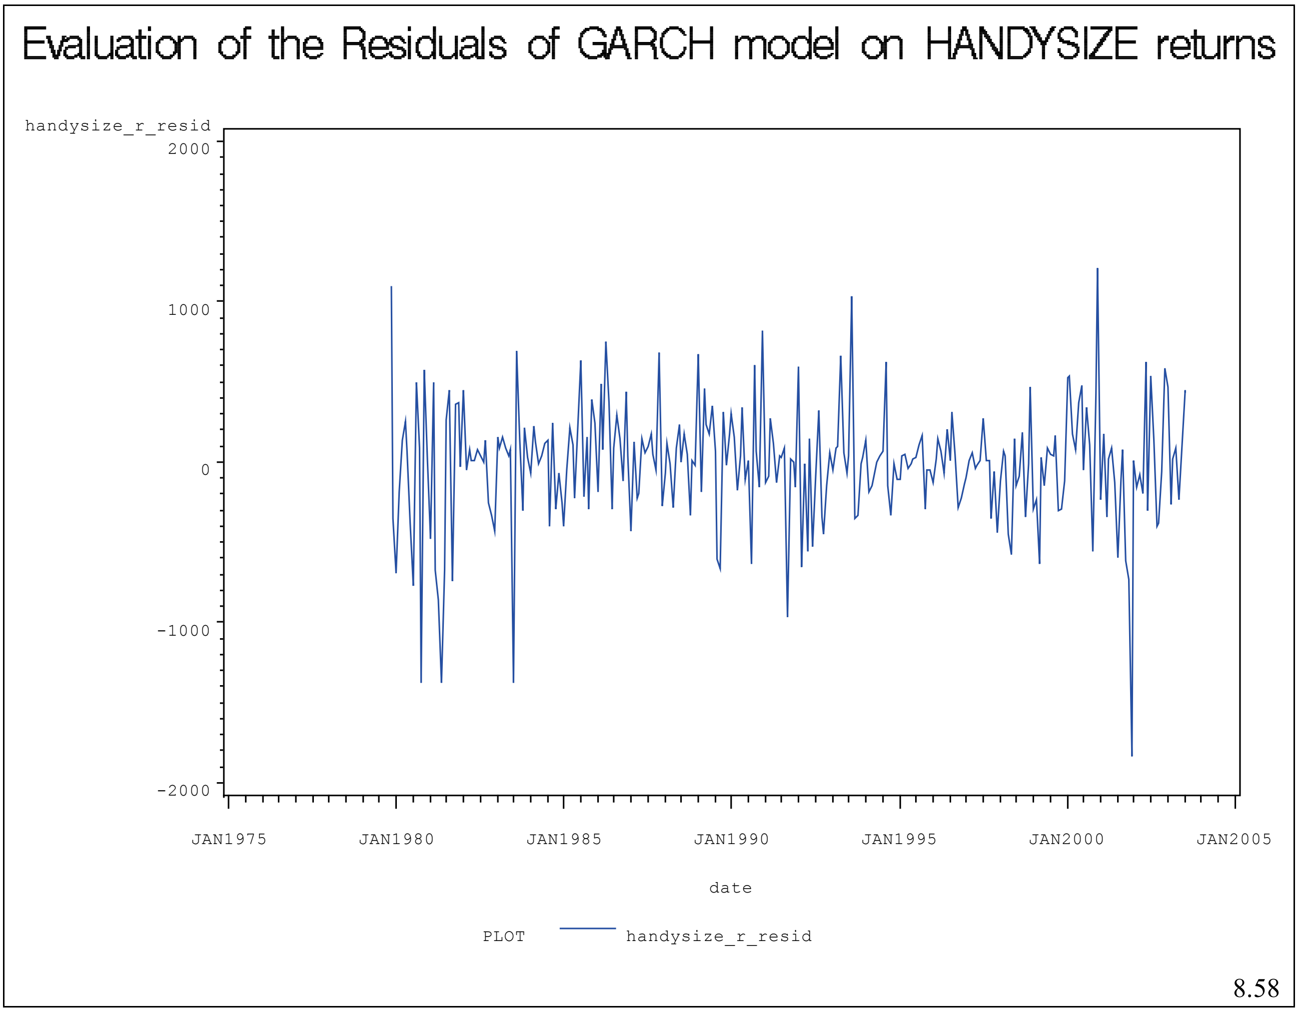

Figure 94 : Evaluation of the Residuals of GARCH model on HANDYSIZE Returns

Using (8.59) we obtain:

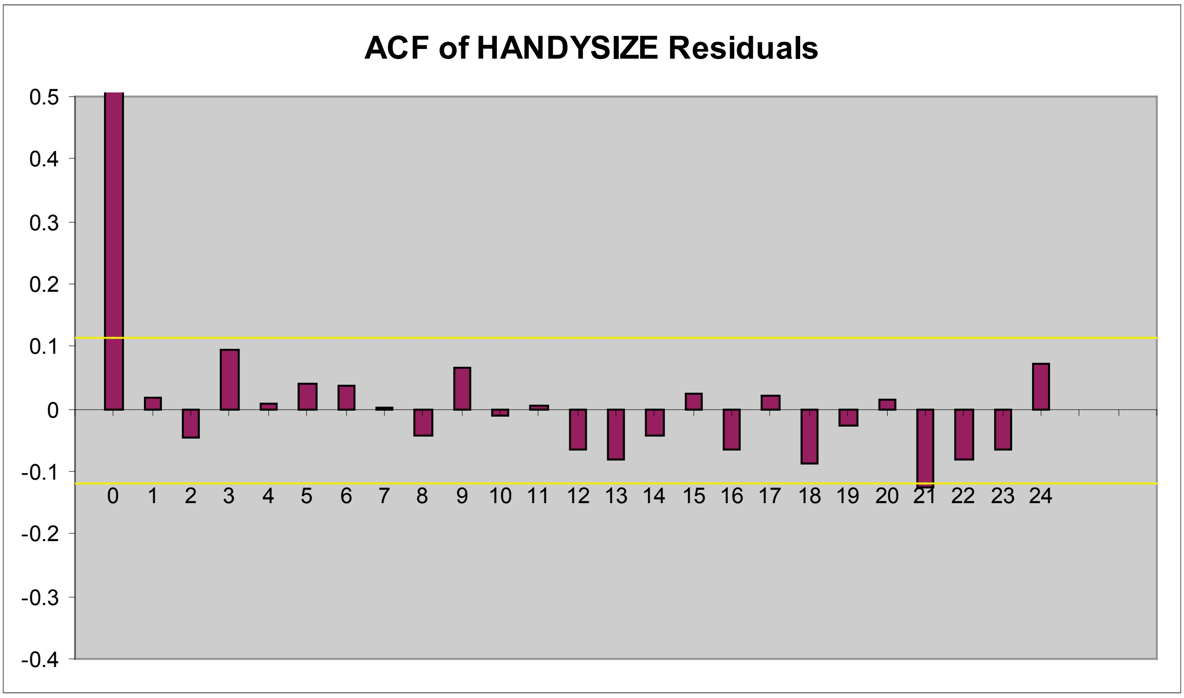

Figure 95 : ACE of HANDYSIZE Residuals

Autocorrelation Check for White Noise

To Lag	Chi-Square	DF	Pr > ChiSq	Autocorrelations					
6	4.28	6	0.6387	0.017	-0.046	0.096	0.008	0.041	0.037
12	7.35	12	0.8336	0.003	-0.042	0.065	-0.011	0.005	-0.065
18	13.75	18	0.7451	-0.080	-0.042	0.025	-0.066	0.022	-0.086
24	23.70	24	0.4788	-0.026	0.013	**-0.124**	-0.080	-0.065	0.071

The identify statement shows that the residuals are uncorrelated white noise which shows that the model has been correctly fitted to the data although there does seem to a some borderline correlation in larger lags, especially at lag 21 (in bold). This may be be simply an outlier variable.

6.7.3 FORECASTING SUEZMAX RETURNS WITH GARCH

The actual and fitted returns are shown superimposed over each other.

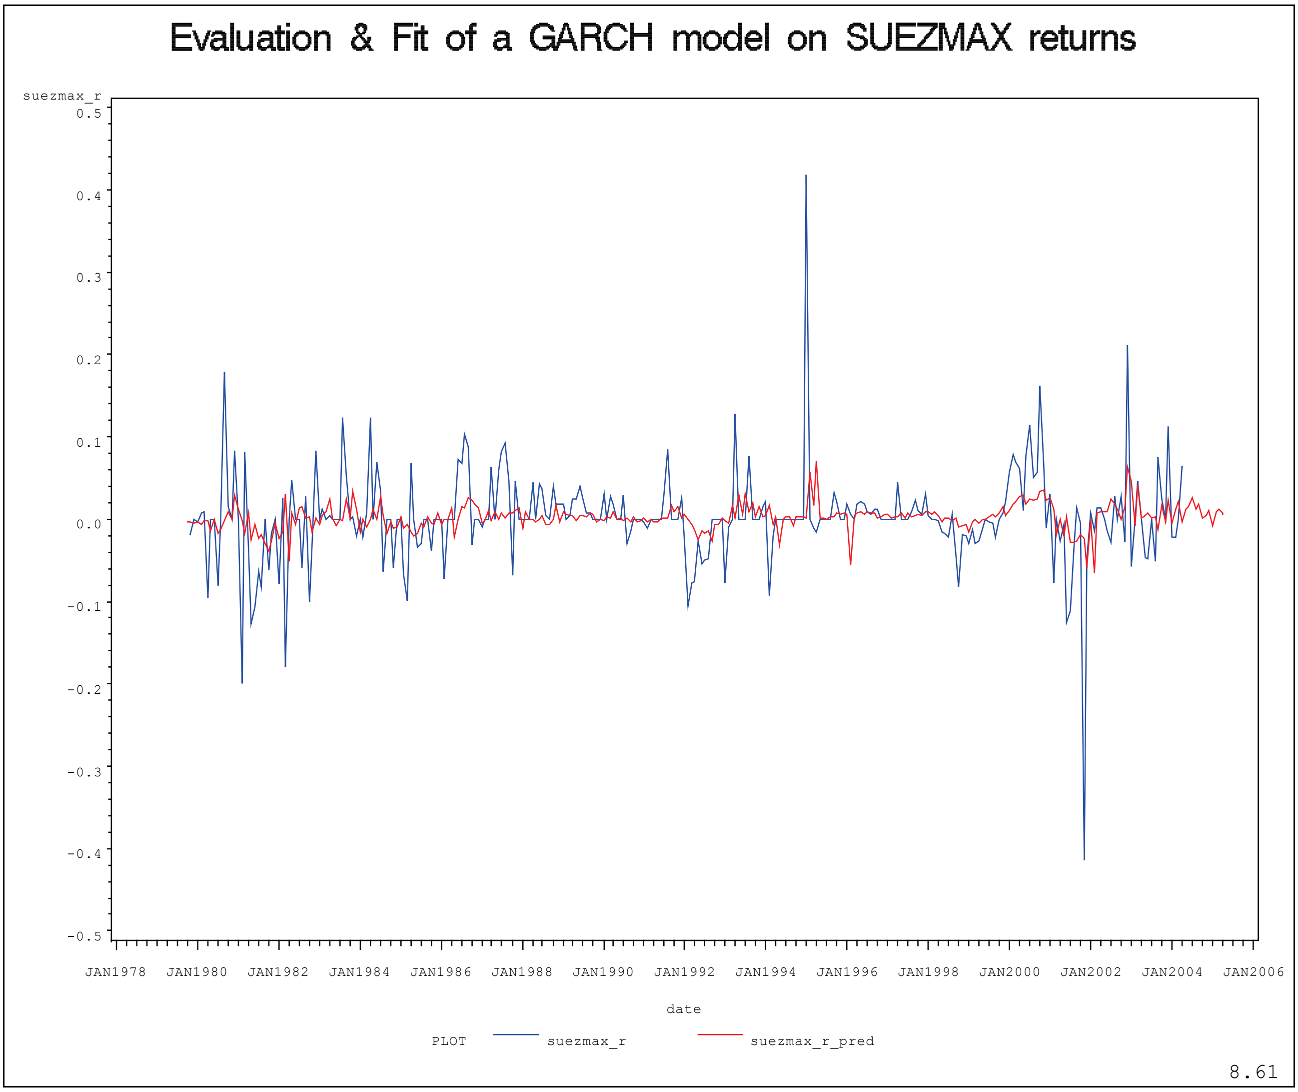

Figure 96 : Evaluation and Fit of a GARCH Model on SUEZMAX Returns

We can then zoom to the end of the model data and see the model prediction (in red). The perpendicular line running along the Y axis indicates the end of the known data and the beginning of the forecasted period.

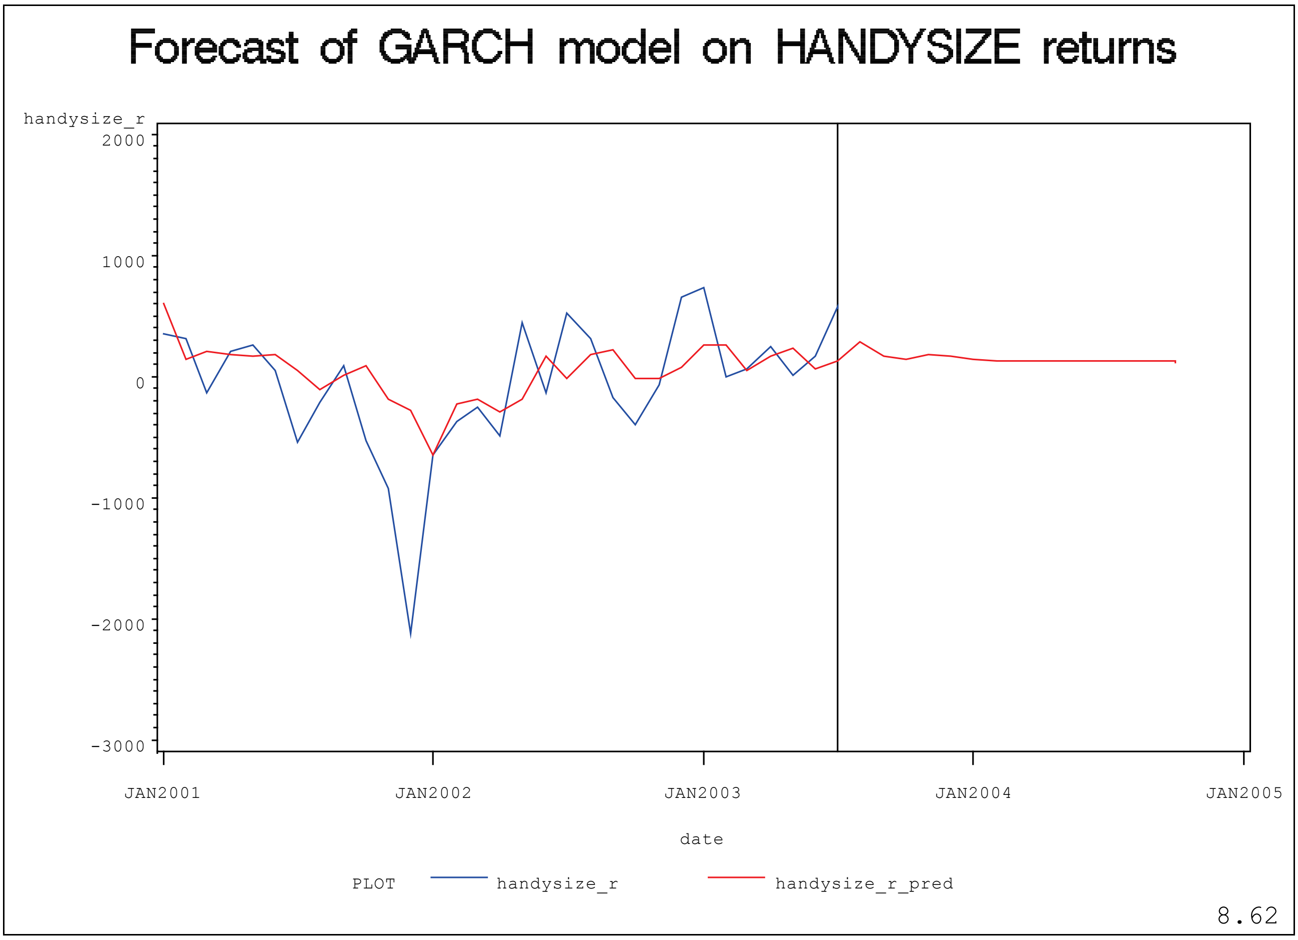

Figure 97 : Forecast of GARCH Model on HANDYSIZE Returns

The forecasted values for the HANDYSIZE time charter rate are shown below:

DATE	FORECAST RETURNS	FORECAST
8/1/2003	291.491	23000
9/1/2003	163.995	22512
10/1/2003	141.956	22759
11/1/2003	184.536	22512
12/1/2003	163.984	21815
1/1/2004	141.903	20433
2/1/2004	132.160	18739
3/1/2004	134.416	17019
4/1/2004	133.093	16061
5/1/2004	130.331	16186
6/1/2004	128.620	16581
7/1/2004	128.701	16538

We now print and evaluate the SUEZMAX residuals.

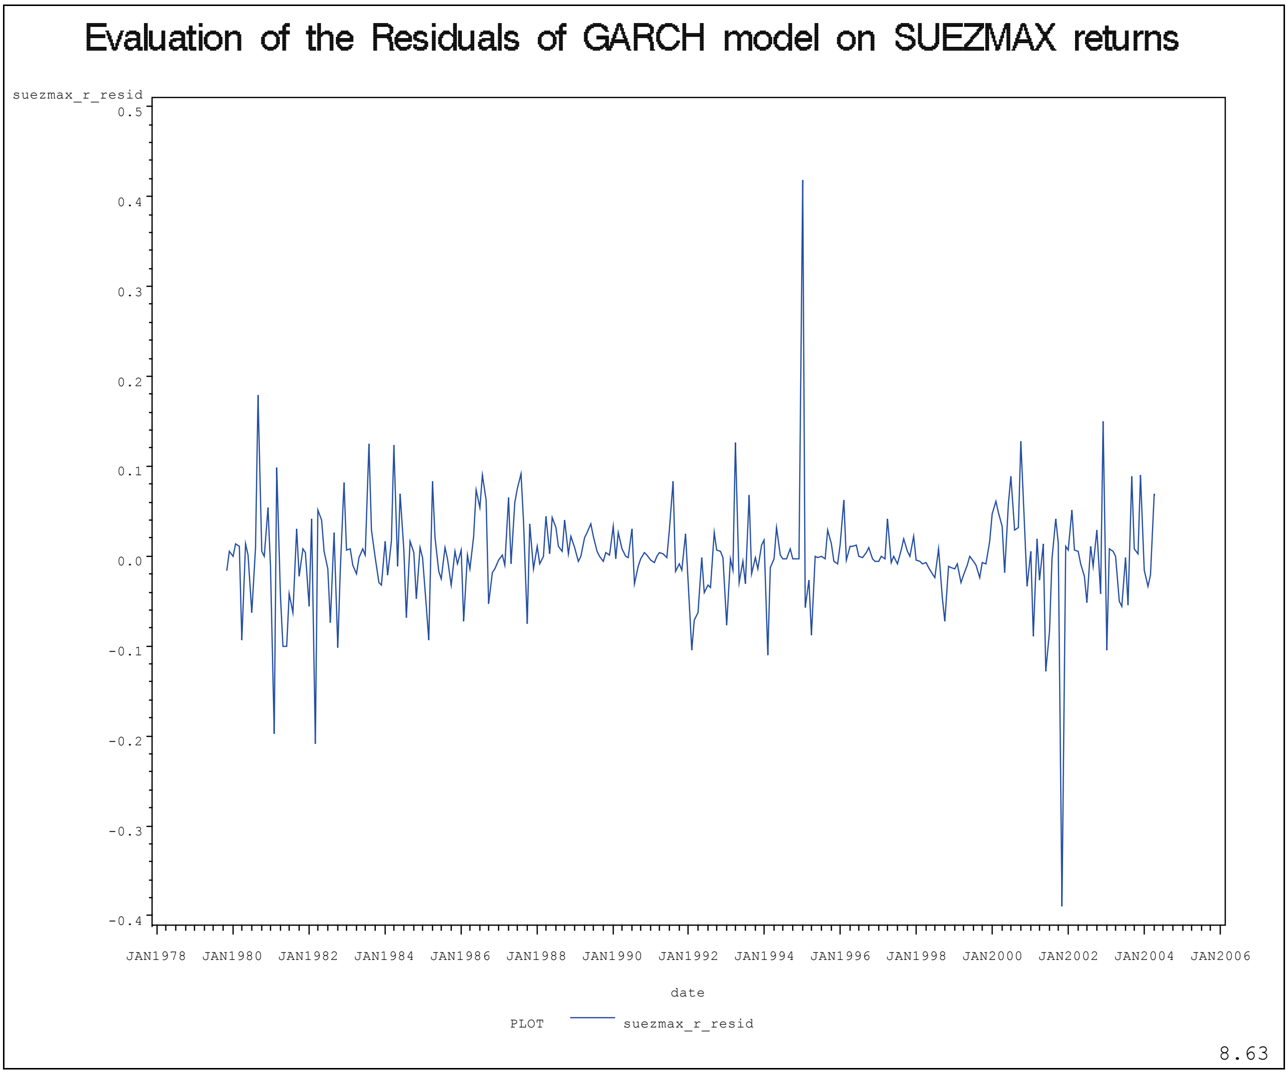

Figure 98 : Evaluation of the Residuals of GARCH Model on SUEZMAX Returns

Using (8.64) we obtain:

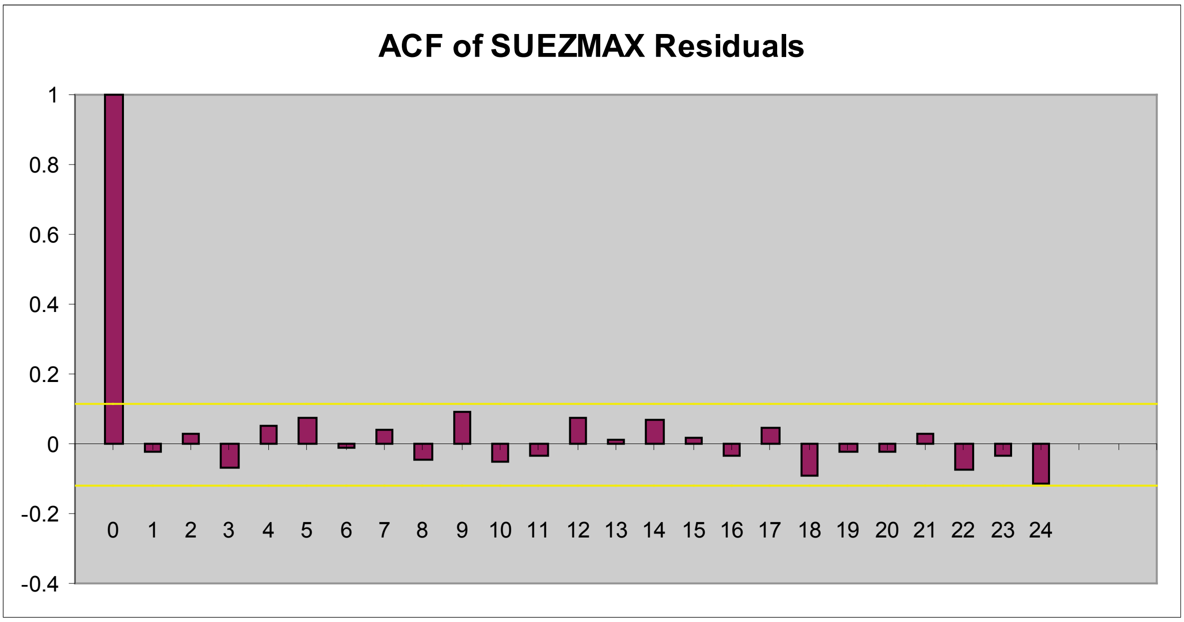

Figure 99 : ACE of SUEZMAX Residuals

Autocorrelation Check for White Noise

To Lag	Chi-Square	DF	Pr > ChiSq	Autocorrelations					
6	4.09	6	0.6641	-0.021	0.027	-0.069	0.050	0.073	-0.014
12	10.46	12	0.5756	0.037	-0.048	0.090	-0.051	-0.035	0.077
18	15.61	18	0.6200	0.011	0.067	0.017	-0.035	0.046	-0.093
24	22.48	24	0.5509	-0.022	-0.022	0.030	-0.074	-0.032	-0.117

The identify statement shows that the residuals are uncorrelated white noise which shows that the model has been correctly fitted to the data.

6.7.4 **ARMAX GARCH Model:**

Using the alternate data transformation, the cross correlation plots of each explanatory variable show a maximum when compared to the VLCC time charter rate:

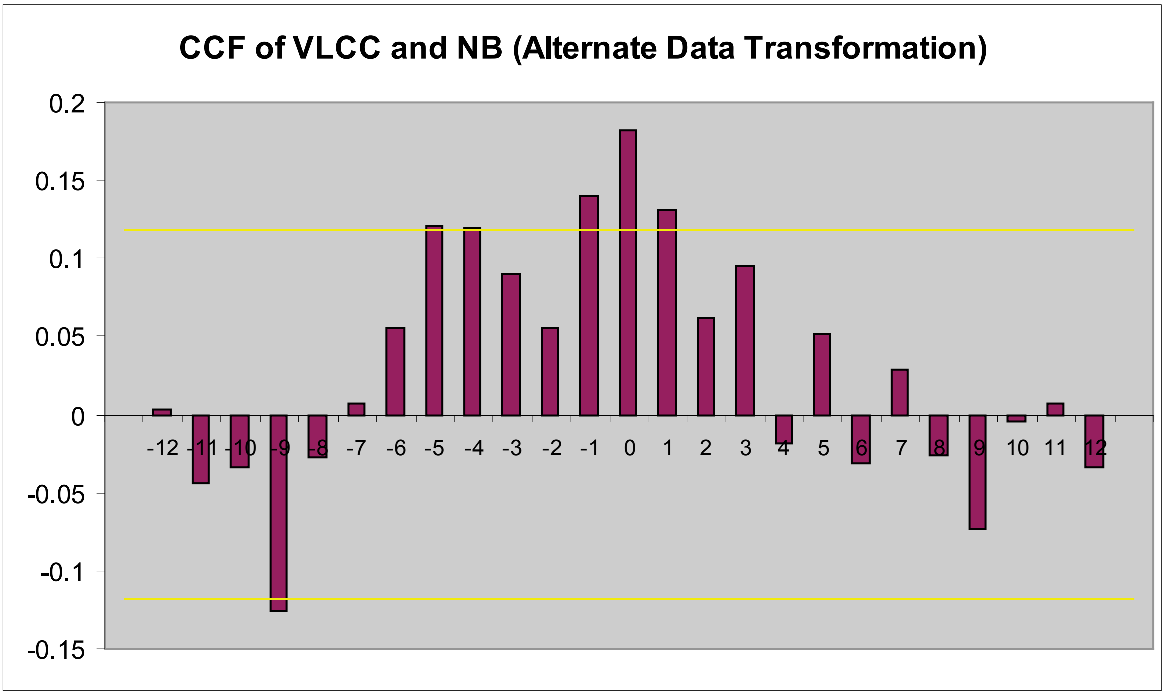

Figure 100: CCF of VLCC and NB (Alternate Data Transformation)

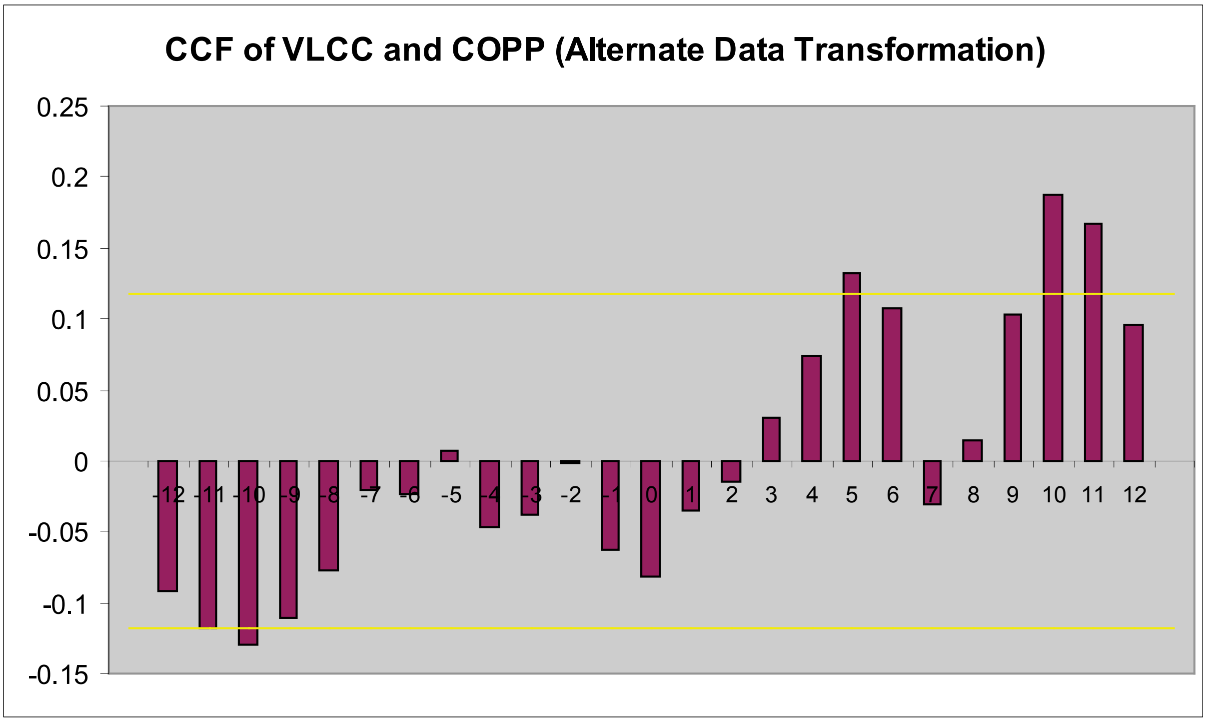

Figure 101 : CCF of VLCC and COPP (Alternate Data Transformation)

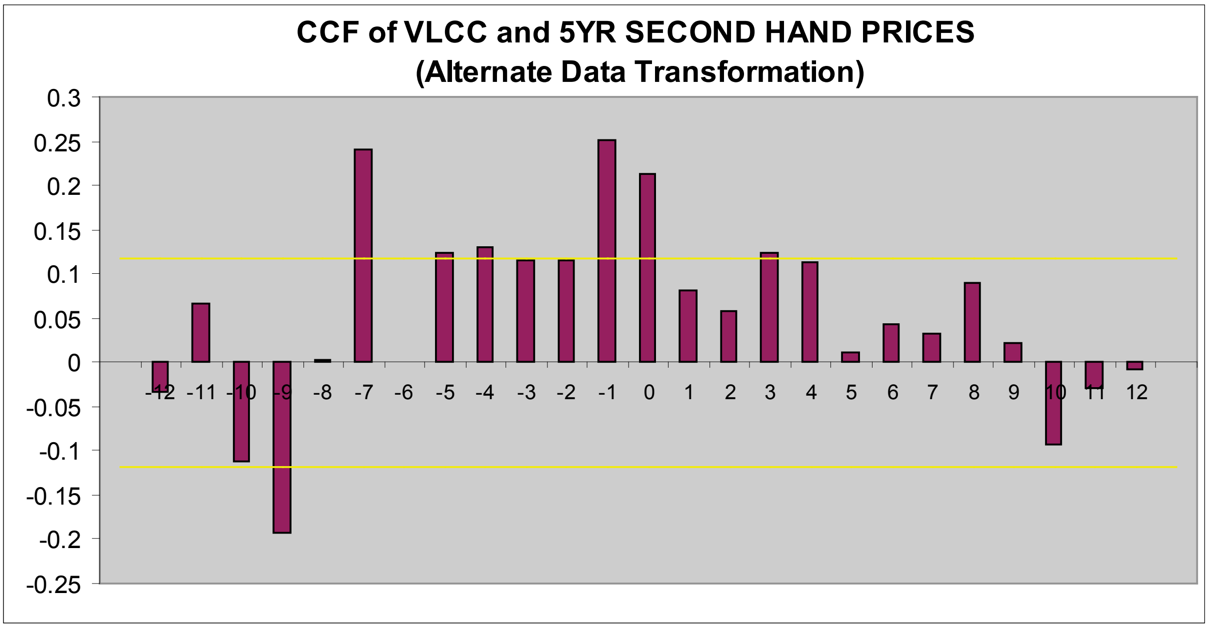

Figure 102 : CCF of VLCC and 5YR SECOND HAND PRICES

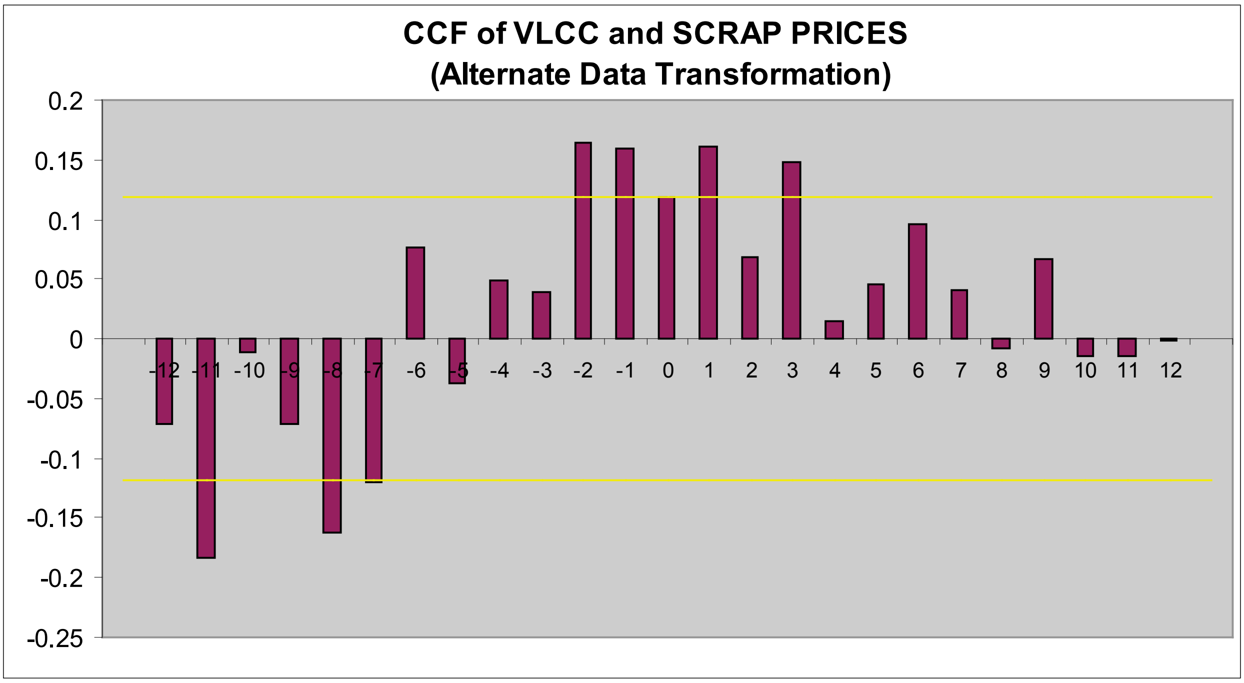

Figure 103 : CCF of VLCC and Scrap Prices

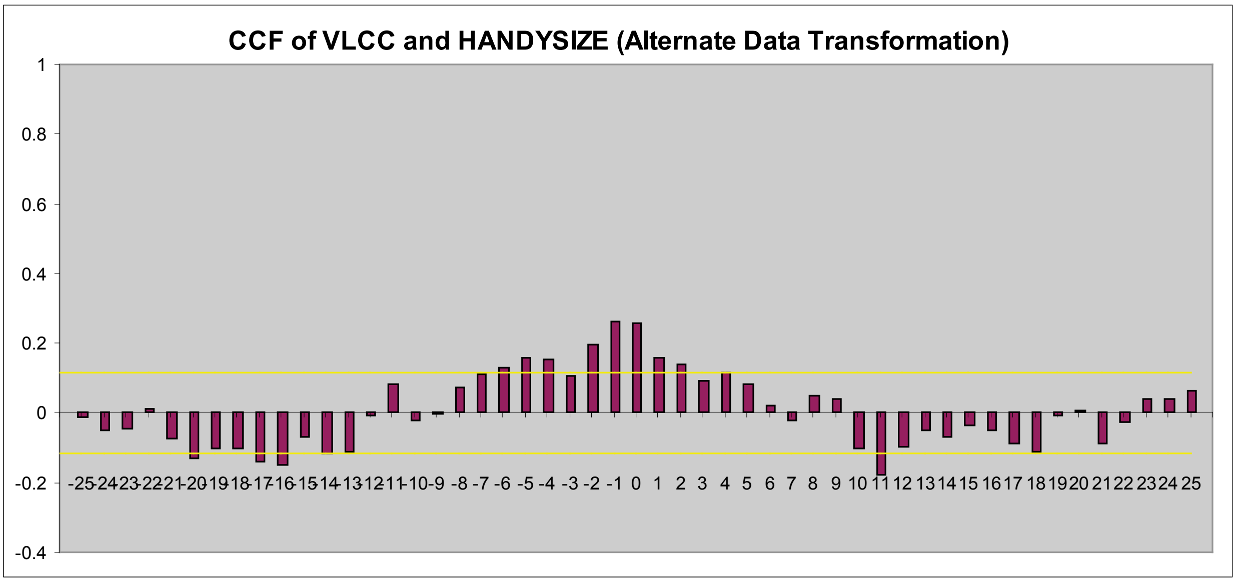

Figure 104 : CCF of VLCC and HANDYSIZE (Alternate Data Transformation)

In order to have enough lagged variables to be able to create a long range forecast, we must choose a relatively large lag combined with a high cross correlation value. For the model pecification we will use only the Crude Oil Purchase price lagged 10 periods.

We eliminate the New Building Prices, 5 Year Second Hand Prices because the CCF shows they are correlated with negative lags, that is, the VLCC TCR influences *them* rather than vice versa.

We also eliminate the Handysize Time Charter Rate from the model because we would like to have a true *causal model.* A causal model can be identified from the CCF by the absence of significant negative lags. The existence of both positive and negative lags leads us to infer that time charter rates may be cyclically influencing each other and may hinder model development.

Durbin Watson and Arch Tests

GARCH Estimates

SSE	274897735	Observations	275
MSE	999628	Uncond Var	1216281.93
Log Likelihood	-2281.3593	Total R-Square	0.2135
SBC	4607.65282	AIC	4578.71865
Normality Test	142.8265	Pr > ChiSq	<.0001

Variable	DF	Estimate	Standard Error	t Value	Approx Pr > |t|
Intercept	1	21.2882	97.5461	0.22	0.8272
copplag10	1	1.3523	0.3570	3.79	0.0002
AR1	1	-0.3767	0.0758	-4.97	<.0001
AR10	1	-0.1158	0.0677	-1.71	0.0872
AR12	1	0.1954	0.0516	3.78	0.0002
AR18	1	0.1027	0.0621	1.65	0.0982
ARCH0	1	1012140	0.0189	5.349E7	<.0001
ARCH1	1	0.1678	0.0764	2.20	0.0281
GARCH1	1	-7.24E-24	2.0819E-9	-0.00	1.0000

8.65

We can see from the significance test that we have a good model. Lags one and twelve, the Crude Oil Purchase price and the ARCH disturbances are all significant. Lags ten and twelve are also relatively significant and have been left as they include valuable information that the model uses. Using (8.66) we obtain:

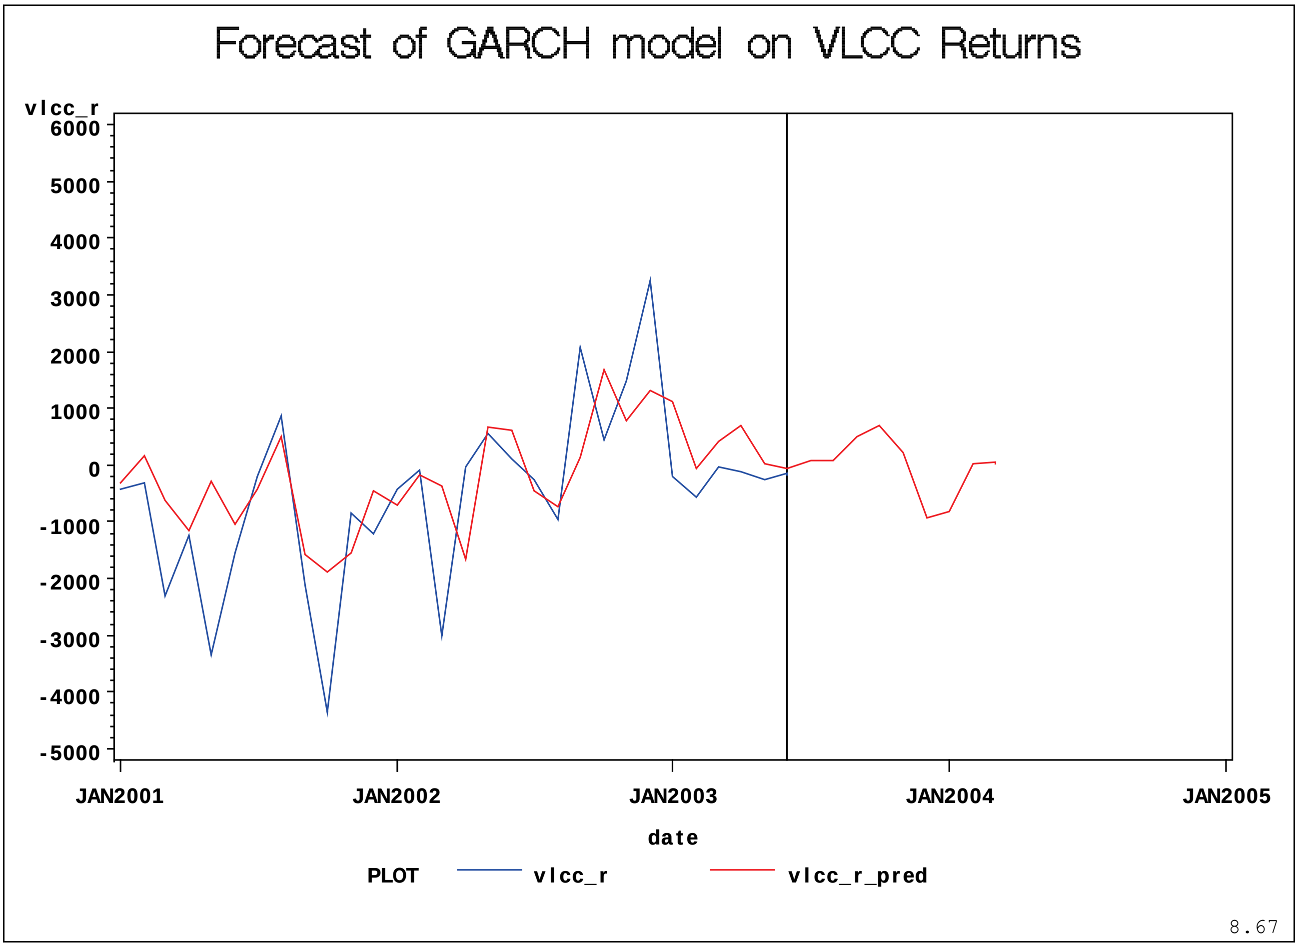

Figure 105 : GARCH ARMAX Model Fit on VLCC Returns

	VLCC FORECAST	
DATE	RETURNS	FORECAST
8/1/2003	90.6132084	16194
9/1/2003	88.8203117	16261
10/1/2003	493.718666	16862
11/1/2003	695.692211	17956
12/1/2003	221.969186	18223
1/1/2004	-943.11116	17232
2/1/2004	-810.30522	16106
3/1/2004	30.3079286	15838

As shown in the autocorrelation function of the residuals, we can safely rule the possibility of an improper model. The residuals show exemplary behavior (no autocorrelation)

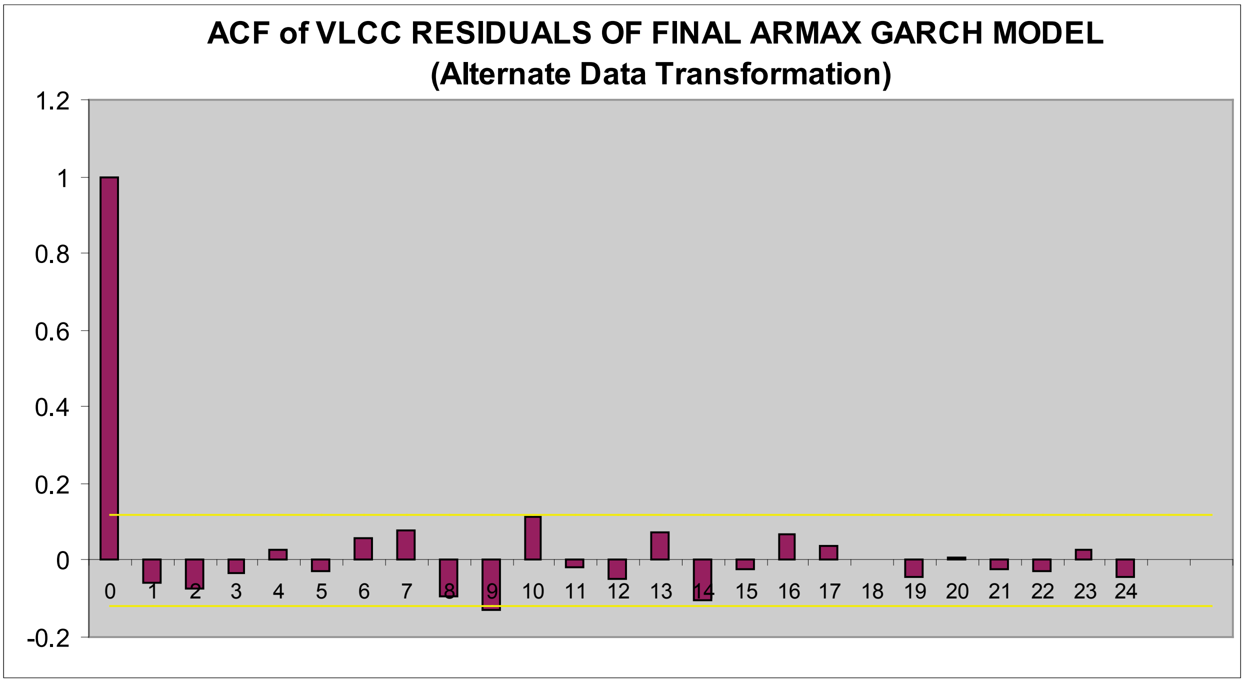

Figure 106 : ACE of VLCC Residuals of Final GARCH Model

Autocorrelation Check for White Noise

To Lag	Chi-Square	DF	Pr > ChiSq	--------------------Autocorrelations--------------------					
6	4.37	6	0.6273	-0.061	-0.073	-0.033	0.029	-0.027	0.058
12	18.53	12	0.1004	0.080	-0.092	-0.131	0.115	-0.017	-0.047
18	25.48	18	0.1123	0.073	-0.105	-0.021	0.069	0.039	-0.000
24	27.39	24	0.2866	-0.046	0.009	-0.025	-0.028	0.027	-0.043

8.68

7 Conclusions

7.1 Model Comparison

The MSE Error and the Akaike R-Squared Statistic for each of the models is:

Model Type	Root Mean Squared Error (10^{-5}) (RMSE) *Smaller is Better*	R-Squared Statistic (R^2) *Larger is Better*
VLCC		
Univariate GARCH	1023	.1791
Lagged Explanatory Variable	999	.2135
AFRAMAX		
Univariate GARCH	695	0.1470
Lagged Explanatory Variable	667	0.1853
HANDYSIZE		
Univariate GARCH	372	0.1741
Lagged Explanatory Variable	365	0.2020
SUEZMAX		
Univariate GARCH	924	0.2273
Lagged Explanatory Variable	915	0.2569

When comparing models it is useful to follow the simple guidelines. The root mean squared error (RMSE) – or the *standard error* – is the statistic which takes precedence over the other statistics.

During the parameter estimation process this statistic's value is minimized and its this statistic which determines the width of the confidence intervals for predictions.

The 95% confidence intervals for forecasts one step ahead are approximately equal to the point forecast "plus or minus 2 standard errors", that is, plus or minus two times the root-mean-squared error. This means that between two models with significantly different RMSEs, the model with the smallest RMSE must be chosen.

Its important to point out that the root mean squared error (and mean absolute error) can only be compared between models whose errors are measured in the same units. If one model's data are transformed or adjusted in any way, the same will be certain of the errors. For instance if one model's errors are in absolute units while the other's are in logged units (from logging the input data) then the error measure cannot be directly compared. In fact, simply unlogging or de-transforming the errors with the same transformation used in the inputdada will not make both models comparable!

There is also no absolute criterion for a 'good' RMSE value – it depends on the units in which the variable is measured and on the degree of accuracy (measured in the same unit) which is sought after in a particular model. It doesn't make sense to say that a model is good or bad because the RMSE is less than or greater than x unless there is a specific degree of accuracy that is sought after in the forecasting application.

In various regression models which use the same dependent variable and the same estimation period, the RMSE goes down as adjusted R-squared goes up. This means that the model with the highest adjusted R-squared will have the lowest RMSE. In simpler cases, one can use the adjusted R-squared as a guide but when comparing regression models in which the dependent variables were transformed in different ways (e.g., differenced in one case and undifferenced in another, or logged in one case and unlogged in another), or which used different sets of observations as the estimation period, R-squared is not a reliable guide to model quality. This can be shown in the GARCH model where the data has been transformed in order to meet model criteria. The resulting R-squared statistics are extremely low.

All things being comparable, when comparing RMSEs between two or more models, their differences should be significant. A 2% difference in the RMSE is not significant and in such cases, models should be chosen based on other factors such as their simplicity. This means that it's not worth adding another independent variable or a lag to a model in order to decrease the RMSE by only a few percent.

Finally, although the forecast confidence intervals are based almost entirely on the RMSE, the confidence intervals for longer-horizon forecasts depend on other model assumptions, particularly concerning the variability of the trend. For some models, the confidence intervals widen relatively slowly as the forecast horizon is lengthened while in other models confidence intervals widen much faster.

The rate at which the confidence intervals widen is *not* a reliable guide to model quality - what is important is the model should be making the correct assumptions about how uncertain the future is. It is very important that the model should pass the various residual diagnostic tests and "eyeball" tests in order for the confidence intervals for longer-horizon forecasts to be taken seriously.

7.2 Model Forecast Comparison

Each time charter rate has been forecasted using the following methods:

- Univariate ARIMA and Multivariate ARIMAX Models
- UNIVARIATE GARCH Models
- ARMAX GARCH Models (Crude Oil Purchase Price Explanatory Variable)

Plotting the results, we get:

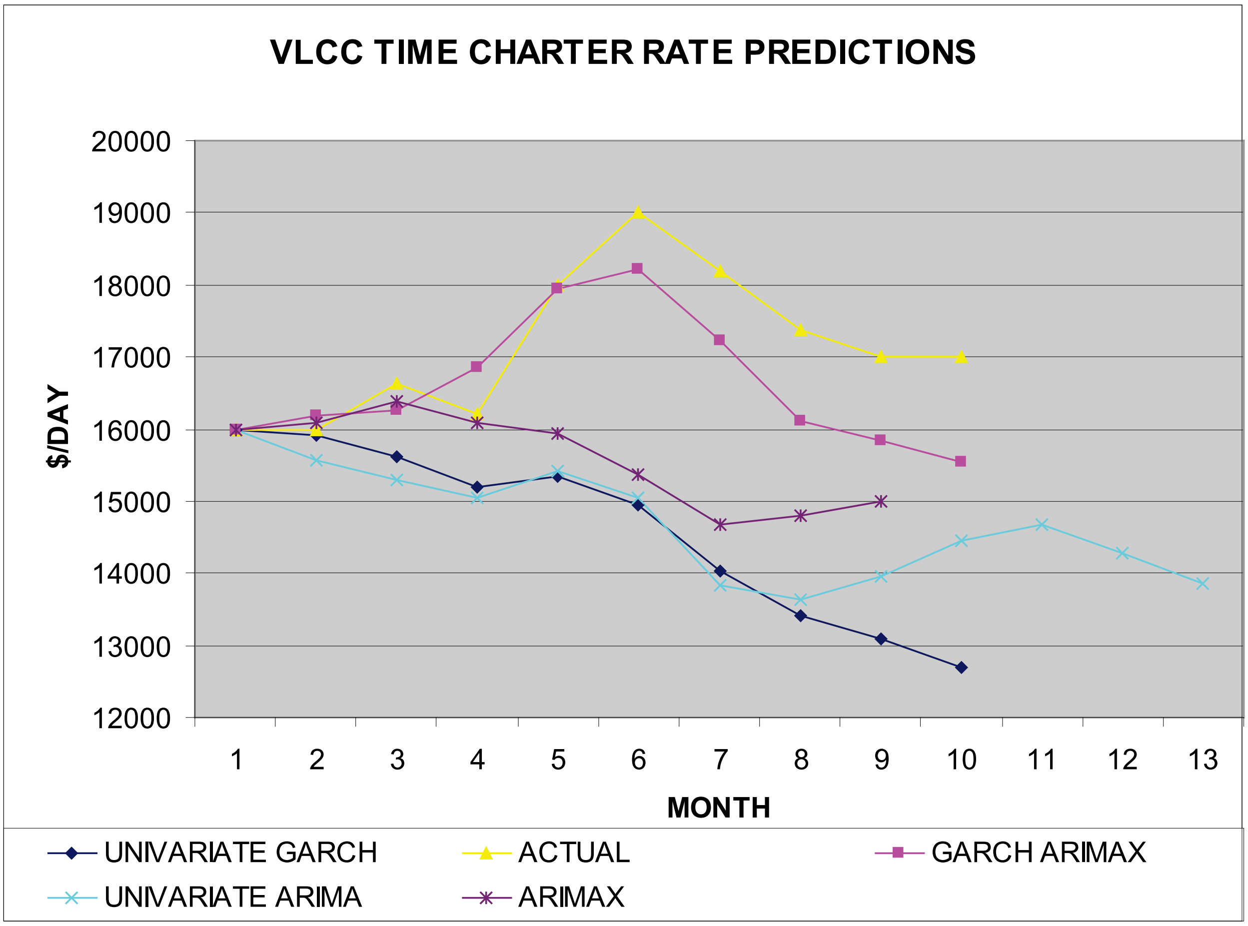

Figure 107 : VLCC Time Charter Rate Predictions

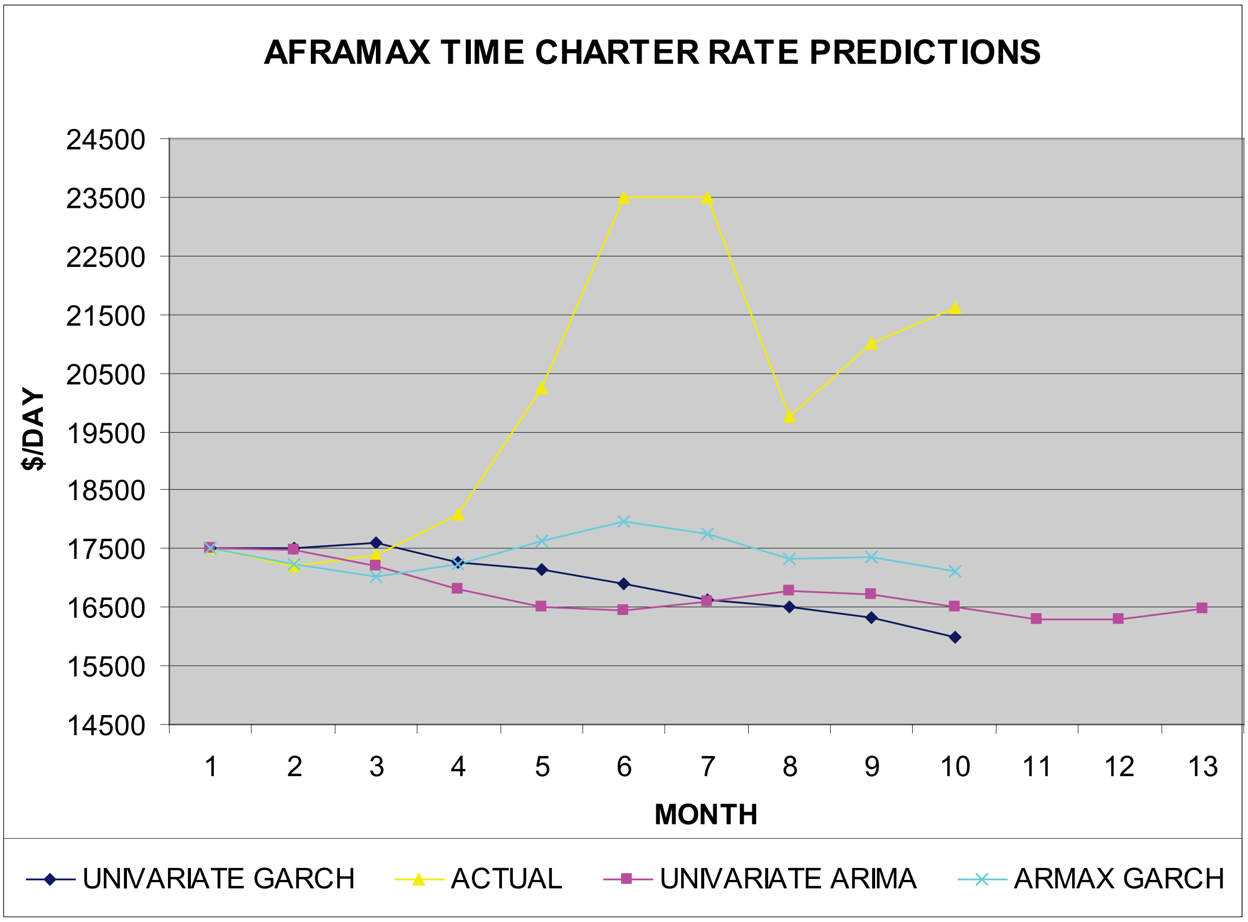

Figure 108 : AFRAMAX Time Charter Rate Predictions

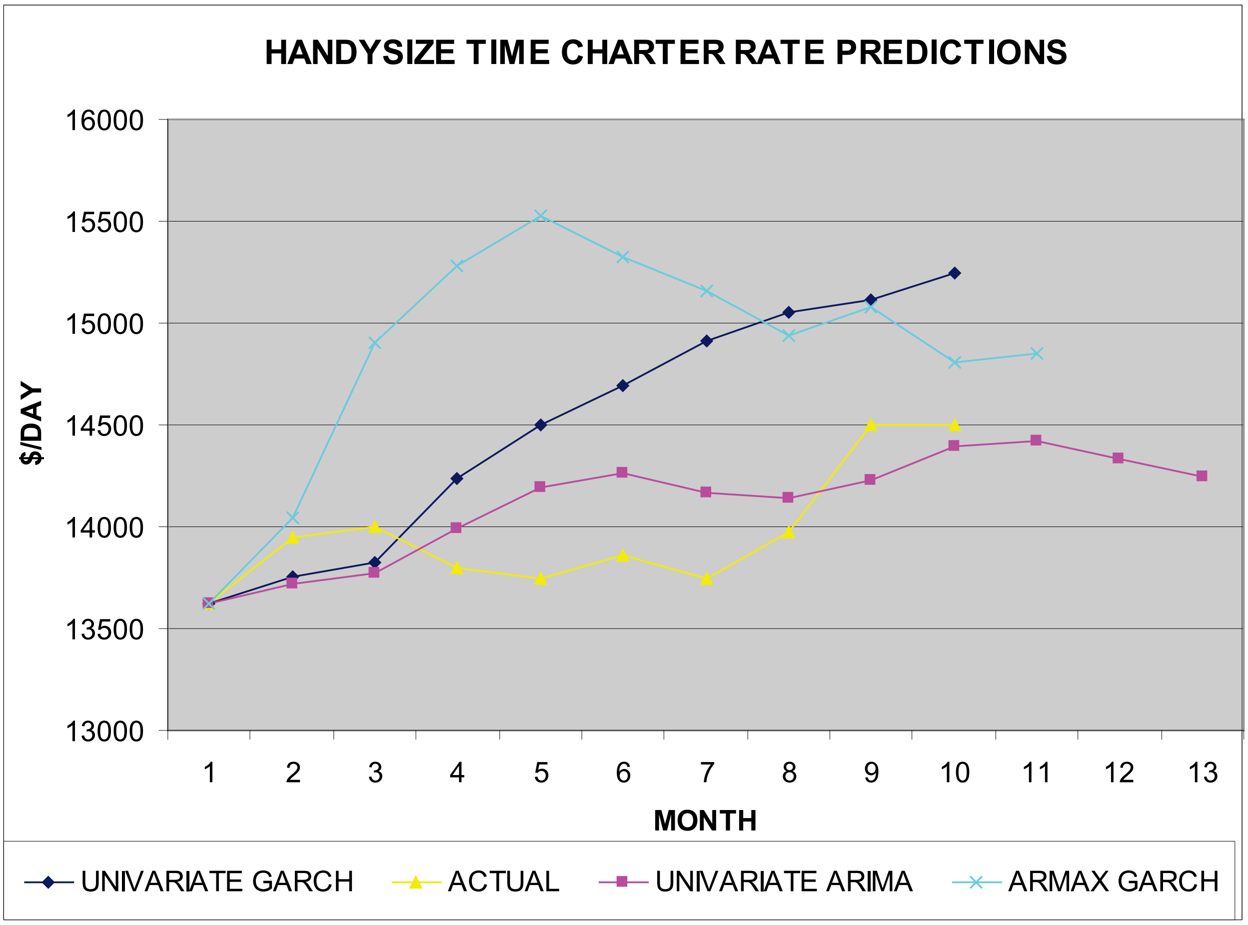

Figure 109 : HANDSYZE Time Charter Rate Predictions

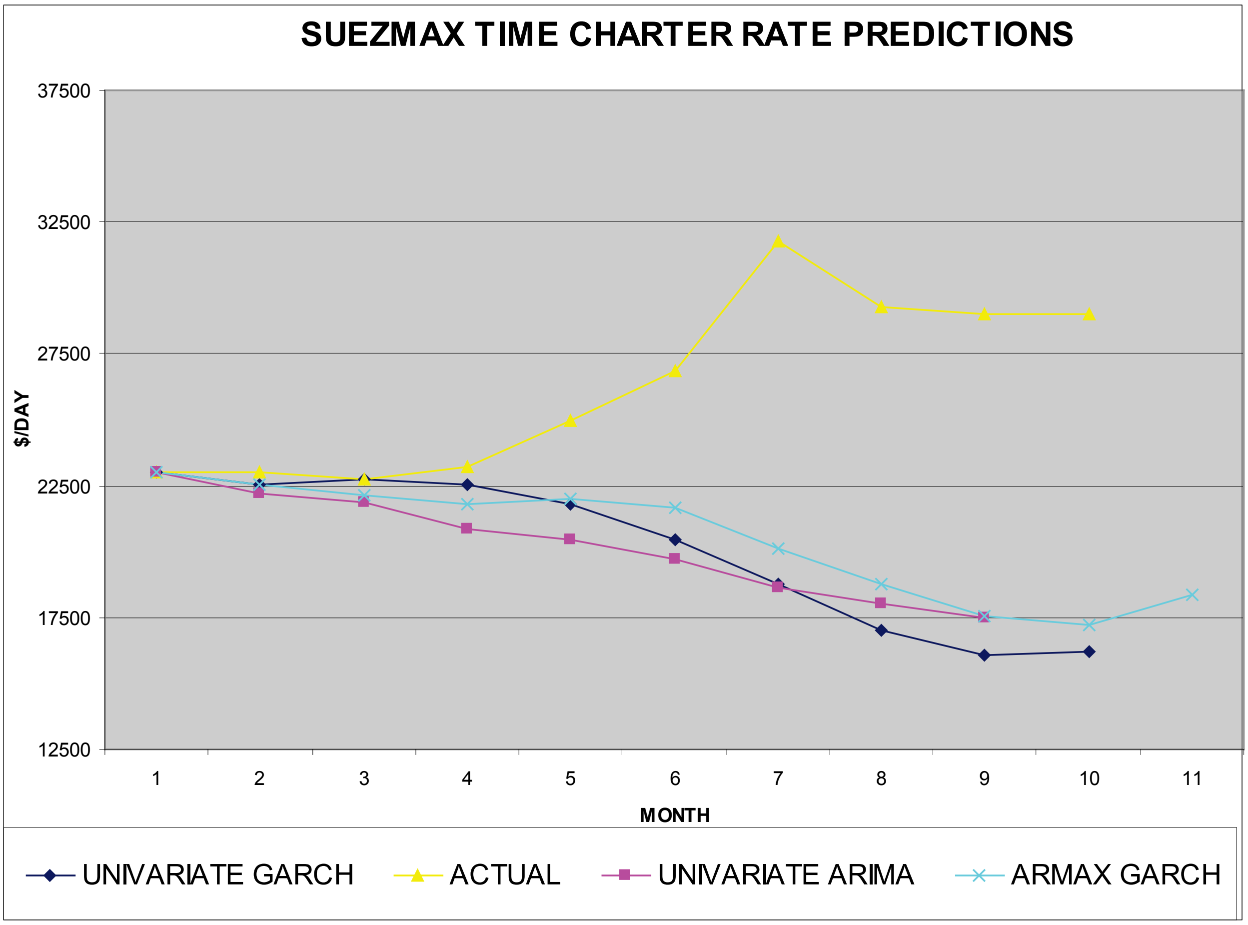

Figure 110 : SUEZMAX Time Charter Rate Predictions

With the completion of this thesis, the following characteristics of time charter rates have been identified.

VOLATILITY

From the data we can each time charter rate may be successfully modeled up to around a four month period during periods of low volatility. On the other hand the sudden changes that are due to volatility are above the ability of the forecasting methods. This is likely because while there are certain seasonal characteristics associated with time charter rates, these characteristics do not in themselves contribute to the volatility of the time charter rate. The volatility of each time charter rate can be considered unique (for each size class) and if there are common factors that manifest themselves in the time charter rate market they have different effects depending on the size class.

DEPENDENT VARIABLES

The statistical significance of lagged dependent variables is also, in the long run, not a factor in model performance when volatility comes into effect. The use of statistically significant lagged dependent variables does not ensure an accurate long range forecast.

UNIVARIATE VS. MULTIVARIATE FORECASTS

While the RMSE is usually minimized when using multivariate (ARMAX) models, the difference in forecasting ability between ARIMAX models (both ARIMA and GARCH) has not been proven in practice as being significantly better, either way.

FORECASTING ABILITY

The forecasting ability of all models seems to decrease significant after the third or fourth forecasted month. The volatile nature of time charter rates, combined with an elusive explanatory variable data set renders their forecasting a difficult undertaking. As far as the VLCC Time Charter Rate results are concerned, the GARCH model with the alternate data transformation and the Crude Oil Purchase Price lagged dependent variable fared the best. While not being a carbon copy of the actual values, it managed to give a better account of the time charter rate's future shift than any other model.

It must be pointed out that the VLCC time charter rate shown above was modeled relatively successfully by the GARCH ARMAX model. Yet the addition of new data will always change the model specification and this month's successful model is next season's disappointment.

CONCLUSION

It has been proven that neither the standard explanatory variables associated with time charter rates, or lagged time charter rates themselves contain useful information to be able to forecast ARIMA, ARIMAX and GARCH models successfully over long time periods. As in the OLS analysis, there is a significant statistical correlation between dependent and independent variables but the low R-squared value of each model indicates that the explanatory variables are not capable of describing the volatility of the VLCC time charter rate. With the help of chapter one as a guide, it may be possible to locate other explanatory variables which may be more successful for forecasting time charter rates.

8 SOURCE CODE

Load Dataset

```
data dimitris;
    input aframax handysize suezmax vlcc;
      date = intnx('month','1oct1979'd, _n_-1);
      format date monyy7.;  /* 7 for MONYYYY */
      vlcc_r = dif(log(vlcc));
      aframax_r = dif(log(aframax));
      handysize_r = dif(log(handysize));
      suezmax_r = dif(log(suezmax));
      datalines;
[values]
.     .     .     .
.     .     .     .
.     .     .     .
;
```

8.1 Ordinary Least Squares Analysis of VLCC and Crude Oil Purchase Price

```
proc autoreg data=dimitris;
          model VLCC = COPP;
run;
```

8.2 Ordinary Least Squares Analysis of SUEZMAX and Crude Oil Purchase Price

```
proc autoreg data=dimitris;
    model SUEZMAX = COPP;
run;
```

8.3 Ordinary Least Squares Analysis of AFRAMAX and Crude Oil Purchase Price

```
proc autoreg data=dimitris;
    model AFRAMAX = COPP;
run;
```

8.4 Ordinary Least Squares Analysis of HANDYSIZE and Crude Oil Purchase Price

```
proc autoreg data=dimitris;
    model HANDYSIZE = COPP;
run;
```

8.5 Ordinary Least Squares Analysis of VLCC and AFRAMAX

```
proc autoreg data=dimitris;
      model VLCC = AFRAMAX;
run;
```

8.6 Ordinary Least Squares Analysis of VLCC and HANDYSIZE

```
proc autoreg data=dimitris;
      model VLCC = HANDYSIZE;
run;
```

8.7 Ordinary Least Squares Analysis of VLCC and SUEZMAX

```
proc autoreg data=dimitris;
      model VLCC = SUEZMAX;
run;
```

8.8 Cross Correlation Function of VLCC Returns and Crude Oil Purchase Price Returns

```
proc arima data=dimitris;
identify var=vlcc_r crosscorr=copp_r nlag=13;
run;
```

8.9 Cross Correlation Function of AFRAMAX Returns and Crude Oil Purchase Price Returns

```
proc arima data=dimitris;
identify var=aframax_r crosscorr=copp_r nlag=13;
run;
```

8.10 Cross Correlation Function of HANDYSZIE Returns and Crude Oil Purchase Price Returns

```
proc arima data=dimitris;
identify var=handysize_r crosscorr=copp_r nlag=13;
run;
```

8.11 Cross Correlation Function of SUEZMAX Returns and Crude Oil Purchase Price Returns

```
proc arima data=dimitris;
identify var=suezmax_r crosscorr=copp_r nlag=13;
run;
```

8.12 Cross Correlation Function of VLCC Returns and VLCC Scrap Prices

```
proc arima data=dimitris;
identify var=vlcc_r crosscorr=vlccscrap_r nlag=13;
run;
```

8.13 Cross Correlation Function of VLCC Returns and VLCC Five Year Second Hand Price Returns

```
proc arima data=dimitris;
identify var=vlcc_r crosscorr=vlcc5yrsechnd_r nlag=13;
run;
```

8.14 Cross Correlation Function of VLCC Returns and Crude Energy Materials Returns

```
proc arima data=dimitris;
identify var=vlcc_r crosscorr=cem_r nlag=13;
run;
```

8.15 Cross Correlation Function of VLCC Returns and New Building Price Returns

```
proc arima data=dimitris;
identify var=vlcc_r crosscorr=nb_r nlag=13;
run;
```

8.16 Cross Correlation Function of VLCC Returns and Arab Group Oil Production Returns

```
proc arima data=dimitris;
identify var=nb_r crosscorr=agop_r nlag=13;
run;
```

8.17 Cross Correlation Function of VLCC Returns and AFRAMAX Returns

```
proc arima data=dimitris;
identify var=vlcc_r crosscorr=aframax_r nlag=13;
run;
```

8.18 Cross Correlation Function of VLCC Returns and SUEZMAX Returns

```
proc arima data=dimitris;
identify var=vlcc_r crosscorr=suezmax_r nlag=13;
run;
```

8.19 Cross Correlation Function of VLCC Returns and HANDYSIZE Returns

```
proc arima data=dimitris;
identify var=vlcc_r crosscorr=handysize_r nlag=13;
run;
```

8.20 Plot VLCC Time Charter Rates

```
proc gplot data=dimitris;
      title "VLCC Time Charter Rates";
      symbol1 v=none c=blue  i=join;
      plot vlcc * date = 1/ legend overlay;
run;
```

8.21 VLCC Time Charter Rate Stationarity Test

```
proc arima data=dimitris;
      identify var=vlcc stationarity=(dickey) nlag=24;
run;
```

8.22 VLCC Time Charter Rate Returns Stationarity Test

```
proc arima data=dimitris;
      identify var=vlcc_r stationarity=(dickey) nlag=12;
run;
```

8.23 Plot AFRAMAX Time Charter Rates

```
proc gplot data=dimitris;
      title "Aframax Time Charter Rates";
      symbol1 v=none c=blue  i=join;
      plot aframax * date = 1/ legend overlay;
run;
```

8.24 AFRAMAX Time Charter Rate Stationarity Test

```
proc arima data=dimitris;
      identify var=aframax stationarity=(dickey) nlag=12;
run;
```

8.25 AFRAMAX Time Charter Rate Returns Stationarity Test

```
proc arima data=dimitris;
identify var=aframax_r nlag=12;
run;
```

8.26 Plot HANDYSIZE Time Charter Rates

```
proc gplot data=dimitris;
      title "Handysize Time Charter Rates";
      symbol1 v=none c=blue  i=join;
      plot handysize * date = 1/ legend overlay;
run;
```

8.27 HANDYSIZE Time Charter Rate Stationairity Test

```
proc arima data=dimitris;
identify var=handysize nlag=13;
run;
```

8.28 HANDYSIZE Time Charter Rate Returns Stationarity Test

```
proc arima data=handy;
      identify var=handysize_r nlag=13;
run;
```

8.29 Plot SUEZMAX Time Charter Rates

```
proc gplot data=dimitris;
      title "Suezmax Time Charter Rates";
      symbol1 v=none c=blue  i=join;
      plot suezmax * date = 1/ legend overlay;
run;
```

8.30 SUEZMAX Time Charter Rate Stationarity Test

```
proc arima data=dimitris;
      identify var=suezmax stationarity=(dickey) nlag=12;
run;
```

8.31 SUEZMAX Time Charter Rate Returns Stationarity Test

```
proc arima data=dimitris;
      identify var=suezmax_r stationarity=(dickey) nlag=12;
run;
```

8.32 OLS Model of VLCC and DATE

```
proc autoreg data=dimitris;
      model vlcc = date;
run;
```

8.33 VLCC Time Charter Rate Returns Durbin-Watson and ARCH Tests

```
proc autoreg data=dimitris;
      model vlcc_r = date /dw=13 archtest dwprob;
run;
```

8.34 AFRAMAX Time Charter Rate Returns Durbin-Watson and ARCH Tests

```
proc autoreg data=dimitris;
      model aframax_r = date /dw=13 archtest dwprob;
run;
```

8.35 SUEZMAX Time Charter Rate Returns Durbin-Watson and ARCH Tests

```
proc autoreg data=dimitris;
      model suezmax_r = date /dw=13 archtest dwprob;
run;
```

8.36 HANDYSIZE Time Charter Rate Returns Durbin-Watson and ARCH Tests

```
proc autoreg data=dimitris;
      model handysize_r = date /dw=13 archtest dwprob;
run;
```

8.37 VLCC Time Charter Rate Returns Backward Elimination of Autoregressive Terms

```
proc autoreg data=dimitris;
      model vlcc_r = date / method=ml nlag=13 backstep;
run;
```

8.38 AFRAMAX Time Charter Rate Returns Backward Elimination of Autoregressive Terms

```
proc autoreg data=dimitris;
      model aframax_r = date / method=ml nlag=13 backstep;
run;
```

8.39 HANDYSIZE Time Charter Rate Returns Backward Elimination of Autoregressive Terms

```
proc autoreg data=dimitris;
      model handysize_r = date / method=ml nlag=13 backstep;
run;
```

8.40 SUEZMAX Time Charter Rate Returns Backward Elimination of Autoregressive Terms

```
proc autoreg data=dimitris;
      model suezmax_r = date / method=ml nlag=13 backstep;
run;
```

8.41 VLCC Time Charter Rate Returns GARCH Model Estimates

```
proc autoreg data=dimitris;
      model vlcc_r = date / nlag=(1 12) garch=(q=1,p=1) archtest dwprob;
      output out=out cev=vhat;
run;
```

8.42 AFRAMAX Time Charter Rate Returns GARCH Model Estimates

```
proc autoreg data=dimitris;
  model aframax_r = date / nlag=(1 5 9) garch=(q=2,p=1) maxit=90000 archtest
dwprob;
  output out=out cev=vhat;
run;
```

8.43 HANDYSIZE Time Charter Rate Returns GARCH Model Estimates

```
proc autoreg data=dimitris;
    model handysize_r = date / nlag=(1 3 4) garch=(q=2, p=1) maxit=200;
    output out=out3 r=handysize_r_resid cev=vhat p=handysize_r_pred lcl=lcl
ucl=ucl;
run;
```

8.44 SUEZMAX Time Charter Rate Returns GARCH Model Estimates

```
proc autoreg data=dimitris;
      model suezmax_r = date / nlag=(1 3 13) garch=(q=1, p=1) archtest dwprob;
      output out=out4 cev=vhat;
run;
```

8.45 Forecasting VLCC Time Charter Rate Returns Using GARCH

```
proc autoreg data=dimitris;
      model vlcc_r = date / nlag=(1 12) garch=(q=2, p=1);
      output out=out1 r=vlcc_r_resid cev=vhat p=vlcc_r_pred lcl=lcl ucl=ucl;
run;
```

8.46 Plot Evaluation & Fit of a GARCH model on VLCC Returns

```
proc gplot data=out1;
      title "Evaluation & Fit of a GARCH model on VLCC returns";
      symbol1 v=none c=blue  i=join;
      symbol2 v=none c=red i=join;
      symbol3 v=none c=brown i=join;
      symbol4 v=none c=green i=join;
      plot vlcc_r * date = 1
      vlcc_r_pred * date = 2
       / legend overlay;
run;
```

8.47 Plot Forecast of a GARCH Model on VLCC Returns

```
proc gplot data=out1;
      title "Forecast of GARCH model on VLCC Returns";
      symbol1 v=none c=blue  i=join;
      symbol2 v=none c=red i=join;
      symbol3 v=none c=brown i=join;
      symbol4 v=none c=green i=join;
      plot vlcc_r * date = 1
      vlcc_r_pred * date = 2
       / href='1may2004'd haxis= '1jan2001'd to '1jan2005'd by yr legend
overlay;
run;
```

8.48 Plot Evaluation of the Residuals of a GARCH Model on VLCC Returns

```
proc gplot data=out1;
      title "Evaluation of the Residuals of GARCH model on VLCC returns";
      symbol1 v=none c=blue  i=join;
      symbol2 v=none c=red i=join;
      symbol3 v=none c=brown i=join;
      symbol4 v=none c=green i=join;
      plot vlcc_r_resid * date = 1
       / legend overlay;
run;
```

8.49 Statistical Analysis of the Residuals of a GARCH Model on VLCC Time Charte Rate Returns

```
proc arima data=out1;
      identify var=vlcc_r_resid;
run;
```

8.50 Forecasting AFRAMAX Time Charter Rate Returns Using GARCH

```
proc autoreg data=dimitris;
      model aframax_r = date / nlag=(1 5 9) garch=(q=2, p=1) maxit=200;
      output out=out2 r=aframax_r_resid cev=vhat p=aframax_r_pred lcl=lcl
ucl=ucl;
run;
```

8.51 Plot Evaluation & Fit of a GARCH model on AFRAMAX Returns

```
proc gplot data=out2;
      title  "Evaluation & Fit of a GARCH model on AFRAMAX returns";
      symbol1 v=none c=blue  i=join;
      symbol2 v=none c=red i=join;
      symbol3 v=none c=brown i=join;
      symbol4 v=none c=green i=join;
      plot aframax_r * date = 1
      aframax_r_pred * date = 2
       / legend overlay;
run;
```

8.52 Plot Forecast of GARCH Model on AFRAMAX Returns

```
proc gplot data=out2;
      title "Forecast of GARCH model on AFRAMAX Returns";
      symbol1 v=none c=blue  i=join;
      symbol2 v=none c=red i=join;
      symbol3 v=none c=brown i=join;
      symbol4 v=none c=green i=join;
      plot aframax_r * date = 1
      aframax_r_pred * date = 2
       / href='1may2004'd haxis= '1jan2001'd to '1jan2005'd by yr legend
overlay;
run;
```

8.53 Plot Evaluation of the Residuals of GARCH Model on AFRAMAX Returns

```
proc gplot data=out2;
      title "Evaluation of the Residuals of GARCH model on AFRAMAX returns";
      symbol1 v=none c=blue  i=join;
      symbol2 v=none c=red i=join;
      symbol3 v=none c=brown i=join;
      symbol4 v=none c=green i=join;
      plot aframax_r_resid * date = 1
       / legend overlay;
run;
```

8.54 Statistical Analysis of the Residuals of a GARCH Model on AFRAMAX Time Charte Rate Returns

```
proc arima data=out2;
      identify var=aframax_r_resid;
run;
```

8.55 Forecasting HANDYSIZE Time Charter Rate Returns Using GARCH

```
proc autoreg data=dimitris;
   model handysize_r = date / nlag=(1 3 4) garch=(q=2, p=1);
   output out=out3 r=handysize_r_resid cev=vhat p=handysize_r_pred lcl=lcl
ucl=ucl;
run;
```

8.56 Plot Evaluation & Fit of a GARCH Model on HANDYSIZE Returns

```
proc gplot data=out3;
      title "Evaluation & Fit of a GARCH model on HANDYSIZE returns";
      symbol1 v=none c=blue  i=join;
      symbol2 v=none c=red i=join;
      symbol3 v=none c=brown i=join;
      symbol4 v=none c=green i=join;
      plot handysize_r * date = 1
      handysize_r_pred * date = 2 / legend overlay;
run;
```

8.57 Plot Forecast of GARCH Model on HANDYSIZE Returns

```
proc gplot data=out3;
      title "Forecast of GARCH model on HANDYSIZE returns";
      symbol1 v=none c=blue  i=join;
      symbol2 v=none c=red i=join;
      symbol3 v=none c=brown i=join;
      symbol4 v=none c=green i=join;
      plot handysize_r * date = 1
      handysize_r_pred * date = 2 / href='1may2004'd haxis= '1jan2001'd to
'1jan2005'd by yr legend overlay;
run;
```

8.58 Plot Evaluation of the Residuals of GARCH model on HANDYSIZE Returns

```
proc gplot data=out3;
      title "Evaluation of the Residuals of GARCH model on HANDYSIZE returns";
      symbol1 v=none c=blue  i=join;
      symbol2 v=none c=red i=join;
      symbol3 v=none c=brown i=join;
      symbol4 v=none c=green i=join;
      plot handysize_r_resid * date = 1 / legend overlay;
run;
```

8.59 Statistical Analysis of the Residuals of a GARCH Model on HANDYSIZE Time Charter Rate Returns

```
proc arima data=out3;
identify var=handysize_r_resid;
run;
```

8.60 Forecasting HANDYSIZE Time Charter Rate Returns Using GARCH

```
proc autoreg data=dimitris;
      model suezmax_r = date / nlag=(1 3 13) garch=(q=1, p=1);
      output out=out4 r=suezmax_r_resid cev=vhat p=suezmax_r_pred lcl=lcl
ucl=ucl;
run;
```

8.61 Plot Evaluation & Fit of a GARCH Model on SUEZMAX Time Charter Rate Returns

```
proc gplot data=out4;
      title "Evaluation & Fit of a GARCH model on SUEZMAX returns";
      symbol1 v=none c=blue  i=join;
      symbol2 v=none c=red i=join;
      symbol3 v=none c=brown i=join;
      symbol4 v=none c=green i=join;
      plot suezmax_r * date = 1
      suezmax_r_pred * date = 2
       / legend overlay;
run;
```

8.62 Plot Forecast of GARCH Model on HANDYSIZE Returns

```
proc gplot data=out4;
      title "Forecast of GARCH model on HANDYSIZE returns";
      symbol1 v=none c=blue  i=join;
      symbol2 v=none c=red i=join;
      symbol3 v=none c=brown i=join;
      symbol4 v=none c=green i=join;
      plot suezmax_r * date = 1
      suezmax_r_pred * date = 2
       / href='1may2004'd haxis= '1jan2001'd to '1jan2005'd by yr legend
overlay;
run;
```

8.63 Plot Evaluation of the Residuals of GARCH Model on SUEZMAX Returns

```
proc gplot data=out4;
      title "Evaluation of the Residuals of GARCH model on SUEZMAX returns";
      symbol1 v=none c=blue  i=join;
      symbol2 v=none c=red i=join;
      symbol3 v=none c=brown i=join;
      symbol4 v=none c=green i=join;
      plot suezmax_r_resid * date = 1
       / legend overlay;
run;
```

8.64 Statistical Analysis of the Residuals of a GARCH Model on SUEZMAX Time Charter Rate Returns

```
proc arima data=out4;
      identify var=suezmax_r_resid;
run;
```

8.65 Final Model Durbin-Watson and ARCH Tests

```
proc autoreg data=dimitris;
      model vlcc  = copplag10 date / nlag=36 backstep;
run;
```

8.66 Final Model Fit

```
proc autoreg data=dimitris;
      model vlcc = copplag10 date / nlag=(1 10 12 18) garch=(q=1, p=1);
      output out=out5 r=vlcc_r_resid cev=vhat p=vlcc_r_pred lcl=lcl ucl=ucl;
run;
```

8.67 Plot Forecast of GARCH model of VLCC Returns

```
proc gplot data=out5;
      title "Forecast of GARCH model on VLCC returns";
      symbol1 v=none c=blue  i=join;
      symbol2 v=none c=red i=join;
      symbol3 v=none c=brown i=join;
      symbol4 v=none c=green i=join;
      plot vlcc * date = 1
      vlcc_r_pred * date = 2
      / href='1jun2003'd haxis= '1jun2001'd to '1oct2005'd by yr legend
overlay;
run;
```

8.68 Statistical Analysis of the Residuals of a GARCH Model on the Final Model

```
proc arima data=out5;
      identify var=vlcc_r_resid;
run;
```

References

SAS Institute Inc. (1993), *SAS/ETS User's Guide, Version 8, Cary*, NC: SAS Institute Inc.

Vandaele, Walter (1983). “Applied Time Series and Box-Jenkins Models”
Academic Press, INC.

Sharma, Namit (1998) “Forecasting Oil Price Volatility” , Master Thesis.

Tsolakis, Stavros (2005) “Econometric Analysis of Bulk Shipping Markets”

Investment Research (2001) “The Shipping Landscape”

Brocklebank, J. C. and Dickey, D. A. (2003) SAS System for Forecasting Time Series, SAS Institute.

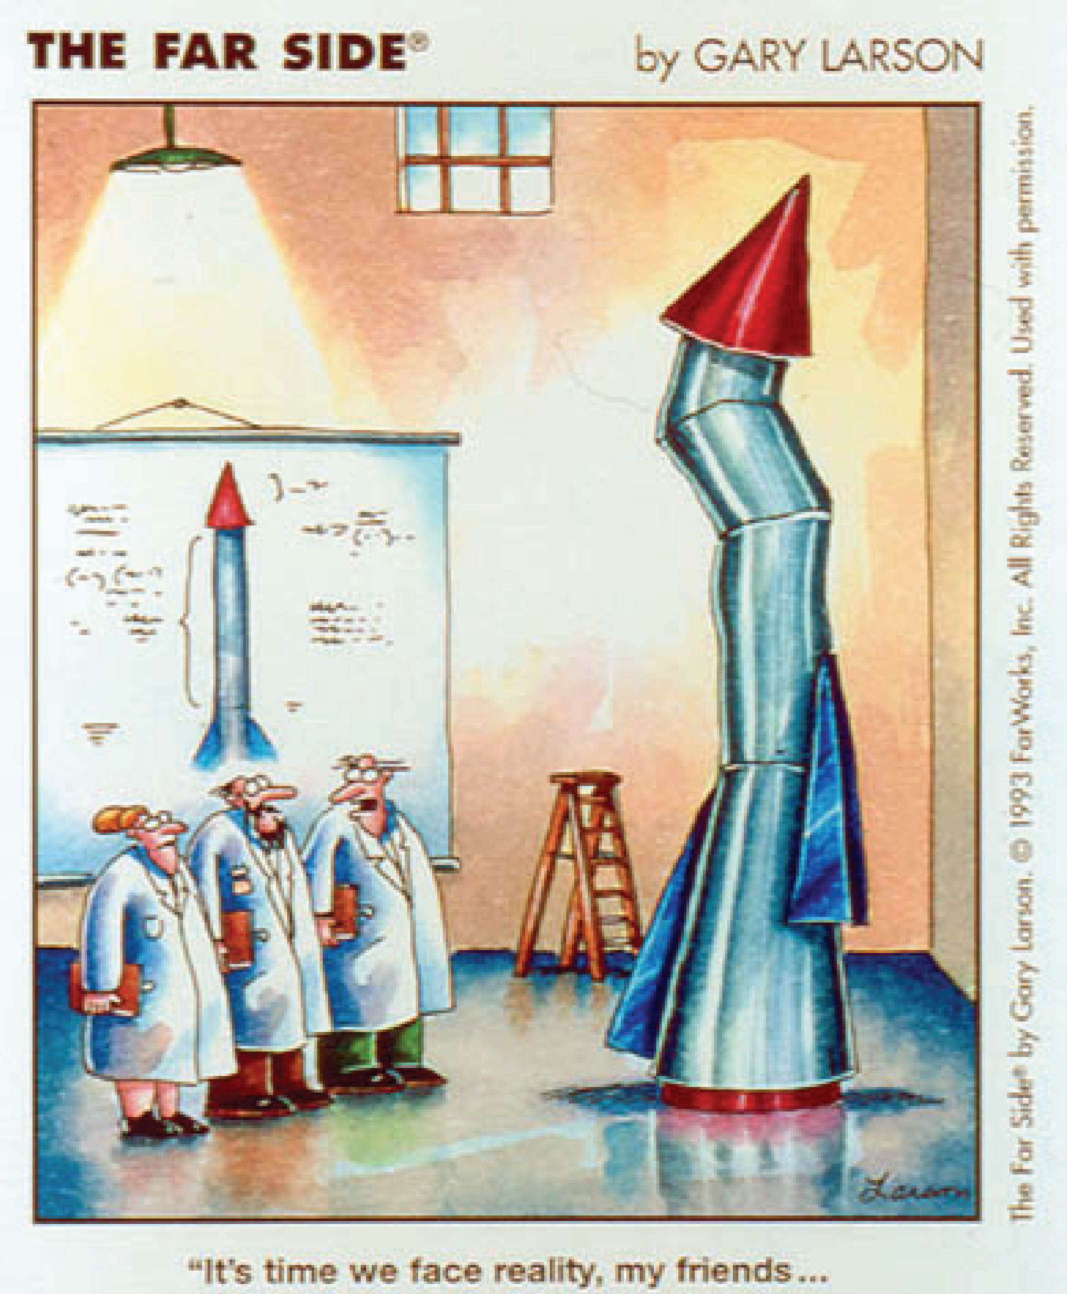

"It's time we face reality, my friends...
We're not exactly rocket scientists."