Modelling and Forecasting Volatility of JSE Sectoral Indices: A Model Confidence Set Exercise

University of the Witwatersrand, Johannesburg

Matthew Song

0604878F

June 2014

Supervisor: Prof. Paul Alagidede

Declaration of Authorship

I, Matthew Song, declare that this research thesis and the work presented herein are my own. This thesis or any part thereof has not been submitted to any other university for a degree or any other qualification. The submission of this document is in partial fulfillment of the requirements for a Masters of Management in Finance and Investment at the University of the Witwatersrand. The consulted work of other authors and resources is clearly acknowledged where appropriate.

	12/06/2014
Matthew Song	Date

Abstract

Volatility plays an important role in option pricing and risk management. It is crucial that volatility is modelled as accurately as possible in order to forecast with confidence. The challenge is in the selection of the 'best' model with so many available models and selection criteria. The Model Confidence Set (MCS) solves this problem by choosing a group of models that are equally good. A set of GARCH models were estimated for several JSE indices and the MCS was used to trim the group of models to a subset of equally superior models. Using the Mean Squared Error to evaluate the relative performance of the MCS, GARCH (1,1) and Random Walk, it was found that the MCS, with an equally weighted combination of models, performed better than the GARCH (1,1) and Random Walk for instances where volatility in the returns data was high. For instances of low volatility in the returns, the GARCH (1,1) had superior 5-day forecasts but the MCS had better performance for 10-days and greater. The EGARCH (2,1) volatility model was selected by the MCS for 5 out of the 6 indices as the most superior model. The Random Walk was shown to have better long term forecasting performance.

Keywords: Model Confidence Set, volatility models, volatility forecasting

Acknowledgments

I would like to thank Professor Paul Alagidede for his invaluable knowledge and patience in reviewing, commenting and offering guidance throughout this research project.

To my family for all the support and wisdom they have gifted me. The sacrifice and hard work they have put in to give me the opportunity and freedom to pursue my interests is lovingly appreciated to no end. The sound advice they have given me throughout my life will be treasured forever.

Contents

1 Chapter 1

1.1 Introduction

Investing can be defined as committing resources in order to achieve later benefits (Luenberger 2009). It is a choice that business entities and consumers make in order to achieve a particular goal in a certain time frame. Volatility plays an important role in the decision making process and is a crucial unknown variable in option pricing and valuations (Ederington & Guan 2006). When estimating volatility, Ederington & Guan (2006) stressed the importance of the measurement of volatility and that it could have 'dramatic' effects on out-of-money options.

The JSE Resource Sector has been highly volatile during the third quarter of 2012 to the present day. Due to the rampant mining strikes throughout the South African mining sector starting in August 2012, output and revenue of the mining companies have become highly uncertain. As an example, on 14 June 2013, Anglo Platinum shares fell by 0.14% on news of a strike but Lonmin shares appreciated by an impressive 6.48% on news of a postponed strike (Nortje 2013). However, during the preceding week of 6 June 2013, Lonmin shares surprisingly appreciated by 4.01% on news that a worker's union were threatening similar strike action (DeIonno 2013). The volatility in the share price seems unpredictable and coupled with the important role that volatility has in option pricing, it is important to understand it and how it is determined.

1.2 What is Volatility?

There are many definitions of volatility but perhaps the best way of describing volatility is to describe its characteristics. Masset (2011) performed a study on both bullish and bearish markets and was able to confirm the following 'stylized facts' about volatility:

1. It is not constant but it is mean reverting.
2. It tends to cluster in that high volatility periods are likely to be followed by more periods of high volatility.
3. The distribution of returns have fat tails. Shocks to the economy have a strong impact on volatility. This feature of volatility is particularly difficult to capture with most financial models.
4. It displays the 'Leverage Effect'. This 'stylized' effect has been widely researched. Aboura & Wagner (2013) and Zhou (2007) arrived at the same conclusion that volatility is more pronounced when returns are negative and that for positive returns, volatility is less dramatic.
5. The explanatory power of volatility is greater when measured at a high frequency when compared to lower sampling periods.

1.3 Why is Volatility Important?

In option pricing theory, the Black-Scholes pricing model has the volatility of the asset directly in the formula (Jordan et al. 2012). Therefore an accurate estimation of the variance will be highly valued by financial professionals working in the options space. Volatility also plays a crucial role in another option pricing technique called the Binomial

Tree Model which can be used to price both American and European options. In this model, the price change at every branch in the tree is determined by the volatility.

For risk management, volatility is used to quantify Value at Risk (VaR) and Expected Shortfall values. The banking industry uses these risk measures to determine the amount of reserve that is kept and has an impact on the profit and loans that can safely be issued.

1.4 Problem Statement

The 'stylized facts' of volatility present certain gaps that need to be studied in order to keep knowledge relevant and up to date. The fact that volatility is not constant and demonstrates clustering behaviour means that it must be studied continuously. The Lonmin example shows that volatility is relevant in the South African market and with the clustering effect, it will most probably persist for some time. With the importance that volatility plays in option pricing the data must be kept current and accurate. Very few papers have used the Model Confidence Set (MCS) to choose the best performing models and in particular no papers have used the MCS for the South African market.

1.5 Objectives

The Problem Statement Section highlighted several objectives for this study. The objectives can be broadly defined as follows:

1. Model several Johannesburg Stock Exchange (JSE) sectoral indices using a variety of univariate volatility models.

2. The in-sample data used to construct the volatility models must be chosen with the most recent data to characterize a suitable model that is up to date.

3. Use the Model Confidence Set framework to select the superior model or set of models for each index.

4. Forecast volatility for the selected models and compare its performance with the GARCH (1,1) and Random Walk model.

2 Chapter 2

2.1 Literature Review

There are two broad categories for volatility modelling. There are historic volatility models and implied volatility models. Implied volatility calculates or derives future volatility using the price of options. Historical volatility is calculated using past price information and suitably estimated models to forecast future volatility. Deciding on the best form of model to use seems to be a point of contention. In a review of 39 papers discussing the merits of each of the two categories, Samouilhan & Shannon (2008) found that half were in favour of implied volatility modelling and the other half viewed the historic models as being superior.

An interesting paper by Venter (2003), showed that implied volatility modelling for JSE warrants compared marginally better than historic models. In his paper, Venter comes to the conclusion that no model can be used as a ‘cookie cutter’. In other words, no matter what category of volatility modelling is used, it must be fit for purpose. This argument by Venter, is challenged by Hansen & Lunde (2005) where they asked the question whether or not anything beats a GARCH (1,1) model. They found no evidence that a GARCH (1,1) model is outperformed by other complex and sophisticated models when analyzing exchange rates. However they did concede that the GARCH (1,1) is inferior when analyzing IBM stock returns.

During the literature review it was found that the evaluation methods in determining the superior volatility model differed between papers. Most papers use the Mean Squared Error (MSE) technique to evaluate models. Ramashala (2011) on the other hand, used hypothesis testing to show that implied volatility models were inferior to historic models. There are other significance testing methods used to evaluate the performance of models. The test for Superior Predictive Ability (SPA) and the Model Confidence Set by Hansen et al. (2003) are examples of such models.

On the South African National Electronic Thesis database, there was no evidence that the Model Confidence Set has been used to evaluate models based on the South African market.

2.2 Volatility Models

2.2.1 ARCH Models

Linear models such as the ARMA class of models assume that several properties in the data are present. One of the properties that is assumed to be constant is the variance of the error terms. When volatility is present in the data, this assumption does not hold. The variance of the error terms change with time and this is called heteroscedasticity. Heteroscedasticity has many undesirable consequences if it is not recognised. Some of the undesirable consequences include inaccurate estimates of coefficients under the Ordinary Least Squares (OLS) method and as a result inaccurate R^2 values will be obtained.

Therefore a non-linear class of models must be used to capture this change in variance in the error terms.

Engle (1982) introduced a model known as the Autoregressive Conditional Heteroscedasticity (ARCH) model in an effort to capture time varying variance. An ARCH(1) is shown in Equation 1.

$$\sigma_t^2 = \omega_0 + \alpha_1 u_{t-1}^2 \tag{1}$$

A general ARCH model can be described as follows:

$$y_t = \mu_t + \epsilon_t \tag{2}$$

$$\epsilon_t = e_t \sigma_t \tag{3}$$

$$e_t \sim N(0,1) \tag{4}$$

The general model consists of a time varying dependant variable y_t which can be described by a conditional mean equation μ_t and residuals ϵ_t. The specification of the mean equation can take any form. In this study an ARMA process will be used to describe the mean. The residuals are further modelled as the product of a random variable e_t and the square root of variance. e_t is normally distributed with mean zero and unit variance. The conditional variance σ_t^2 can take many different forms and is sometimes denoted as h_t^2. The form of the model given above is a convenient presentation of an ARCH model for simulation purposes.

2.2.2 GARCH Models

Engle's work laid down the framework for other variations of the model such as the Generalised ARCH (GARCH) model which has the current variance dependent on past values of variance.

$$\sigma_t^2 = \omega_0 + \alpha_p u_{t-1}^2 + \beta_q \sigma_{t-1}^2 \tag{5}$$

There are many denominations of the GARCH and ARCH model that range in complexity and that have various unique properties:

1. **Glosten-Jagannathan-Runkle GARCH (GJR-GARCH)**- This group of models is similar to the GARCH whereby the future variance depends on past lagged variance values, but the difference is that it includes a term that takes asymmetry into account. In doing so, positive and negative shocks have different effects on future variance.

2. **Exponential GARCH (EGARCH)**- The log of the variance is used. This has several advantages over the GARCH models in that there are no constraints that the parameters of the model need to be positive. This class of model can also account for asymmetry.

3. **Asymmetric GARCH (AGARCH)**- The AGARCH is similar to to the GARCH except that the future variance depends on the non-squared previous shocks.

4. **Nonlinear Asymmetric (NAGARCH)**- The NAGARCH is similar to to the AGARCH except that it is non-linear where the past residuals and volatility are squared together.

2.3 Superior Predictive Ability

Superior Predictive Ability (SPA) tests whether a particular set of alternative forecasts is superior to a benchmark forecast. There are many variations of the SPA. White (2000) introduced a framework for SPA which he calls the Reality Check (RC). RC is based on hypothesis testing whereby the predictive abilities of a test forecast is defined by the expected loss between the forecast and a benchmark. The framework encompasses m number of alternative forecasts of which the number of forecasts remain the same.

2.4 Model Confidence Set

The Model Confidence Set overcomes the problem of selecting the 'best' model. Instead the MCS seeks to find the group of models that are equally likely to be superior. It may be that the MCS only selects a single model as being superior to the others. The methodology used by the MCS is illustrated in Figure 1. The selection process begins with the allocation of the initial set of models M_0 to the set M. A hypothesis test for equal predictive ability (EPA) is performed on the set of models M using equivalence testing with a confidence level $1 - \alpha$. If the null hypothesis that the models are equally good is rejected, an elimination rule is employed to remove an inferior model from the group. The process is then repeated with the test for EPA. When the null is not rejected, the remaining set of models is the Model Confidence Set $M^*_{1-\alpha}$.

2.5 Difference between SPA and MCS

The SPA and MCS frameworks are very similar because they both employ hypothesis testing. However, the MCS has the following advantages over SPA:

1. MCS does not compare the test models to a benchmark.

2. MCS characterizes the entire set of models whereas SPA is an individual relative comparison between the benchmark and each test forecast.

3. The hypothesis tests for MCS are relatively simple compared to SPA where com-

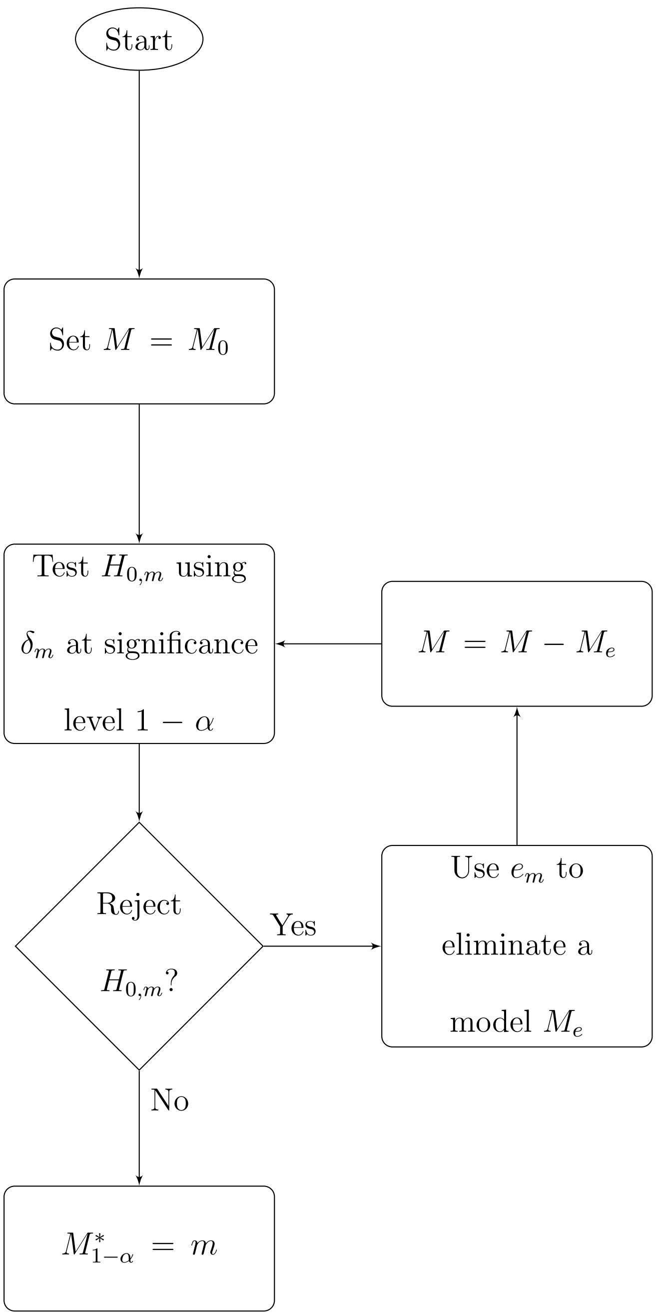

Figure 1: Flowchart for MCS process.

posite hypotheses can be used.

4. It eliminates models if it is found that the current set of models are not 'equally as good'. SPA does no eliminate models.

2.6 Combining Several Models

A question that arises from the MCS is what should be done with the set of models that it produces? Samuels & Sekkel (2013) presented a paper that discusses the various ways of combining models; equal weights and inverse MSE weights. The results of a MCS exercise selects a natural way of combining the models. The MCS presents a group of models that can not be significantly rejected as being equally good. Therefore the most natural way of combining the models is to use equal weighting.

2.7 Forecast Evaluation Tools

There are numerous techniques that are used to evaluate the out-of-sample performance of volatility models. A popular technique that is used is a Mincer-Zarnowitz regression, $r_t^2 = a + bh_t^2 + u_t$. The performance of the variance model, h_t^2 is quantified by the R^2 value which is a measure of how well the dependant variable and estimated coefficients 'fit' the independent variable. Engle & Patton (2001) highlighted a limitation of the Mincer-Zarnowitz regression by explaining that the heteroscedasticity in the returns are further amplified when the returns are squared. Therefore parameters and coefficients are inefficiently estimated and a good R^2 could be misleading through artificially generated coefficients.

As a simple measure of performance, the Mean Squared Error (MSE) shall be used as the loss function to determine the relative performance of the MCS, Random Walk and GARCH (1,1) to the true variance.

$$MSE = \frac{1}{n}\sum_{t=1}^{n}(\sigma_t^2 - h_t^2)^2 \tag{6}$$

Where σ_t^2 is the realized variance and h_t^2 is the forecasted volatility series.

2.8 Realized Variance

As a substitution for the true variance, the realized variance shall be calculated using :

$$RV_t = \sigma_t^2 = \sum_{i=1}^{m} r_{i,t}^2 \tag{7}$$

Where $r_{i,t}^2$ is the inter-day returns. This is a simplification of the realized variance equation which uses the intra-day returns. A more accurate realized variance may be calculated using interpolation and assuming a distribution function to estimate the intra-day returns. A Fourier Method can also be used to calculate daily volatility [1]. For simplicity and to keep the number of observations in the simulation to a reasonable number, Equation 7 will be used.

3 Chapter 3

3.1 Methodology

The following procedure was used to determine the performance between the MCS, Random Walk and GARCH (1,1):

1. The daily closing prices for several JSE indices using McGregorBFA were obtained.
2. The log return for each index was determined.
3. The mean process of the returns were estimated using ARMA models. The order of the model was chosen to minimize the Akaike information criteria.

[1]The Fourier method involves the integration of the time series data using the Fourier Inversion formula to reconstruct the volatility series between price samples.

4. After simulating the the ARMA model, the residuals were extracted and analysed. The following tests were conducted on the residuals:

   (a) A test for heteroscedasticity.

   (b) A test for normality.

5. After the residual tests were satisfied, several GARCH and ARCH models were estimated on the residuals using maximum likelihood techniques. The results of the normality test determined the distribution that was used to maximize the likelihood estimate.

6. For each volatility model, a one step-ahead out-of-sample forecast was made up to 30 days ahead.

7. The forcasted time series were then used in a MCS function with a confidence level of 95%.

8. The remaining models forming the MCS were averaged (equally weighted) and then compared to a forecast of the returns using a Random Walk process and GARCH (1,1) process.

The entire exercise was conducted in MATLAB R2012a. The volatility model estimation functions and the MCS function is part of the MFE toolbox created by Kevin Sheppard. The one-step-ahead forecasting functions were created by the the author. EViews was used to determine the specification of the mean process for each index.

3.2 Data

The indices chosen were: JSE General Mining, JSE Allshare, JSE Top 40, JSE Midcap, JSE Industrial 25 and JSE Gold. The data set consists of the daily closing prices for each of the selected indices. The time span of the data is from 11 December 2003 to 11 December 2013 which includes 2502 observation in each index. The returns were calculated using $r_t = \log(P_t) - \log(P_{t-1})$. The data was split into two major categories; an estimation period and an evaluation period. The first 2400 observations were used to estimate the volatility models. The remaining 101 observations were used to calculate the realized volatility which forms a basis for comparison in Equation 6. The descriptive statistics for the log returns are shown in Table 1. The statistics show that none of the indices have normally distributed returns. The mean return across the board is close to zero. All of the indices have around the same kurtosis value indicating some fatness in the tail. The standard deviation, which is a measure of the average volatility of the indices over 10 years shows that the General Mining and Gold indices have higher volatility than the other indices.

The line graphs for the closing prices and log return data for each index is displayed in Appendix C. The closing prices for all the indices, except Gold, show an overall exponential increase in the value of the index between the period 2003 to 2008. The global financial crisis can be seen in the dramatic fall in the indices towards the second half of 2008. Since the event of the crisis, most of the indices including Top 40, Industrial 25, Midcap and Allshare have seen a steady recovery. Post crisis, the General Mining index has stagnated due to frequent mining sector strikes starting with the Marikana strikes in August 2012. This is in contrast with the JSE Allshare index which has seen significant growth since the 2008 crisis. The growth of the Allshare index for 2013 was 21.3%. Throughout the past 10 years, the Gold index has followed a seesaw movement that averages around the 2250 mark.

Table 1: Descriptive Statistics for JSE Indices

Characteristic	General Mining	Allshare	Top 40	Mid Cap	Industrial 25	Gold
Median	0.0433	0.1090	0.1263	0.1066	0.1377	-0.0944
Maximum	15.3771	6.8340	7.7069	4.7119	7.1729	16.3170
Minimum	-13.7654	-7.5807	-7.9594	-5.6325	-6.7940	-13.2032
Std. Dev.	2.2450	1.2971	1.4196	0.8177	1.1878	2.3106
Skewness	0.1169	-0.1995	-0.1353	-0.6262	-0.1619	0.2227
Kurtosis	7.1878	6.6027	6.4589	7.4072	6.0429	6.2145
Jarque-Bera	1759	1313	1203	2099	936	1053
Probability	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
Sum	114.6920	142.5484	139.8860	155.6476	195.7084	-62.3284
Sum Sq. Dev.	12091	4036	4834	1604	3384	12807

The return graphs for all the indices appear to fluctuate around a constant level. Volatility clustering is evident in all the data with periods of large changes grouped together and small changes tending to also group together. In 2008 the indices had the largest fluctuations both positive and negative. From the the visual observation there are indications of ARCH type effects in the data. This will be confirmed with ARCH tests.

3.3 ARMA Models

To determine the the mean process for the returns, an iterative approach was used to minimize the Akaike Information Criteria. The ARMA specification for each of the JSE indices is shown in Table 2. There are other information criteria that may be used which may yield a different specification for the mean process. The Schwarz Bayesian criterion usually produces a specification that is at least as small as Akaike as it imposes stiffer penalty terms. The residuals for the specification were extracted and used in a series of tests to better understand the characteristics of the residuals and to help model volatility as accurately as possible.

Table 2: ARMA Specification for each Index

Index	ARMA Model(p,q)
JSE General Mining	(8,8)
JSE Allshare	(7,10)
JSE Top 40	(10,10)
JSE Mid cap	(8,11)
JSE Ind 25	(7,8)
JSE Gold	(5,6)

3.4 Statistical Tests

If there are any volatility effects in a set of data, they would appear in the residuals or innovations of a model. Thus it is important to determine if there are indeed ARCH effects in the residuals. Once this has been determined the distribution of the residuals need to be determined in order to ensure that the volatility models that are estimated are as accurate as possible. The tests that were performed on the residuals are as follows:

1. A heteroscedasticity test (Engle's (1982) ARCH test).

2. A normality test (Jarque-Bera Test).

3.4.1 ARCH Test

Engle's ARCH test involves a regression on the the squared residuals and the null hypothesis that all the coefficients are equal to zero. The alternative is that all the coefficients are non-zero.

$$H_0 : \gamma_1 = 0 \; and \; \gamma_2 = 0 \; and \; \gamma_3 = 0 \; and \; ... \; and \; \gamma_q = 0$$

$$H_1 : \gamma_1 \neq 0 \; and \; \gamma_2 \neq 0 \; and \; \gamma_3 \neq 0 \; and \; ... \; and \; \gamma_q \neq 0$$

The results of the ARCH tests are given in Table 3. It can be concluded that the null hypothesis for all the indices can be confidently rejected and the residuals contain ARCH effects.

3.4.2 Normality Test

The test for normality involves the Jarque-Bera test. Figure 2 shows the distribution density of the residuals for the General Mining Index. Superimposed on the histogram is a normal distribution. From the descriptive statistics in Table 1, all the indices have

Table 3: ARCH Test Results

Index	F-Statistic (p-value)	TR^2 (p-value)
General Mining	112.9399 (0.000)	601.3529 (0.000)
Allshare	76.36306 (0.000)	580.3400 (0.000)
Top 40	72.14724 (0.000)	555.6122 (0.000)
Mid cap	62.37589 (0.000)	534.7348 (0.000)
Industrial 25	53.93682 (0.000)	366.5598 (0.000)
Gold	48.40998(0.000)	259.6517 (0.000)

high Jarque-Bera values and low p-values. Thus the null hypothesis that the residuals are normally distributed is strongly rejected to four significant figures. A Student's T distribution with mean 0.0118, variance 5.012 and 5.39 degrees of freedom was fitted to the General Mining residuals and it can be seen that it offers a better fit. A similar process was followed for all the indices (See Appendix D) and the Student's T distribution gave a superior fit compared to the Normal distribution. Therefore during the estimation of the volatility models, a Student's T distribution was used in the maximum likelihood function.

3.5 Volatility Model Estimation

GARCH and ARCH models were estimated using the extracted residuals from the conditional mean process. Most of the models are specified using two lag parameters p and q. The general specification of each of the models was over a range of $p, q = 1, 2$. The following models were estimated using a Student's T distribution and the log likelihood function:

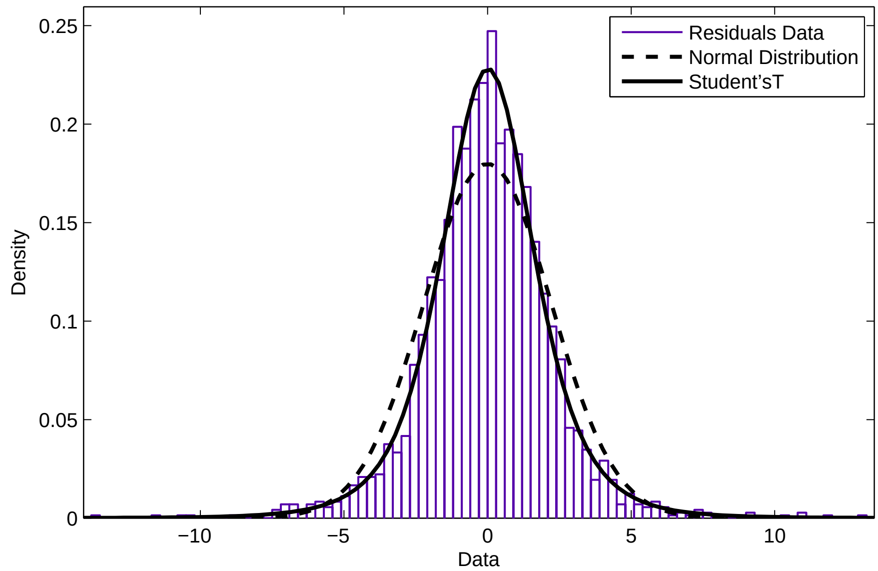

Figure 2: Histogram of ARMA (8,8) residuals

1. GARCH
2. GJR-GARCH
3. EGARCH
4. AGARCH
5. NAGARCH

There are a total of 20 estimated models for each index.

3.6 Forecasting

To forecast the volatility of each index using the various estimated models, functions were developed in MATLAB to calculate the one-step-ahead out-of-sample forecast. Using expected value theory, the first step ahead $E[\sigma_{t+1}^2]$ was determined with the in-sample residuals. Then, using an iterative process and updating the residuals for each advancement, 30 days of forecasted volatility was estimated for each model. It must be noted that for the one-step-ahead forecasting algorithm, a normal distribution was assumed. This assumption greatly simplifies the calculations. Under normality, E[e_{t-1}] =1 and for models that incorporate asymmetry such as the GJR-GARCH model, E[$I_{e_{t-1}} < 0$] can be expected to take place 1/2 of the time.

The Random Walk forecasts were made using a random generator to produce small excursions proportional to the standard deviation of the in-sample returns. As a starting point the last realized variance value was used as a platform for the excursions. The Random Walk in this simulation is akin to Brownian motion in the positive direction due to the square operator when calculating variance.

4 Chapter 4

4.1 MCS Empirical Results

The MCS function was initiated on all of the models with a confidence level of 95% (1-α, $\alpha = 0.05$) and 1000 iterations. The models indicated in Table 4 were in the MCS, $M^*_{95\%}$ after the inferior models were removed.

The results of the MCS selection process shows that for the General Mining index, all four estimated GARCH models were included in the MCS. The standard GARCH models create symmetrical responses of volatility for both negative and positive shocks e_{t-1}. This is in contrast to the GJR-GARCH and EGARCH which can account for possible asymmetric shocks. Since all four GARCH models were included it would appear that symmetrical models dominate asymmetric models for the JSE General Mining index.

Table 4: MCS Results

Index	Model contained in $M^*_{95\%}$
JSE General Mining	GARCH(1,1), GARCH (1,2), GARCH (2,1), GARCH (2,2), EGARCH(1,2)
JSE Allshare	EGARCH(2,1)
JSE Top40	EGARCH(2,1)
JSE Midcap	EGARCH(2,1)
JSE Ind25	EGARCH(2,1)
JSE Gold	EGARCH(2,1)

The inclusion of the EGARCH (1,2) model in the MCS for General Mining seems an oddity as it is the only model in the group that incorporates asymmetry. The coefficients for the EGARCH(1,2) included a negative γ term which indicates asymmetry was taken into account when the parameters were estimated. It is also unusual for the EGARCH model to be grouped with GARCH models as it can provide superior fits when compared to GARCH models. The EGARCH does not suffer from the parameter restrictions that are imposed on GARCH models. Since the log of the variance is modelled, the parameters can be negative. So it would be expected that the EGARCH not be classified together with any GARCH models in a group of equally superior models. The question is whether or not the inclusion of the EGARCH model will hamper the performance of the MCS models as a forecasting group for the General Mining index?

For the remaining JSE indices that were modelled, the EGARCH (2,1) was the only model selected by the MCS. Intuitively, an EGARCH model would be the most likely model to be selected as it does not suffer the same parameter restrictions as the other volatility models and it can incorporate asymmetry. This indicates that models that incorporate leverage effects are favoured for modelling volatility in the chosen JSE indices.

4.2 Forecast Evaluation

The MSE for the three forecasting models; Random Walk, GARCH(1,1) and MCS are shown in Table 5 to 10 . The results show that, in general, the equally weighted combination of MCS models performs better than the GARCH (1,1) for time horizons greater than 10 days. The GARCH (1,1) model has a lower MSE value than the MCS for a 5-day time horizon in the case of the Top40, Midcap, Industrial and Allshare. Interestingly, for all time horizons, the MCS performs better than the GARCH (1,1) for indices that have high volatility. The two cases where this was evident was for the General Mining and Gold indices. From the descriptive statistics, these two indices had the highest standard deviation values.

As the MSE is calculated for different time horizons, it can be seen that the MSE increases at an accelerated rate. A general property of volatility is that it increases with $\sqrt{t}$ and variance increases with t. The iterative process of the one-step-ahead explains the exponential increase in MSE as past values are successively squared. This suggests that the power of volatility forecasting greatly diminishes with time. However, the Random Walk[2] shows resilience over time and is better than both the GARCH(1,1) and MCS by orders of magnitude at the 30-day forecast horizon in some cases.

[2]Due to the random number generation, the simulation was performed several times.

A possible reason for the better general performance of the MCS is the inclusion of an asymmetrical EGARCH model, which is a property that the GARCH models do not have. In addition, the natural logs used in an EGARCH model gives it a property that slows the acceleration of volatility as time increases. This is the reason that for time horizons greater than 10 days, the EGARCH has a lower MSE than the GARCH (1,1).

The MSE results lend credibility to the GARCH (1,1) as a popular volatility model used by analysts. For four out of the six simulated indices, the GARCH (1,1) performed the best for 5-day-ahead forecasts.

Table 5: MSE for JSE General Mining

	JSE General Mining		
Days	Random Walk	MCS	GARCH (1,1)
5	23.66	1.61	2.35
10	110.79	2.77	3.08
15	221.82	3.31	2.95
20	393.67	28.79	29.50
25	664.47	115.32	141.80
30	970.58	1.35E+33	7.41E+33

Table 6: MSE for JSE Allshare

	JSE Allshare		
Days	Random Walk	MCS	GARCH (1,1)
5	6.93	54.21	5.54
10	14.54	61.67	288.92
15	33.62	80.28	4.31E+47
20	56.16	102.53	Inf
25	94.69	170.65	Inf
30	132.56	387.62	Inf

Table 7: MSE for JSE Top40

	JSE Top40		
Days	Random Walk	MCS	GARCH (1,1)
5	9.16	102.14	8.70
10	21.60	134.04	11354.06
15	35.57	228.36	2.65E+97
20	47.67	503.47	Inf
25	82.44	2965.39	Inf
30	127.52	305281.80	Inf

Table 8: MSE for JSE Midcap

	JSE Midcap		
Days	Random Walk	MCS	GARCH (1,1)
5	1.59	20.09	1.51
10	2.59	23.27	2.46
15	6.46	26.72	5.92
20	9.95	31.29	10919.27
25	15.14	42.38	1.46E+106
30	20.14	73.34	Inf

Table 9: MSE for JSE Industrial25

	JSE Industrial 25		
Days	Random Walk	MCS	GARCH (1,1)
5	12.94	482.02	44.07
10	18.16	616.00	7.02E+16
15	39.61	1071.09	Inf
20	61.13	3102.98	Inf
25	88.34	66454.37	Inf
30	113.44	1.16E+10	Inf

Table 10: MSE for JSE Gold

	JSE Gold		
Days	Random Walk	MCS	GARCH (1,1)
5	565.33	536.93	557.57
10	325.81	293.66	309.75
15	349.84	216.87	232.59
20	1092.55	1256.24	1293.52
25	1160.23	1038.58	1075.86
30	1214.62	978.17	1036.73

5 Chapter 5

5.1 Future Work

As an exercise for future work, the impact of the number of models used in the the MCS and the resulting performance can be simulated. In this project only 20 volatility models were used per index and it is hypothesized that if a greater number of models with a wider variety of properties and lags are used in the initial population of models, the MCS would offer even greater forecasting power.

The MCS offers an exciting platform for model manipulation and forecasting. The output of the MCS process is a set of models that are 'equally good' and this opens up opportunities on how to combine the models to obtain the best forecasting performance. A simple equally weighted combination of the MCS models was used in this project but another possible way of combining the models is to use the p-values of the MCS function to weight each model.

5.2 Conclusion

A full simulation exercise was conducted whereby the returns of several JSE indices were tested and modelled using a variety of GARCH models. The Model Confidence Set was used to trim the total group of models down to subset which represented the models that were equally good. The EGARCH family of models was shown to be the most consistently chosen by the MCS as the superior model. The Mean Squared Error model evaluation revealed that for time horizons greater than 10 days, the MCS performed better than the GARCH (1,1) model. The GARCH (1,1) had superior 5-day-ahead forecasting power for most of the indices. An interesting result is that the MCS performed better than the GARCH (1,1), for all time horizons, when the index had high volatility as in the Gold and General Mining index. This suggests that the power of the MCS to forecast is greater when the modelled quantity (in this case volatility) is more evident in the data.

References

Aboura, S. & Wagner, N. (2013), 'Extreme asymmetric volatility, leverage, feedback and asset prices', *Université de Paris Dauphine* .

DeIonno, P. (2013), 'Lonmin share price shuns amcu strike threats', *Business Report* p. 17.

Ederington, L. H. & Guan, W. (2006), 'Measuring historical volatility', *Journal of Applied Finance* **16**, 5.

Engle, R. F. (1982), 'Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation', *Econometrica* **50**, 987.

Engle, R. F. & Patton, A. J. (2001), 'What good is a volatility model?', *Quantative Finance* **1**, 237–245.

Hansen, P. R. & Lunde, A. (2005), 'A forecast comparison of volatility models: Does anything beat a garch(1,1)?', *Journal of Applied Economics* **20**, 873–889.

Hansen, P. R., Lunde, A. & Nason, J. M. (2003), 'Choosing the best volatility models: The model confidence set approach', *Brown University Department of Economics* **Working Paper No. 2003-05**.

Jordan, B. D., Jr, T. W. M. & Dolvin, S. D. (2012), *Fundamentals of Investments*, 6th edn, McGraw-Hill.

Luenberger, D. G. (2009), *Investment Science*, international edn, Oxford University Press.

Masset, P. (2011), 'Volatility stylized facts', *Université de Fribourg and Ecole Hôtelière de Lausanne* .

Nortje, B. (2013), 'Jse ends high despite volatility', *Business Day* .

Ramashala, T. (2011), 'Option-implied volatility as a predictor of realized volatility in derivative markets', *University of Pretoria* .

Samouilhan, N. & Shannon, G. (2008), 'Forecasting volatility on the jse', *Investment Analysts Journal* **67**.

Samuels, J. D. & Sekkel, R. M. (2013), 'Forecasting with many models: Model confidence sets and forecast combination', *Bank of Canada* **Working Paper**.

Venter, R. G. (2003), 'Pricing options under stochastic volatility', *University of Pretoria* .

White, H. (2000), 'A reality check fo data snooping', *Econnmetrica* **68**, 1097–1126.

Zhou, P. (2007), 'Essays on financial asset return volatility', *University of California* .

A Appendix A: MCS Simulation

Table 11: MCS and p-values

No.	Model	General Mining	Allshare	Top40	Midcap	Ind25	Gold
1	AGARCH11	0	0	0	0	0	0
2	AGARCH12	0	0	0	0	0	0
3	AGARCH21	0	0	0	0	0	0
4	AGARCH22	0	0	0	0	0	0
5	EGARCH11	0.041	0	0	0.005	0	0
6	EGARCH12	0.268*	0	0	0	0	0
7	EGARCH21	0	1*	1*	1*	1*	1*
8	EGARCH22	0.041	0.001	0.037	0	0.037	0
9	GARCH11	0.268*	0	0	0	0	0
10	GARCH12	1.00*	0	0	0	0	0
11	GARCH21	0.268*	0	0	0	0	0
12	GARCH22	0.268*	0	0	0	0	0
13	GJRGARCH11	0	0	0	0	0	0
14	GJRGARCH12	0	0	0	0	0	0
15	GJRGARCH21	0	0	0	0	0	0
16	GJRGARCH22	0	0	0	0	0	0
17	NAGARCH11	0	0	0	0	0	0
18	NAGARCH12	0	0	0	0	0	0
19	NAGARCH21	0	0	0	0	0	0
20	NAGARCH22	0.003	0	0	0	0	0

Table 11 represents the Model Confidence Set simulation results. The p-values that are marked with one $*$ indicate models that are included in $M^*_{95\%}$. The MCS was simulated with the R method.

B Appendix B: GARCH Models

Table 12: Alternative GARCH models

ARCH:	$\sigma_t^2 = \omega + \sum_{i=1}^q \alpha_i \epsilon_{t-i}^2$
GARCH:	$\sigma_t^2 = \omega + \sum_{i=1}^q \alpha_i \epsilon_{t-i}^2 + \sum_{j=1}^p \beta_j \sigma_{t-j}^2$
AGARCH	$\sigma_t^2 = \omega + \sum_{i=1}^q [\alpha_i \epsilon_{t-i}^2 + \gamma_i \epsilon_{t-i}] + \sum_{j=1}^p \beta_j \sigma_{t-j}^2$
NAGARCH	$\sigma_t^2 = \omega + \sum_{i=1}^q \alpha_i (\epsilon_{t-i} + \gamma_i \sigma_{t-i})^2 + \sum_{j=1}^p \beta_j \omega_{t-j}^2$
GJRGARCH	$\sigma_t^2 = \omega + \sum_{i=1}^q [\alpha_i + \gamma_i I_{\{\epsilon_{t-i} > 0\}}] \epsilon_{t-i}^2 + \sum_{j=1}^p \beta_j \sigma_{t-j}^2$
EGARCH	$\ln(\sigma_t^2) = \omega + \sum_{i=1}^q [\alpha_i \epsilon_{t-i} + \gamma_i (\lvert e_{t-i} \rvert - E \lvert e_{t-i} \rvert)] + \sum_{j=1}^p \beta_j \ln(\sigma_{t-j}^2)$

C Appendix C: Graphs

C.1 JSE General Mining Index

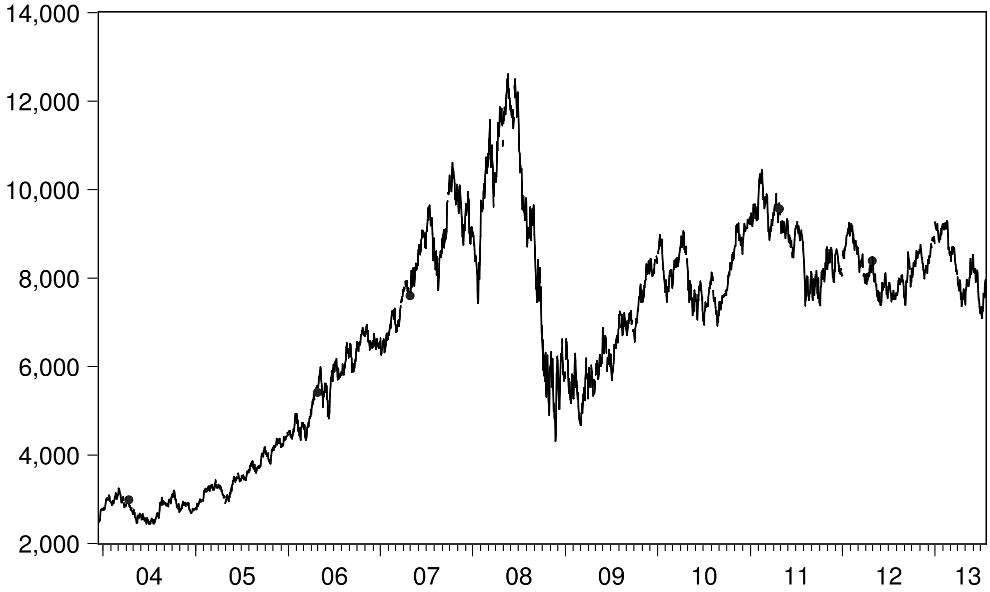

Figure 3: Price data for JSE General Mining Index

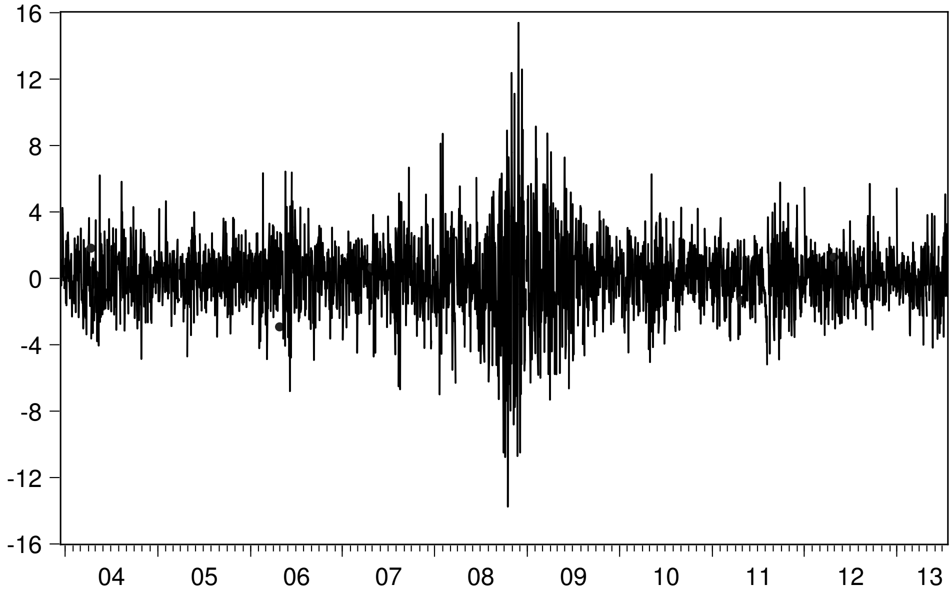

Figure 4: Log returns for JSE General Mining Index

C.2 JSE Allshare

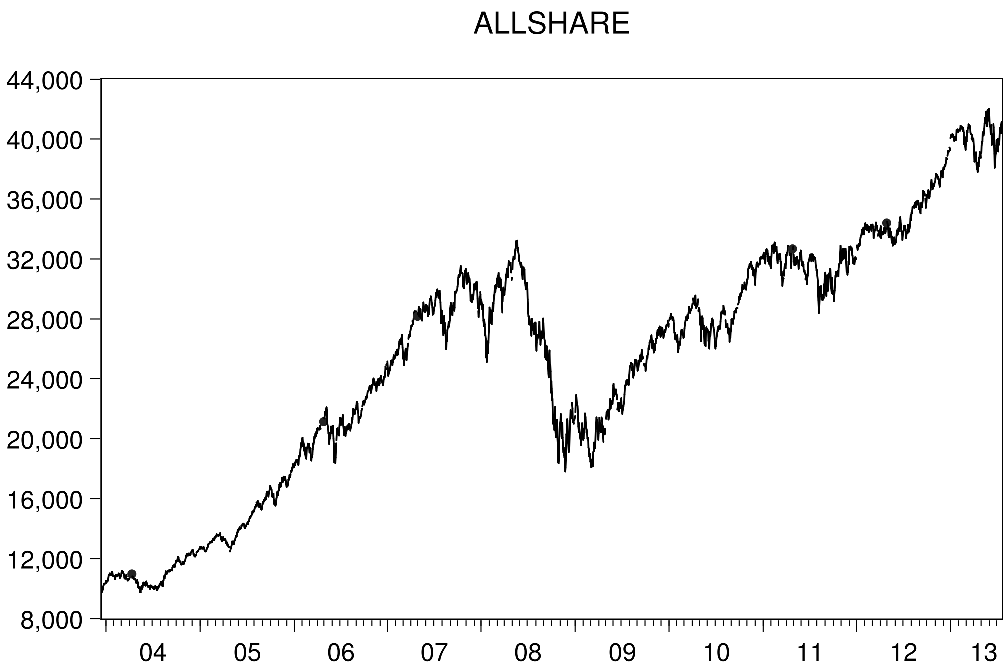

Figure 5: Price data for JSE Allshare Index

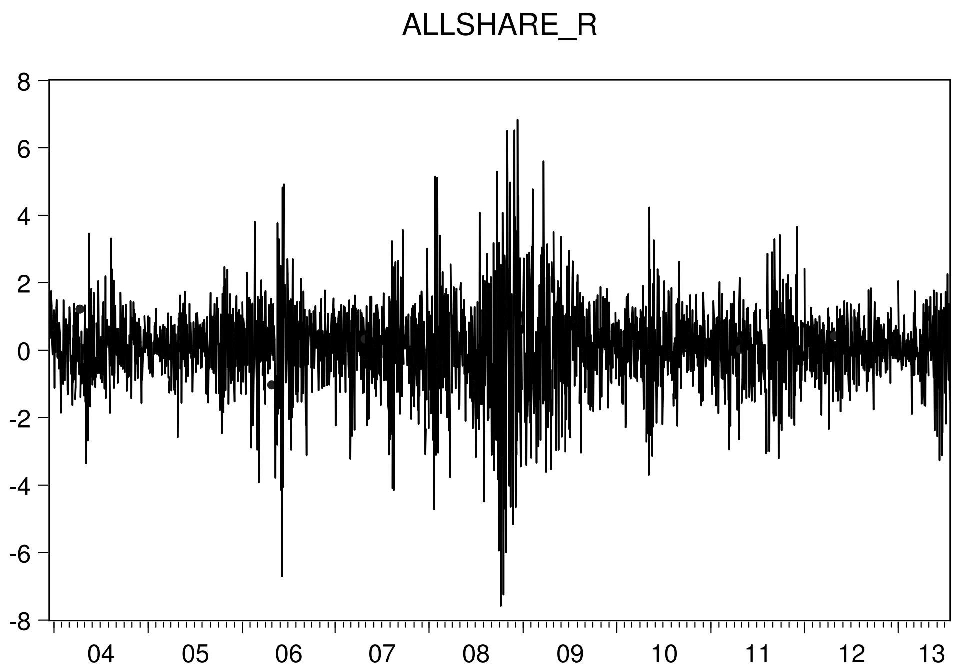

Figure 6: Log returns data for JSE Allshare Index

C.3 JSE Top 40

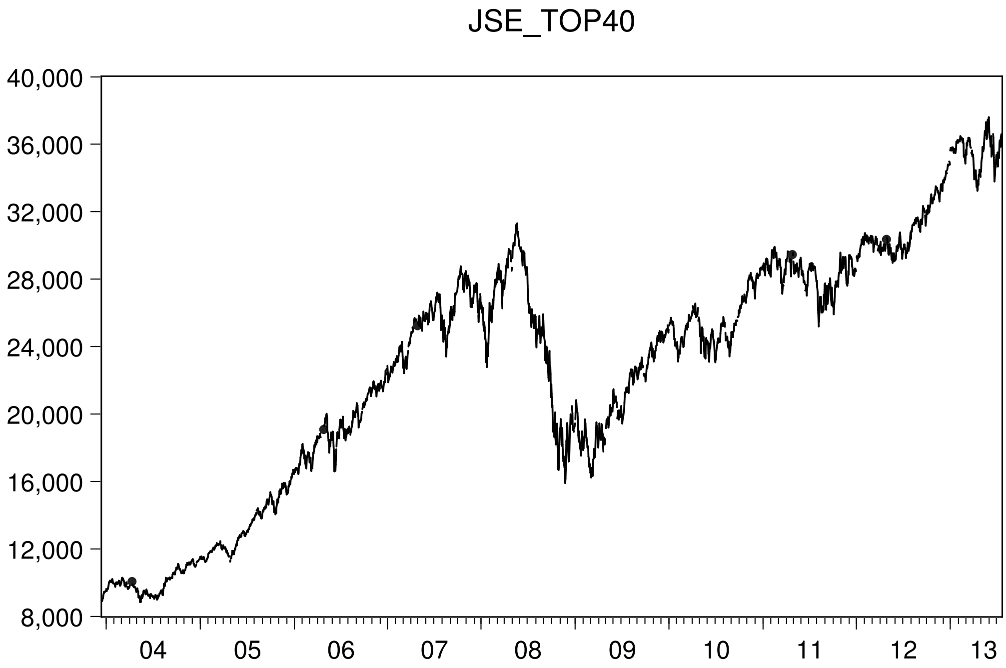

Figure 7: Price data for JSE Top 40 Index

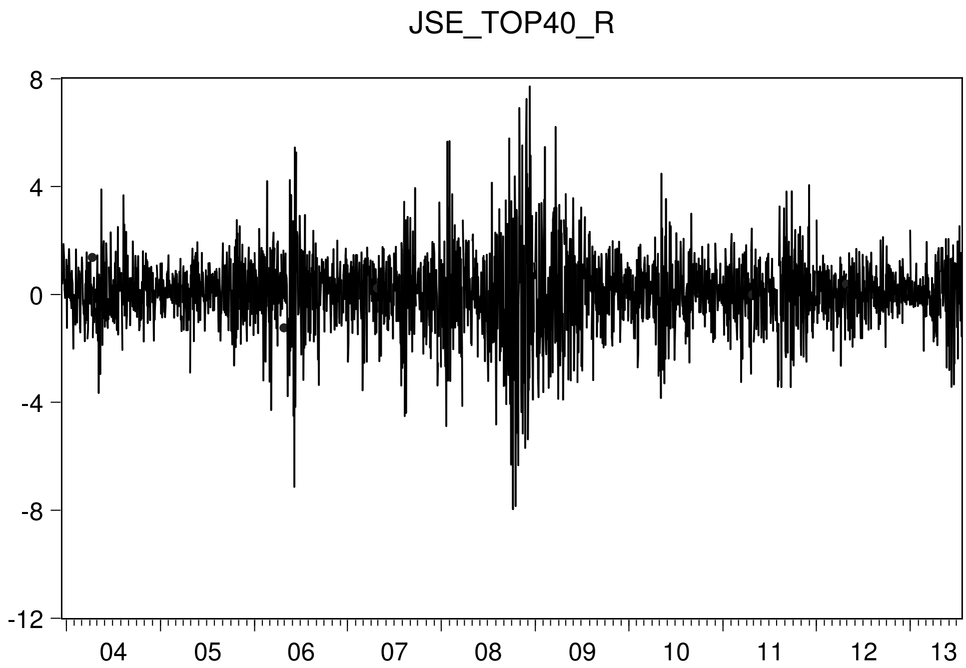

Figure 8: Log returns data for JSE Top 40 Index

C.4 JSE Mid Cap

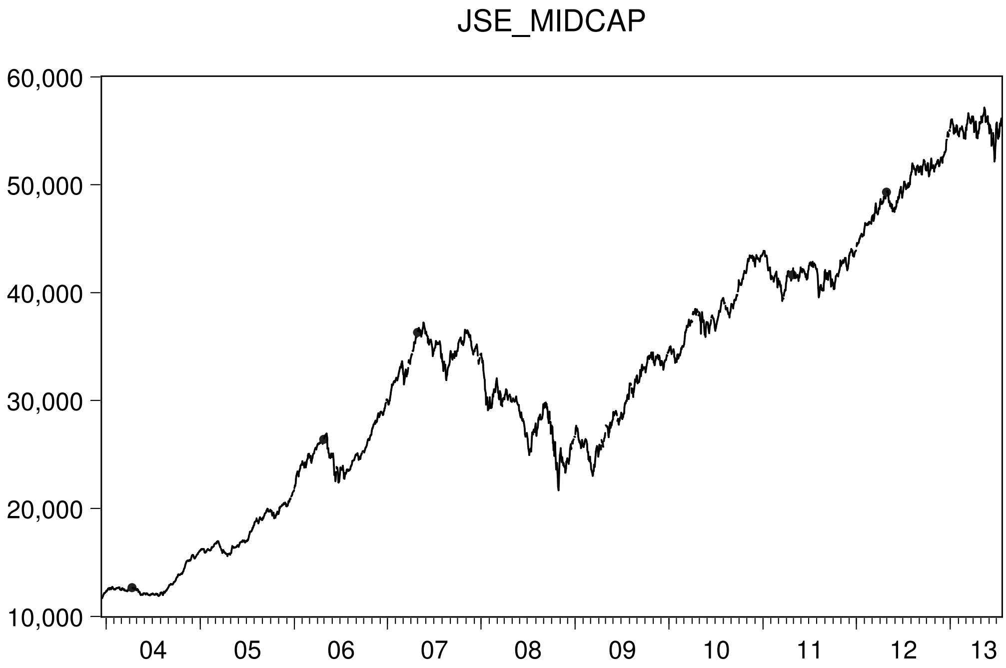

Figure 9: Price data for JSE Midcap Index

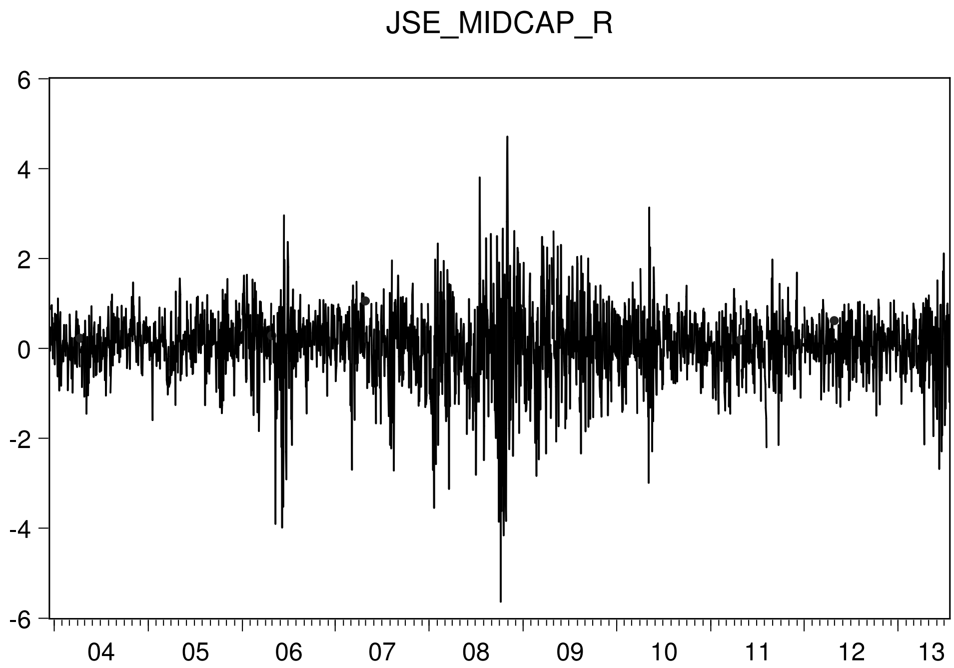

Figure 10: Log returns data for JSE Midcap Index

C.5 JSE Industrial 25

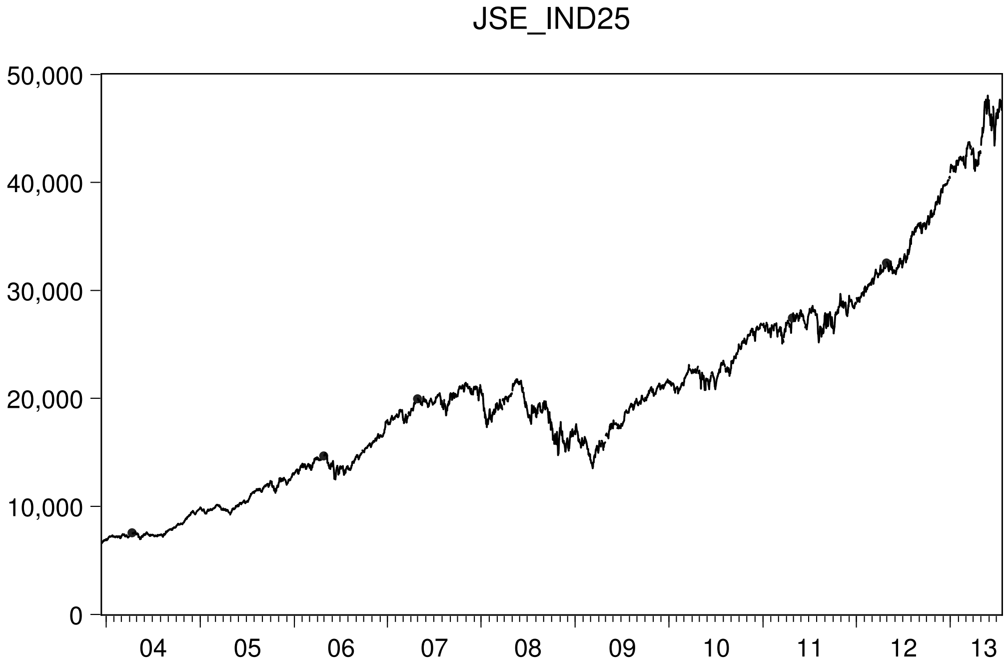

Figure 11: Price data for JSE Industrial 25 Index

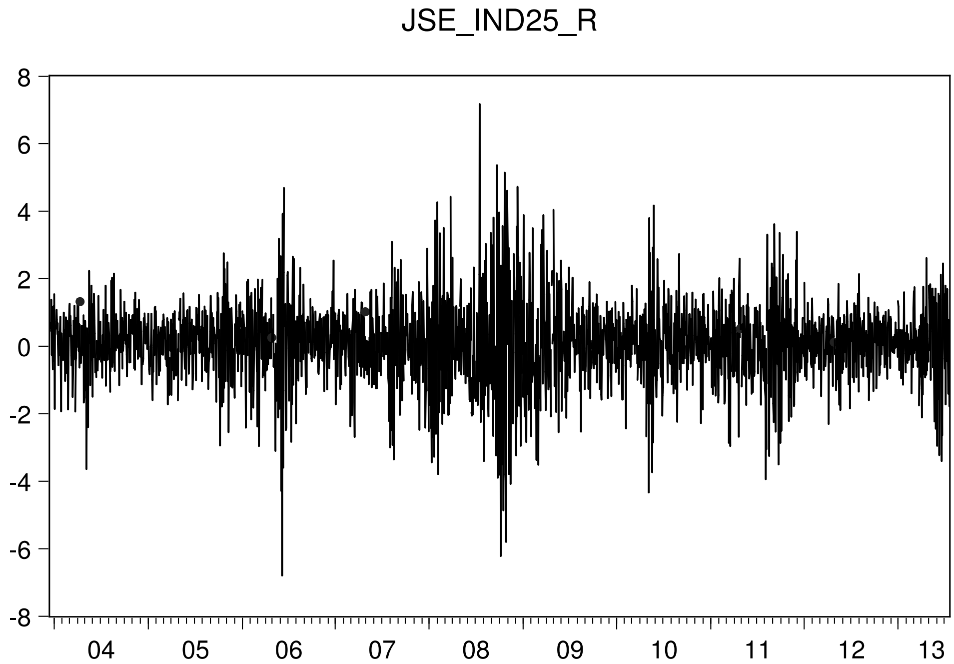

Figure 12: Log returns data for JSE Industrial 25 Index

C.6 JSE Gold

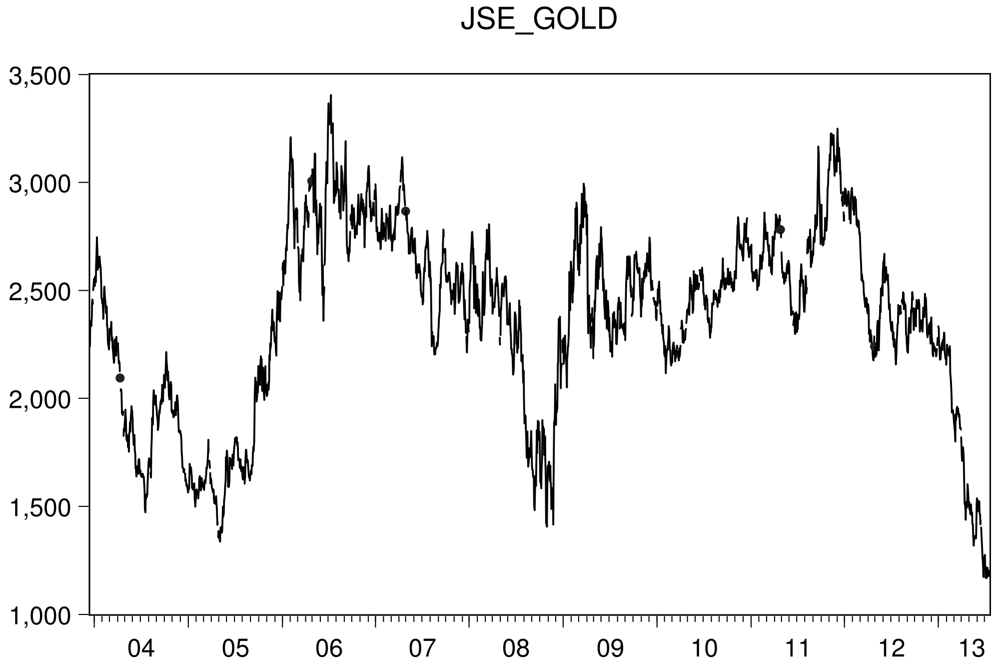

Figure 13: Price data for JSE Gold Index

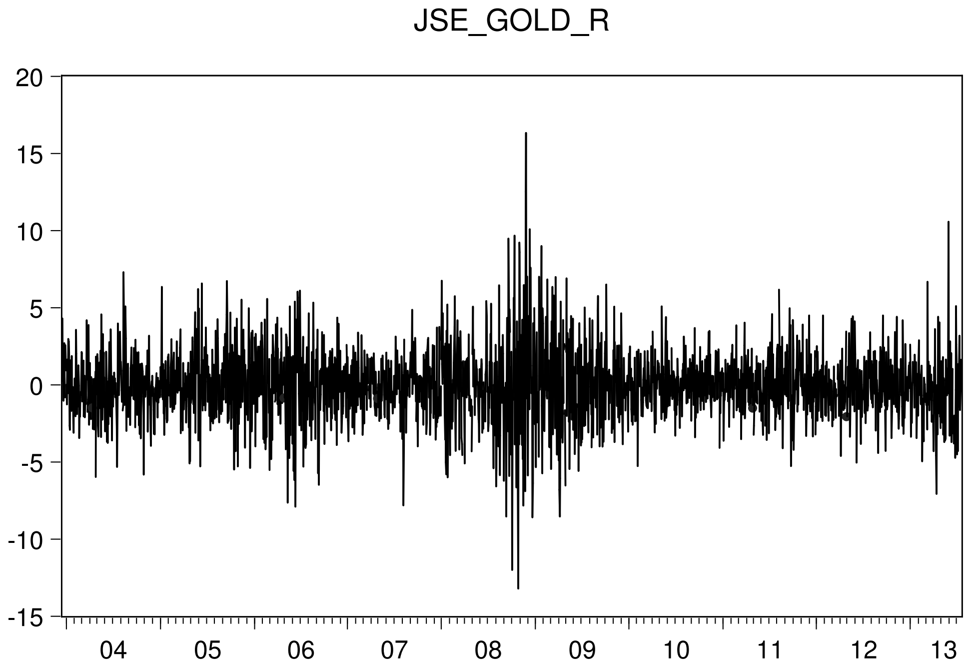

Figure 14: Log returns data for JSE Gold Index

D Appendix D: Normality Tests

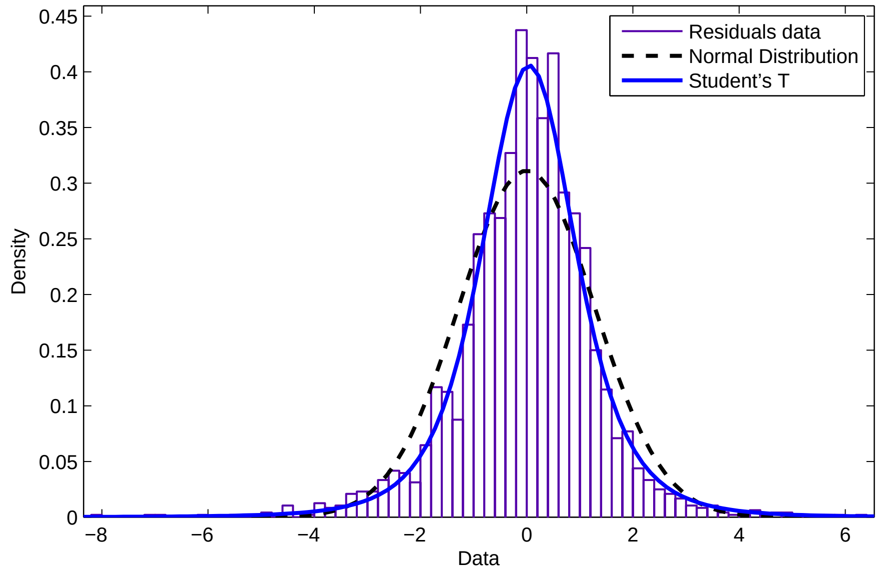

Figure 15: Histogram of ARMA (7,10) residuals for JSE Allshare returns. Superimposed on the histogram is a normal distribution and a Student’s T distribution with mean = 0.0423164, variance = 1.73323, degrees of freedom = 3.96455.

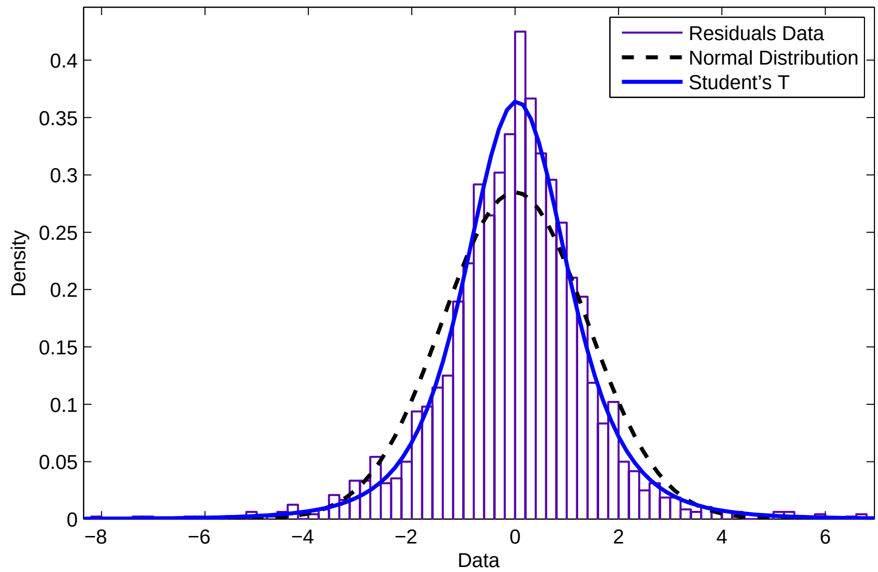

Figure 16: Histogram of ARMA (10,10) residuals for JSE Top 40 returns. Superimposed on the histogram is a normal distribution and a Student’s T distribution with mean = 0.0322511, variance = 2.04206, degrees of freedom = 4.18478.

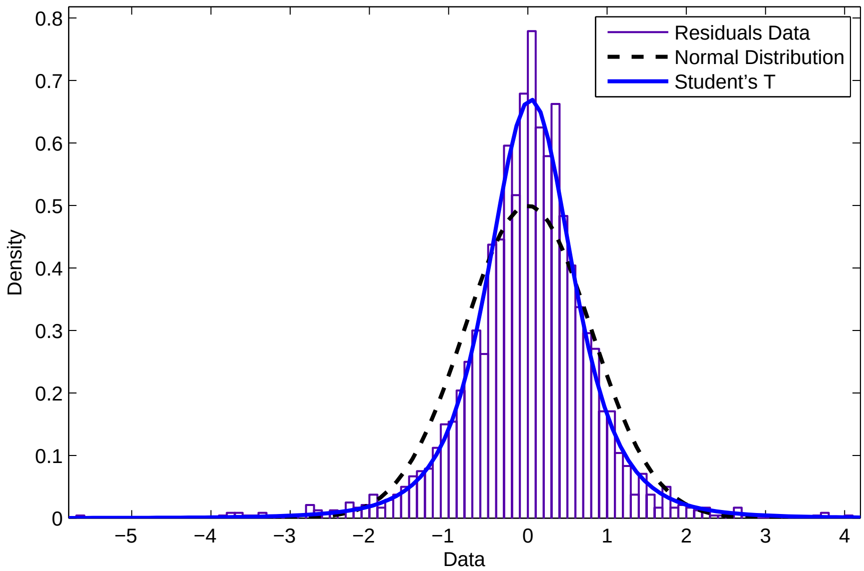

Figure 17: Histogram of ARMA (8,11) residuals for JSE Midcap returns. Superimposed on the histogram is a normal distribution and a Student's T distribution with mean = 0.035681, variance = 0.681355, degrees of freedom = 3.66929.

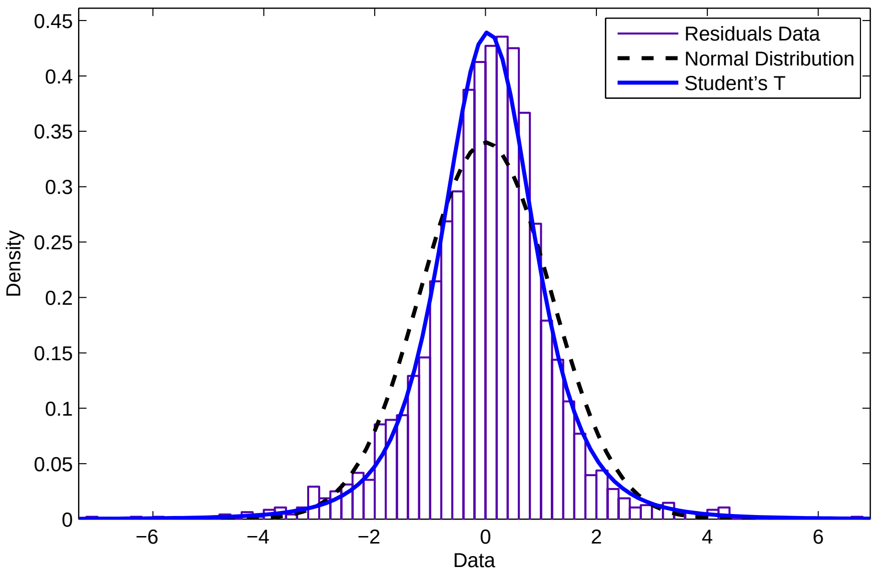

Figure 18: Histogram of ARMA (7,8) residuals for JSE Industrial 25 returns. Superimposed on the histogram is a normal distribution and a Student's T distribution with mean = 0.0447531, variance = 1.46776, degrees of freedom = 3.96316.

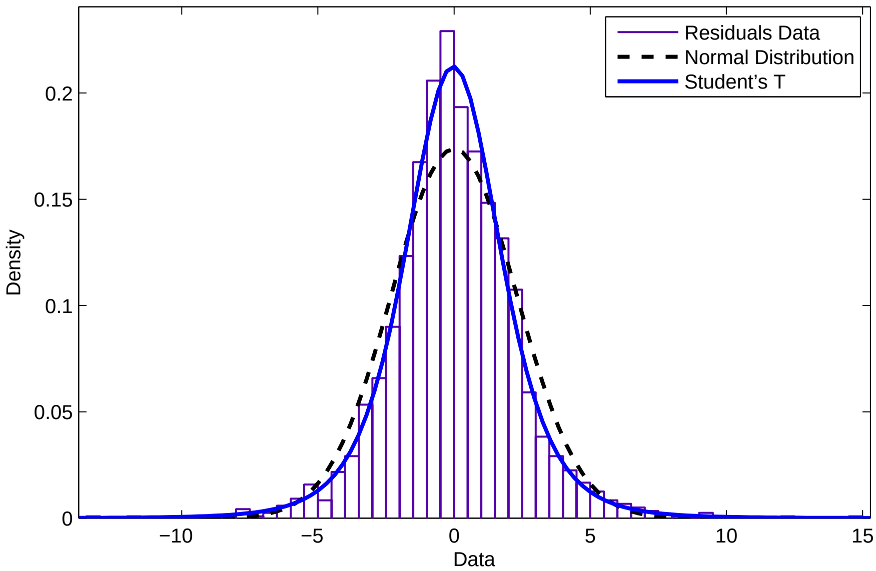

Figure 19: Histogram of ARMA (5,6) residuals for JSE Gold returns. Superimposed on the histogram is a normal distribution and a Student's T distribution with mean = -0.0288665, variance = 5.35379, degrees of freedom = 4.93471.