

https://dspace.org/

The dynamics of market efficiency: testing the adaptive market hypothesis in South Africa

Seetharam, Yudhvir

2016

http://hdl.handle.net/10539/21982

Seetharam, Yudhvir (2016) The dynamics of market efficiency: testing the adaptive market hypothesis in South Africa, University of the Witwatersrand, Johannesburg, <http://wiredspace.wits.ac.za/handle/10539/21982>

Downloaded from a DSpace repository. https://dspace.org/

The dynamics of market efficiency:
Testing the Adaptive Market Hypothesis in South Africa

Yudhvir Seetharam

School of Economic and Business Sciences

University of the Witwatersrand

Johannesburg

South Africa

The dynamics of market efficiency:
Testing the Adaptive Market Hypothesis in South Africa

Yudhvir Seetharam
0705581W

Supervisors: Prof. C. Auret

School of Economic and Business Sciences

Prof. T. Celik

School of Computer Science and Applied Mathematics

A thesis submitted to the School of Economic and Business Sciences, Faculty of Commerce, Law and Management, University of the Witwatersrand in fulfilment of the requirements for the degree of Doctor of Philosophy (Ph.D).

Johannesburg, South Africa

June 2016

DECLARATION

I, Yudhvir Seetharam, declare that this thesis is my own unaided work. It is submitted in fulfilment of the requirements for the degree of Doctor of Philosophy (Ph.D) at the University of the Witwatersrand, Johannesburg. It has not been submitted before for any degree or examination at this or any other university.

_____

Yudhvir Seetharam
June 2016

ACKNOWLEDGEMENTS

"If I have seen further it is by standing on the shoulders of giants,"
(Isaac Newton)

Reading towards a doctorate marks an important milestone in the life of the student - it is both rewarding to embark on a journey of discovery; yet equally frightening to not know where your route will lead. I have only been able to complete mine by having a strong foundation of guidance, support, confusion and sometimes apprehension from various individuals. I would like to thank those who have helped me grow throughout this process.

To my supervisors, past and present, Prof. Christo Auret, Mr. James Britten, CFA, Prof. Turgay Celik and Prof. Christian Omlin. I have matured immensely under your tutelage. Prof. Auret, I hope that I now know the "what and the why". I will always appreciate the role you have played in my student and academic lives (especially the lesson to never argue with one's supervisor). Prof. Celik, I hope that I now know how to think analytically.

A special mention must be made to Prof. Robert Vivian, whose way of thinking will always astound me. Thank you for guiding me to make a contribution to the academic community and for being a mentor.

To my dad, my late mum and my siblings. You have showed me what it means to dedicate your life to a passion and a pursuit of happiness. Mum, while it has been a few months since you passed, I know you will be smiling down on me when I graduate, relieved that your youngest has finally finished university after nine long years.

To my close friends, Sipe Bacela, David McClelland, Craig Nichol, Daniel Page and Magen Ponnan. You either compelled me to work harder (smarter), nagged me to get this done or argued with me to improve the idea.

Last, to the colleagues I have grown to know over the years either as FNB staff, Wits staff, or my own postgraduate students. Those "small" acts of teaching me, questioning me or simply asking me to explain what I am doing or why I am pursuing this degree has borne fruit.

Definitions of Terms and Abbreviations

Adaptive Market Hypothesis (AMH) - An alternative theory of market efficiency, which posits that market efficiency follows a dynamic (cyclical) pattern. The agents in such a market are subject to the principles of behavioural biases, competition, adaptation and natural selection.

Akaike Information Criterion (AIC) - A measure of relative quality of one regression model to another. The criterion measures the trade off between a model's goodness of fit and complexity.

ALSI - All Share Index, given by the share code J203.

ARIMA *(p,d,q)* model - A particular time series model, the Auto Regressive Integrated Moving Average model is used to better understand the existing relationship in a dataset and to forecast that relationship. The parameters p, d and q refer to the order of the autoregressive, integrated and moving average components, respectively.

Artificial Neural Network (ANN) - Computational models that were inspired by the processing capabilities of the human brain. These models are from the field of computer science and are capable of learning and performing pattern recognition (if the pattern is captured in a time series, then network is referred to as a dynamic artificial neural network).

Autocorrelation - The correlation of a series with itself. Autocorrelation is also referred to as serial correlation as the observations in the time series are correlated across time.

Bayes Information Criterion (BIC) - Closely related to the AIC, the Bayes Information Criterion is based on the maximum likelihood function of a model and measures the trade-off between a model's goodness of fit and complexity.

Bayes' Theorem - A theorem derived from the axioms of probability, with emphasis on conditional probability. The theorem is used to describe how a subjective degree of belief should rationally change to account for the evidence observed.

Cointegration - When examining two or more series that are individually integrated, if a linear combination of those series has a lower order of integration, then the series are said to be cointegrated.

Efficient Market Hypothesis (EMH) - A cornerstone of traditional investment theory, the EMH asserts that prices at all times reflect relevant available information. One consequence of this is that no consistent, abnormal profits can be made in financial markets. There are three forms of market efficiency, each a stricter definition of the other, with the strong form being characterised as market prices reflecting all available information.

Falsifiable - A statement is referred to as falsifiable if there exists some observation or argument that proves the statement false. According to philosophy, falsifiability is often the criterion to distinguish the scientific from the unscientific.

Feed-forward neural networks (FFNN) - The simplest type of neural network, a FFNN is characterised by the uni-directional flow of information.

Forecasting - The formal process of employing statistical methods to create observations on a time series that have not yet occurred. Automatic forecasting refers to the selection of the appropriate time series model and the generating of forecasts without human intervention.

Gaussian Random Walk - A Gaussian random walk is one where the successive steps have an underlying Gaussian distribution – an enhancement of the discovery by Regnault (1863).

Hurst exponent - A measure of long term memory in a time series.

Jensen's market efficiency - Jensen's (1978) defines market efficiency as an extension of a zero profit competitive equilibrium condition bridging the gap between the certainty world of classical price theory to the uncertain world of dynamic market behaviour.

JSE - Johannesburg Stock Exchange

Kurtosis - A measure of the peakedness of a probability distribution. If this peakedness is beyond levels of a normal distribution, then the resulting kurtosis is referred to as excess kurtosis.

Learning - The process of training a neural network, learning can either be supervised by the researcher or unsupervised. In supervised learning, the network is given output data and attempts to match it to inputs as accurately as possible. In unsupervised learning, there is no output data given to the network - it attempts to create sample outputs based on minimising some error function.

Levenberg-Marquadt Algorithm - A method to solve non-linear least squares problems that minimises the distance between the error and output data.

Long term memory - Where a single shock will have a noticeable and persistent impact on future volatility.

Ljung Box test - A type of test to determine if a group of autocorrelations from a time series are statistically different from zero.

Market - The earliest definition of a market is provided by Gibson (1889) in that "when shares become publicly known in an open market, the value which they acquire may be regarded as the judgement of the best intelligence concerning them".

Market efficiency - The traditional definition of market efficiency is given by the EMH (discussed above), where the current price reflects all available information such that no abnormal profits can be sustained in the long term. Other forms of market efficiency can include informational efficiency, where information is assimilated instantaneously into the stock price; and allocative efficiency in which capital is allocated in a manner that benefits all market participants.

Martingale - A word with various (related) definitions, the traditional meaning of a martingale refers to a bettering strategy where a gambler doubles his (her) bet after every loss made, resulting in the first win recovering all prior losses as well as the cost of entering the

gamble. In mathematics, a martingale is a stochastic process where the next value in the sequence is equal to the present observed value given all prior observed values.

Moments of the distribution - A quantitative measure of the shape of a set of points, the first four central moments of a distribution often refer to the mean, variance, skewness and kurtosis respectively.

Mean Squared Error (MSE) criterion - The MSE criterion measures the average of the squared error terms. Typically, the better regression is that which has a lower MSE. The square root of the MSE criterion gives rise to the Root MSE (RMSE) criterion.

Non-linear Auto-Regressive with Exogenous (NARX) Neural Networks (NARX NN) - A recurrent neural network, the NARX NN allows for lagged dependent and lagged independent values of each variable to have an influence in explaining the contemporaneous dependent variable.

Neural networks (NN) - Computational models inspired by the processing of the human brain, NNs are capable of matching inputs to outputs under a variety of different learning techniques.

Neurons - A neuron is a biological cell that processes and transmits information through electrical and chemical signals. They form part of the nervous system.

Order of integration - The minimum number of differences required for a time series to be stationary.

Perceptron - An algorithm of supervised learning in which an input variable is transformed into one of several possible non-binary outputs according to a linear classifier.

Random Walk - According to Pearson (1905), a random walk is a mathematical description of a path of successive random steps.

Rational agents - According to Muth (1961), the expectations of agents tend to be distributed for the same information set about the objective probability distribution of

outcomes. Thus, rational agents do not waste information, they form expectations based on the structure of the relevant system describing the economy and public opinion has no substantial effect on an agent's expectation.

Recurrent Neural Networks (RNN) - A class of neural networks where there is a bi-directional flow of information.

Samuelson's dictum - A hypothesis that the EMH should apply more closely at a micro-level than at an aggregated, macro-level.

Skewness - The measure of asymmetry in a probability distribution about its mean. If this value is in excess to that of a normal distribution, it is referred to as excess skewness.

Self-Exciting Threshold Autoregressive Model (SETAR) - A time series model that extends the typical autoregressive model to allow for regime changes.

Student's t-Test - A statistical hypothesis test in which the test statistic follows a Student's t distribution.

Subjective Expected Utility Theory - A category of decision theory that combines the concepts of a personal utility function and a personal probability distribution to make decisions in the presence of risk.

Type I and Type II error - A Type I error is when the null hypothesis is incorrectly rejected; whereas a Type II error is when the null hypothesis should be rejected but is not rejected.

White noise - A term referring to randomly generated errors (noise).

Wiener process - A continuous time stochastic process that is used to describe the random behaviour of share prices.

The dynamics of market efficiency:
Testing the Adaptive Market Hypothesis in South Africa

ABSTRACT

In recent years, the debate on market efficiency has shifted to providing alternate forms of the hypothesis, some of which are testable and can be proven false. This thesis examines one such alternative, the Adaptive Market Hypothesis (AMH), with a focus on providing a framework for testing the dynamic (cyclical) notion of market efficiency using South African equity data (44 shares and six indices) over the period 1997 to 2014. By application of this framework, stylised facts emerged. First, the examination of market efficiency is dependent on the frequency of data. If one were to only use a single frequency of data, one might obtain conflicting conclusions. Second, by binning data into smaller sub-samples, one can obtain a pattern of whether the equity market is efficient or not. In other words, one might get a conclusion of, say, randomess, over the entire sample period of daily data, but there may be pockets of non-randomness with the daily data. Third, by running a variety of tests, one provides robustness to the results. This is a somewhat debateable issue as one could either run a variety of tests (each being an improvement over the other) or argue the theoretical merits of each test befoe selecting the more appropriate one. Fourth, analysis according to industries also adds to the result of efficiency, if markets have high concentration sectors (such as the JSE), one might be tempted to conclude that the entire JSE exhibits, say, randomness, where it could be driven by the resources sector as opposed to any other sector. Last, the use of neural networks as approximators is of benefit when examining data with less than ideal sample sizes. Examining five frequencies of data, 86% of the shares and indices exhibited a random walk under daily data, 78% under weekly data, 56% under monthly data, 22% under quarterly data and 24% under semi-annual data. The results over the entire sample period and non-overlapping sub-samples showed that this model's accuracy varied over time. Coupled with the results of the trading strategies, one can conclude that the nature of market efficiency in South Africa can be seen as time dependent, in line with the implication of the AMH.

Keywords: Market efficiency; neural networks; SETAR models; emerging markets

JEL Classification: C45, C58, G02, G14

Table of Contents

List of Figures ..... **xiii**

List of Tables ..... **xv**

1 Introduction **1**

1.1 From statistics to machine learning 3

1.2 Describing (and modelling) a financial market 6

1.3 Feasibility of study 7

1.4 Hypothesis and theoretical framework 7

1.4.1 Problem statement 7

1.4.2 Primary hypothesis 8

1.4.3 Secondary hypotheses 8

1.5 Chapter Outline 9

2 Literature Review **10**

2.1 A qualitative view of market efficiency 10

2.1.1 The history of market efficiency 12

2.1.2 Defining the Efficient Market Hypothesis 16

2.1.3 An axiomatic approach to informationally efficient markets 18

2.2 A market participant's view on efficiency 20

2.2.1 Tests of the speed of adjustment of prices to new information 22

2.2.2 Discrepancies between price and fundamental value 22

2.3 A statistical view of market efficiency 24

2.3.1 Models of market efficiency 27

2.3.2 Weak form tests 29

2.3.3 Assumptions underlying share returns 32

2.3.4 The behaviour of share returns 36

2.4 Modelling share prices 38

2.5 Time Series Methods 39

2.5.1 Exponential Smoothing 40

2.5.2 Prediction intervals 41

2.5.3 ARIMA models 41

2.5.4 Univariate models 42

2.5.5 Non-linear models 43

2.6 Considerations in asset pricing 47

2.6.2 International asset risk 55
2.6.3 Variable selection in the APT 56
2.7 Considerations in testing market efficiency 61
2.7.1 Risk aversion, uncertainty and market efficiency 63
2.7.2 The Market Fraction Hypothesis 68
2.7.3 Rationality and Market Efficiency 70
2.7.4 The response from the rational investor 72
2.7.5 The Adaptive Market Hypothesis 73
2.7.6 The behavioural view of market efficiency 78
2.8 Frontiers in finance 86
2.8.1 Bounded Rationality 88
2.8.2 Emotion 92
2.8.3 Neuro-finance 93
2.9 Summary 94
3 Data and Methodology 97
3.1 Data 97
3.2 Methodology 101
3.2.1 Testing for normality 101
3.2.2 Testing for linearity 103
3.2.3 Testing for stationarity 104
3.2.4 Testing for random walk behaviour 105
3.2.5 Modelling the return generating process 114
3.2.6 Building an ANN 116
3.3 Neural network hierarchy 118
3.3.1 Non-Linear Autoregressive Models with Exogenous inputs (NARX) 118
3.4 Issues in ANN development 120
3.4.1 Database size and partitioning 121
3.4.2 Data pre-processing, balancing and enrichment 121
3.4.3 Data normalisation 121
3.4.4 Input and Output representation 122
3.4.5 Network weight initialisation 122
3.4.6 The backpropagation learning rate, η 122
3.4.7 The backpropagation momentum coefficient, μ 123

3.4.8 The activation function, σ 123
3.4.9 Convergence criteria 124
3.4.10 Number of training cycles 124
3.4.11 Modes of training 125
3.4.12 Size of the hidden layer 125
3.5 Learning algorithms 126
3.6 Summary 127
4 Results 129
4.1 Preliminary statistics 129
4.1.1 Full sample results 129
4.1.2 ALSI Sub-sample results 132
4.2 Tests for normality 133
4.2.1 Full sample results 133
4.2.2 ALSI Sub-sample results 145
4.3 Tests for non-linearity 148
4.3.1 Full sample results 149
4.3.2 ALSI Sub-sample results 162
4.4 Tests for stationarity 166
4.4.1 Full sample results 166
4.4.2 ALSI Sub-sample results 173
4.5 Testing the random walk hypothesis 177
4.5.1 Full sample results 178
4.5.2 ALSI Sub-sample results 200
4.5.3 Summary of results for the test of random walk behaviour 208
4.6 Modelling the data generating process without additional factors 210
4.6.1 ARIMA models 210
4.6.2 SETAR Test of linearity 212
4.6.3 SETAR model 215
4.6.4 Neural network modelling results 221
4.6.5 ALSI Sub-sample results 230
4.6.6 Summary of results for modelling returns without additional factors 233
4.7 Modelling the data generating process with additional factors 234
4.7.1 Daily return sample results 235

4.7.2 Weekly return sample results 238
4.7.3 Monthly return sample results 240
4.7.4 Quarterly return sample results 243
4.7.5 Semi-annual return sample results 245
4.7.6 ALSI Sub-sample results 249
4.7.7 Summary of results for modelling returns using additional factors 251
4.8 A practical test of market efficiency 252
4.9 Overall summary of results 255
4.10 A combined view of market efficiency 279
5 Conclusion 281
References 289
Appendix A 305
The BDS Test 305
Appendix B 307
Plagiarism Report 357

List of Figures

Figure 1 – Share price changes in response to favourable and unfavourable information (Lewellen and Shanken, 2002)....66
Figure 2 – The phases of ANN development....117
Figure 3 – A Non-linear Autoregressive with Exogenous Inputs Network....119
Figure 4 – A Multi-layer perceptron....120
Figure 5 - Q-Q plot of daily frequency data (1)....136
Figure 6 – Q-Q plot of daily frequency data (2)....137
Figure 7 - Q-Q plot of weekly frequency data (1)....138
Figure 8 - Q-Q plot of weekly frequency data (2)....139
Figure 9 - Q-Q plot of monthly frequency data (1)....140
Figure 10 - Q-Q plot of monthly frequency data (2)....141
Figure 11 - Q-Q plots for each exogenous variable to be used in the network (1)....141
Figure 12 - Q-Q plots for each exogenous variable to be used in the network (2)....142
Figure 13 - Q-Q plots for each sub-sample (1)....146
Figure 14 - Q-Q plots for each sub-sample (2)....147
Figure 15 - Rolling autocorrelation at lag 1 for daily (top right), weekly (top left) and monthly (centre) ALSI returns....172
Figure 16 - Rolling autocorrelation for each sub-sample (1)....175
Figure 17 - Rolling autocorrelation for each sub-sample (2)....176
Figure 18 - Rolling autocorrelation for each sub-sample (3)....177
Figure 19 - Graphical representation of variance ratio for daily returns (1)....186
Figure 20 - Graphical representation of variance ratio for daily returns (2)....187
Figure 21 - Graphical representation of variance ratios for weekly returns (1)....188
Figure 22 - Graphical representation of variance ratios for weekly returns (2)....189
Figure 23 - Graphical representation of variance ratios for monthly returns (1)....190
Figure 24 - Graphical representation of variance ratios for monthly returns (2)....191
Figure 25 - Hurst Exponent (Peng method) for daily returns (1)....193
Figure 26 - Hurst Exponent (Peng method) for daily returns (2)....194
Figure 27 - Hurst Exponent (Whittle Estimator) for weekly returns (1)....195
Figure 28 - Hurst Exponent (Whittle Estimator) for weekly returns (2)....196
Figure 29 - Hurst Exponent (R/S method) for monthly returns (1)....198
Figure 30 - Hurst Exponent (R/S method) for monthly returns (2)....199
Figure 31 - Graphical representation of variance ratios of each sub-sample (1)....203
Figure 32 - Graphical representation of variance ratios of each sub-sample (2)....204
Figure 33 - Hurst exponents for the first set of four sub-samples....206
Figure 34 - Hurst exponents for the second set of four sub-samples....207
Figure 35 - Hurst exponents for the last set of sub-samples....208
Figure 36 - Forecasts of daily returns....211
Figure 37 - Performance of the NAR network using daily returns....222
Figure 38 - R^2 of the NAR network using daily returns....223
Figure 39 - Autocorrelation of the error terms of the NAR network using daily returns....223
Figure 40 - Performance of the NAR network using weekly returns....225

Figure 41 - R^2 of the NAR network using weekly returns 225
Figure 42 - Autocorrelation of the error terms of the NAR network using weekly returns 226
Figure 43 - Performance of the NAR network using monthly returns 227
Figure 44 - R^2 of the NAR network using monthly returns 228
Figure 45 - Autocorrelation of the error terms of the NAR network using monthly returns 228
Figure 46 – Representation of the NARX network in Matlab 235
Figure 47 - Depiction of MSE over time, using daily ALSI data 236
Figure 48 - Comparison of R^2 during training, validation and testing phases 237
Figure 49 - Autocorrelation function over time 237
Figure 50 - Depiction of MSE over time, using weekly ALSI data 239
Figure 51 - Comparison of R^2 during training, validation and testing phases 239
Figure 52 - Autocorrelation function over time 240
Figure 53 - Depiction of MSE over time, using monthly ALSI data 241
Figure 54 - Comparison of R^2 during training, validation and testing phases 242
Figure 55 - Autocorrelation function over time 242
Figure 56 - Depiction of MSE over time, using quarterly ALSI data 243
Figure 57 - Comparison of R^2 during training, validation and testing phases 244
Figure 58 - Autocorrelation function over time 245
Figure 59 - Depiction of MSE over time, using semi-annual ALSI data 246
Figure 60 - Comparison of R^2 during training, validation and testing phases 247
Figure 61 - Autocorrelation function over time 248
Figure 62 - Comparison of hidden and delay nodes for each sub-sample 250
Figure 63 - Comparison of training and overall R^2 for each sub-sample 250

List of Tables

Table 1 - Shares and indices used 98
Table 2 – Variables considered 100
Table 3 - Hurst exponent confidence intervals for the Peng estimate 112
Table 4 - Hurst exponent confidence intervals for the Whittle estimate 113
Table 5 - Hurst exponent confidence intervals 113
Table 6 – Summary of network parameters 125
Table 7 – Descriptive statistics 130
Table 8 – Descriptive statistics of the inputs to the network 132
Table 9 - Descriptive statistics of each sub-sample 133
Table 10 – Jarque Bera Test for normality using daily frequency data 134
Table 11 - Jarque Bera Test for normality using weekly frequency data 134
Table 12 - Jarque Bera Test for normality using monthly frequency data 135
Table 13 – Jarque Bera test for normality on the input variables 135
Table 14 - K-S results for daily returns 143
Table 15 - K-S results for weekly returns 143
Table 16 - K-S results for monthly returns 144
Table 17 - K-S results for exogenous variables 144
Table 18 - Jarque Bera results for each sub-sample 145
Table 19 - K-S results for each sub-sample 148
Table 20 – BDS test for non-linearity of BIL daily returns 149
Table 21 – BDS test for non-linearity of MTN daily returns 149
Table 22 – BDS test for non-linearity of SOL daily returns 150
Table 23 – BDS test for non-linearity of FSR daily returns 150
Table 24 – BDS test for non-linearity of SAB daily returns 150
Table 25 – BDS test for non-linearity of NPN daily returns 151
Table 26 – BDS test for non-linearity of AGL daily returns 151
Table 27 – BDS test for non-linearity of J200 daily returns 151
Table 28 – BDS test for non-linearity of GFI daily returns 152
Table 29 – BDS test for non-linearity of ALSI daily returns 152
Table 30 – BDS test for non-linearity of BIL weekly returns 152
Table 31 – BDS test for non-linearity of MTN weekly returns 153
Table 32 – BDS test for non-linearity of SOL weekly returns 153
Table 33 – BDS test for non-linearity of FSR weekly returns 153
Table 34 – BDS test for non-linearity of MDC weekly returns 154
Table 35 – BDS test for non-linearity of PIK weekly returns 154
Table 36 – BDS test for non-linearity of HYP weekly returns 154
Table 37 – BDS test for non-linearity of ILV weekly returns 155
Table 38 – BDS test for non-linearity of WBO weekly returns 155
Table 39 – BDS test for non-linearity of ALSI weekly returns 155
Table 40 – BDS test for non-linearity of BIL monthly returns 156
Table 41 – BDS test for non-linearity of MTN monthly returns 156
Table 42 – BDS test for non-linearity of SOL monthly returns 156

Table 43 – BDS test for non-linearity of FSR monthly returns 157
Table 44 – BDS test for non-linearity of SAB monthly returns 157
Table 45 – BDS test for non-linearity of NPN monthly returns 158
Table 46 – BDS test for non-linearity of AGL monthly returns 158
Table 47 – BDS test for non-linearity of J200 monthly returns 158
Table 48 – BDS test for non-linearity of GFI monthly returns 159
Table 49 – BDS test for non-linearity of ALSI monthly returns 159
Table 50 – BDS test for non-linearity on Oil 159
Table 51 - BDS test for non-linearity on Gold 160
Table 52 - BDS Test results for the ALSI Dividend Yield 160
Table 53 - BDS Test results for the ALSI Earnings Yield 160
Table 54 - BDS Test results for S&P 500 161
Table 55 - BDS Test results for the Hang Seng 161
Table 56 - BDS Test results for the FTSE 161
Table 57 - BDS Test results for the sub-sample 1 162
Table 58 - BDS Test results for the sub-sample 2 162
Table 59 - BDS Test results for the sub-sample 3 163
Table 60 - BDS Test results for the sub-sample 4 163
Table 61 - BDS Test results for the sub-sample 5 164
Table 62 - BDS Test results for the sub-sample 6 164
Table 63 - BDS Test results for the sub-sample 7 164
Table 64 - BDS Test results for the sub-sample 8 165
Table 65 - BDS Test results for the sub-sample 9 165
Table 66 - BDS Test results for the sub-sample 10 165
Table 67 - ADF Test Results for daily returns 167
Table 68 - ADF Test Results for weekly returns 167
Table 69 - ADF Test Results for monthly returns 168
Table 70 - ADF Test Results for exogenous variables 168
Table 71 - KPSS results for daily returns 169
Table 72 - KPSS results for weekly returns 169
Table 73 - KPSS results for monthly returns 170
Table 74 - KPSS results for exogenous variables 170
Table 75 - ADF results for each sub-sample 173
Table 76 - KPSS results for each sub-sample 174
Table 77 - Runs test for daily returns 178
Table 78 - Runs test for weekly returns 179
Table 79 - Runs test for monthly returns 179
Table 80 – Wright test on daily returns 180
Table 81 - Wright test on weekly returns 181
Table 82 - Wright test on monthly returns 182
Table 83 – Chow Denning test on daily returns 184
Table 84 - Chow Denning test on weekly returns 184
Table 85 - Chow Denning test on monthly returns 185
Table 86 - Runs test result for each sub-sample 201

Table 87 - Wright results for each sub-sample 201
Table 88 - Chow Denning results for each sub-sample 202
Table 89 - ARIMA for daily returns 210
Table 90 - ARIMA models for weekly returns 211
Table 91 - ARIMA models for monthly returns 212
Table 92 - SETAR test for daily returns 213
Table 93 - SETAR test for weekly returns 214
Table 94 - SETAR test on monthly returns 215
Table 95 - SETAR model for daily ALSI returns 216
Table 96 - SETAR model diagnostics for selected shares using daily returns 216
Table 97 - SETAR model for weekly ALSI returns 217
Table 98 - SETAR model diagnostics for selected shares using weekly returns 218
Table 99 - SETAR model for monthly ALSI returns 218
Table 100 - SETAR model diagnostics for selected shares using monthly returns 219
Table 101 - SETAR model for quarterly ALSI returns 219
Table 102 - SETAR model diagnostics for selected shares using quarterly returns 220
Table 103 - SETAR model for semi-annual ALSI returns 220
Table 104 - SETAR model diagnostics for selected shares using semi-annual returns 221
Table 105 - NAR model diagnostics for selected shares using daily returns 224
Table 106 - NAR model diagnostics for selected shares using weekly returns 226
Table 107 - NAR model diagnostics for selected shares using monthly returns 229
Table 108 - ARIMA model for each sub-sample 230
Table 109 - SETAR test on each sub-sample 231
Table 110 - SETAR model diagnostics for each sub-sample 232
Table 111 - NAR model diagnostics for each sub-sample 233
Table 112 – Overview of methods used in training the daily NARX network 235
Table 113 – Overview of methods used in training the weekly NARX network 238
Table 114 – Overview of methods used in training the monthly NARX network 240
Table 115 – Overview of methods used in training the quarterly NARX network 243
Table 116 – Overview of methods used in training the semi-annual NARX network 245
Table 117 – Split sample results of the NARX network 249
Table 118 - Comparison of returns across active and passive strategies (1) 254
Table 119 - Comparison of returns across active and passive strategies (2) 255
Table 120 – BDS test results for all shares (1) 257
Table 121– BDS test results for all shares (2) 258
Table 122– BDS test results for all shares (3) 259
Table 123 – Runs test results for all shares (1) 260
Table 124 – Runs test results for all shares (2) 261
Table 125 – Chow Denning test results for all shares (1) 262
Table 126 – Chow Denning test results for all shares (2) 263
Table 127 – Chow Denning test results for all shares (3) 264
Table 128 – Wright test results for all shares (1) 264
Table 129 – Wright test results for all shares (2) 265
Table 130 – Wright test results for all shares (3) 266

Table 131 – Variance Decomposition test results for all shares (1) 267
Table 132 – Variance Decomposition test results for all shares (2) 268
Table 133– Variance Decomposition test results for all shares (3) 269
Table 134 – Hurst exponent results for all shares (1) 270
Table 135– Hurst exponent results for all shares (2) 271
Table 136 - Summary of results for the random walk hypothesis 274
Table 137 - Summary of results for each sub-sample for the random walk hypothesis (1) ..276
Table 138 - Summary of results for each sub-sample for the random walk hypothesis (2) ..277
Table 139 - Statistical and trading results (1) 279
Table 140 - Statistical and trading results (2) 280

1 Introduction

There is a longstanding debate amongst academics and practitioners on the question of market efficiency. Since the seminal work of Fama (1970) setting out the Efficient Market Hypothesis (EMH), empirical tests have been conducted to determine whether markets are efficient or not. In recent years, the *status quo* has shifted towards providing alternate hypotheses of market efficiency which are testable and can thus be proven to be true or false. This thesis examines one such alternative, the Adaptive Market Hypothesis (AMH) of Lo (2004, 2005) with a focus on providing a framework for testing this different form of market efficiency. Given the development of the AMH, no formal means of testing cyclical efficiency has been established in the literature. Therefore, this thesis offers one such set of ideas to testing cyclical efficiency. In particular, it will be determined whether share returns are deterministic or not, with further extensions to modelling the returns generating process. By examining both individual shares as well as indices (a total of 50 traded assets), as well as across different frequencies of returns, a holistic view of market efficiency can be obtained. If the South African equities market does indeed exhibit cyclical efficiency, this result would be found based on the behaviour (deterministic or not) of share returns and the ability to model said returns over time. A range of tests (including a practical trading application) are conducted on equity returns in South Africa to comprehensively determine if these returns follow a random walk or not. Further, both traditional econometric methods as well as artificial intelligence models are used to determine if the returns generating process can be specified. By the application of a variety of methods, to permutations of the data frequency, this thesis attempts to comprehensively examine market efficiency on the South African equities market. It is conjectured that if the success of this modelling procedure varies over time, then the equities market can be seen as adaptively efficient, in line with the AMH.

The AMH emerged from principles in evolutionary biology, psychology and sociology (Lo, 2004). This Adaptive Market Hypothesis would describe efficiency as the interaction of market participants. In a market with supply- and demand- side participants, the interaction between these two groups determines an equilibrium price. Lo (2004) argues that a market with scarce resources would be more efficient than one with abundant resources. As these two groups are driven by an instinct to survive in the market place, the individuals can learn and make informed decisions on whether to purchase or sell the good in question. Over time,

innovation in the market place can "reset" the current market dynamics, leading to individuals finding a new means of survival. This process can be seen as an evolutionary one, where an innovation is seen to disrupt the current equilibrium; and the learning experience of individuals lead to the formation (over time) of new groups of individuals with similar characteristics leading to a new equilibrium. Once a new equilibrium has been reached, this process is seen as efficient. Hence, efficiency would be seen as cyclical, limited by the nature of said participants and the environment this interaction occurs within. The implications of the AMH suggest that efficiency can be viewed to be a relative measure - there would be times when the market is efficient and times when the market is inefficient. While Lo (2004, 2005) both describes the abstract and practical implications of the AMH, little attention is given to testing the implications thereof. This thesis aims to provide a practical means of testing the core implication of the AMH – that of cyclical market efficiency.

Borrowing from the discipline of computer science, the concept and application of neural networks is used to model the efficiency of the market. A neural network, in its simplest form, can be represented by a set of processes or "nodes" (some of which can be unknown or “hidden”), that would convert an input to the desired output. In contrast, a more traditional time series model would be specified in advance and then applied to the financial problem at hand. Thus, the application of neural networks to solving financial problems can be seen as an extension of an econometrics method, with the difference being that the network's processes do not necessarily have to be specified in advance. The foundation of neural networks rests in scientists’ attempts to map the processing capability of the human brain. Specifying a neural network requires the selection of input data, selecting the appropriate network architecture, training the network based on a particular algorithm, and measuring performance of the network. To allow comparison between traditional econometric methods, the performance of neural networks can be evaluated by examining the error term. Garth, Rollins, Zhu and Chen (1996) show that network performance rests on two variables - the number of hidden layers and thus nodes in the network and the standard error. The optimum network would be the network that balances the errors generated with the number of hidden nodes. Thus, if the error term begins to increase after a particular point, then the network exhibits decreasing performance at its task of learning and forecasting the data series. The basic idea of participants learning as the network develops intrinsically harmonises with that

of an adaptive market, suggesting intuitively that neural networks can provide a method of testing the AMH.

While much related literature exists on the application of neural networks to finance, no published works link the application of neural networks to the AMH. Literature discussed in Chapter 2, shows, for example, the application of a neural network to forecast an additional term in a data series. Using financial data, this forecasting ability can be turned into a portfolio strategy, where the investor will buy if the share is underpriced and sell if the share is overpriced as determined by the difference between the network's output and the actual share price, much like an investor would use asset pricing models to determine if a share is under- or over-priced. This example shows that whilst the foray into neural networks is somewhat distanced from finance, these tools can be used to solve problems in the field of finance. This thesis adopts this cross-disciplinary approach in attempting to show that market efficiency can be considered over time as opposed to at a single point in time. In other words, if one uses the entire sample data over a particular time period, the conclusion reached regarding market efficiency can differ if the sample period was divided into smaller intervals (or differing frequencies). This series is conceptually described by the AMH.

1.1 From statistics to machine learning

Conceptually, there are both similarities and differences between a statistical regression and a neural network. Consider a statistical regression represented as $\hat{Y} = f(A, X)$. It provides an estimate of a dependent variable, $\hat{Y}$, a function of a vector of independent variables (X) and their associated regression coefficients (A) according to some function f. The regression technique rests in minimising the error term of the regression as well as specifying the function and independent variables *a priori*. While a regression aims to minimise the difference between the actual and predicted values of the dependent variable, there is no assurance that a particular regression for a particular problem statement and dataset is the optimal one. Comparatively, a neural network can be represented as $\hat{Y} = g(W, X)$. The dependent variable, $\hat{Y}$, is estimated by a set of independent variables (X) and their connection

weights (W) according to some complex[1] function g. A neural network is trained to produce the optimal output. It can be likened to implementing a stepwise regression in that the process is repeated until the optimal parameters (referred to as connection weights) are obtained. In a multiple linear regression, f represents the set of linear operators whereas in a neural network, g represents a linear combination of a number of non-linear functions. A particular form of network referred to as a feed-forward neural network with no hidden layers can be viewed as a generalisation to a statistical model. In other words, the input data is passed to the complex function and then transformed to an output - there are no additional computations (hidden layers) between the complex function and the output. The error term (the difference between the actual and observed output) is not passed back to the complex function - the data is *fed forward* only. In a statistical model, input data is fed into a (complex) function to produce an output, thus it is similar to the network described above. The process of stepwise regression is analogous to the learning algorithm in a neural network. Many such learning algorithms exist, some of which, such as Hebbian learning, are closely related to statistical modelling (Hebb, 1949). Further, the design of the complex function is referred to as the network architecture. As in econometrics, there is a multitude of "model architectures" to choose from.

The similarities between neural networks and the more traditional regression model are quite striking, yet the application of the former has not yet been fully incorporated as a mainstream approach to solving financial problems. Wythoff (1993) classifies a neural network as a generalisation of a classical regression, where the Artificial Neural Network (ANN) is trained adaptively using non-linear learning laws (such as a sigmoid activation function) compared to matrix inversion in regression modelling. Indeed, the network is trained to produce the lowest error term, so it is unrestricted in choosing whether to provide a linear or non-linear function to map inputs to outputs. Hanson (1995) defines a back-propagation network as a “multi-variate, non-linear, non-parametric, stochastic approximation with dynamic feature extraction and selection, which makes them capable of learning arbitrary mapping”. In other words, in a back-propagation network (a network where the data, including the error term, flows both

[1] A complex function can either be defined as a non-real function or where the complexity of the function is computationally expensive. In other words, the calculation of the function uses a large amount of resources and is time consuming.

forwards and *backwards*), the function or independent variables do not need to be pre-specified as in the case with regression modelling.

From a practical perspective, the choice between two equally appropriate modelling techniques is primarily made by examining the costs of running each model. As the dimensionality and non-linearity of the problem increases, the neural network becomes superior to a regression in producing accurate approximations. In a regression, an increase in the number of independent variables, N, increases the number of polynomial parameters by N^m, where m is the order of the polynomial. However, in a neural network, the number of parameters grows either linearly or quadratically (N^2) for a given m hidden layers. Thus, neural networks are considered (in some cases) to be less computationally expensive than traditional models. Basheer and Hajmeer (2000) recommend that when following an ANN approach, the researcher considers weighing the costs of a higher accuracy, more complex ANN to the increase in development time and lost characteristics of a statistical model. While the network may produce more accurate results (as given by a lower error term), the lack of interpretation and increase in production time may not be feasible for the research question at hand.

According to Ruck, Rogers, Kabrisky, Oxley and Suter (1990), typical neural network models estimate Bayesian *a posteriori* probabilities when given an appropriately defined problem. As such, when given noisy financial data, a delay embedding of previous inputs is usually suggested. However, Giles, Lawrence and Tsoi (2001) outline two reasons that make prediction difficult for noisy, non-stationary time series data. First, as the network will learn from examples, there will exist infinitely many models that can work as well or better by learning from the same example. It is thus desirable to have a larger training set to enable better generalisation of results. Yet, as the training set size increases, the chance of non-stationarity also increases. Second, small datasets that contain much noise makes the ANN prone to overfitting. Typical ANNs will thus often overlook the temporal relationships between the input variables and the output variable. Thus, it is suggested to use a form of Recurrent Neural Network (RNN) to maintain the temporal relationship between variables and to represent certain computational structures in a more parsimonious manner (Elman, 1991). In particular, this thesis utilises a Non-linear Autoregressive with Exogenous inputs (NARX) neural network - in effect a time series model with lagged dependent and

independent variables. The selection of the NARX network is motivated by both empirical studies as well as tests of alternatives conducted in this thesis.

1.2 Describing (and modelling) a financial market

A financial market can be regarded as a complex, dynamic system. The interaction between the forces of supply and demand leads to an ever-changing environment, given by the change in asset prices per tick; whereas the arrival and assimilation of information coupled with the varied reaction of market participants to information lends credence to a market's complex nature. Whether in the search for additional profits by the trader, or the representation of this system for academic exploration, it remains a question of whether the complexities of the financial market can indeed be captured by a neural network. Takens' Theorem (Takens, 1981), states that any chaotic dynamic system can be modelled from a sequence of observations of the state of a dynamical system. At face value, this should imply that the dynamics of the market could be represented by a neural network. However, while one can easily argue that the financial market is dynamic, it is difficult to argue that the financial market is chaotic.

A chaotic system is defined as one where a minor change in the initial conditions can give rise to significantly different outcomes. At the moment of new information arriving and being assimilated by market participants, it can be inferred that the initial condition of, say, a particular share, has been altered. This would culminate in a reaction that is *prima facie* unexpected, implying an element of randomness when it is in fact simply chaotic. If one can incorporate this new information in a model, one can provide output solutions that are in line with the actual target output. The biological link between the neural network model and the processing capability of the human brain leads to an interesting - albeit tenuous - perspective regarding market efficiency and profitability. If the market participant (here represented by a neural network) views share prices as non-random, he would be motivated by the need to earn additional profits. Further, if he were to predict the next observation (share price), using existing and new information, he would also view the market as being inefficient as described by the EMH. Given the tendency of investors to be overconfident (Barber and Odean, 2001) in their investment decisions, they would perpetually apply the same portfolio strategy based on their own predictions of the share price. If other investors follow the same behaviour, then the excess profits earned per investor would diminish to the point of being eliminated. Once eliminated, most investors (exhibiting arguably rational behaviour) would modify their strategy and

predictions whereas some investors (exhibiting arguably not fully rational behaviour) will continue to apply the same strategy and methodology in predicting a share's future price. Thus, the once eliminated profits now begin to re-emerge and a cycle is formed. The cycles of profitability and efficiency are therefore offered as the characteristics of an adaptively efficient market, described by Lo (2004, 2005).

1.3 Feasibility of study

A study of this nature has not been published in a South African context, primarily as the present literature on testing the EMH focuses on tests of randomness in share returns. The literature groups these tests into predictability of security returns and profitability of trading strategies. Under the first grouping, the tests are further subdivided into constant or time-varying parameters with permutations to the full sample, non-overlapping samples and overlapping samples. While employing these tests on the JSE is far from unique, they nonetheless provide the foundation to examine the returns generating process. The returns process will be attempted to be modelled by a regime changing time series model without exogenous factors as well as models from artificial intelligence (some of which make provision for exogenous factors). Here, exogenous factors are informed by the application of Arbitrage Pricing Theory (APT) and can range from macroeconomic, microeconomic or even behavioural (sentiment) based. Further, while artificial intelligence can be considered an enhancement to traditional econometric models, various data considerations, discussed later, need to be considered to ensure that issues such as structural breaks do not influence the results negatively. Thus, the contribution of this thesis is twofold - by combining elements from other disciplines, such as computer science and biology, approaches to solve or explain financial problems using inter-disciplinary approaches are enhanced; and an emerging framework for testing adaptive market efficiency is developed.

1.4 Hypothesis and theoretical framework

1.4.1 Problem statement

Within the context of the South African equity market, is market efficiency, described by the AMH, indeed cyclical?

1.4.2 Primary hypothesis

H_0: Market efficiency is not cyclical.
H_1: Market efficiency is cyclical.

1.4.3 Secondary hypotheses

$H_{0,A}$: Share price behaviour, in the South African market, does not follow a random walk.
$H_{1,A}$: Share price behaviour, in the South African market, does follow a random walk.

$H_{0,B}$: Share price behaviour, in the South African market, cannot be modelled by an autoregressive function *with no* exogenous inputs.
$H_{1,B}$: Share price behaviour, in the South African market, can be modelled by an autoregressive function *with no* exogenous inputs.

$H_{0,C}$: Share price behaviour, in the South African market, cannot be modelled by an autoregressive function *with* exogenous inputs.
$H_{1,C}$: Share price behaviour, in the South African market, can be modelled by an autoregressive function *with* exogenous inputs.

This thesis aims to provide a practical means of testing the AMH over the period of 1997 to 2014. Whilst the objective is not to prove nor disprove the EMH, an indirect comparison should be expected to emerge from the results. The hypothesis of cyclical efficiency will be tested through three phases. Firstly, it is necessary to establish whether share price changes follow a random walk or not. If price changes are random, they cannot be predicted, thus enforcing the notion of weak form market efficiency. However, if price changes are not random, secondly, it is then viable to establish whether they can be modelled. In the simplest case, one can model current share prices based on prior values. If this model is found to be inadequate, then lastly, one can model share prices based on both prior values and exogenous factors.

It is important to note that this thesis does not test the EMH *per se*, as the EMH is not considered to be a falsifiable theory. In other words, one cannot reject or fail to reject the EMH as a test of the EMH requires a pre-specified model of price determination. If one

rejects the EMH, it is not possible to determine whether the pricing model is rejected or whether the hypothesis itself is rejected. This is referred to as the joint hypothesis problem (Campbell, Lo, MacKinlay and Whitelaw, 1998) and is a consequence of the axiomatic (first principles) approach to the definition of an informationally efficient market (a market where any new anticipated information is already priced into the market and thus has no impact on price movements). This approach does not provide testable criteria on what an efficient market is nor its counterpart. Thus, it can be said that there has not been any proof against the EMH, as the EMH cannot be refuted. The tests of the EMH are thus simply descriptions of statistical facts about the behaviour of financial markets.

This thesis provides an analysis of the behaviour of the South African equities market. In analysing different frequency data as well as trying to explain past price behaviour, the aim of the thesis is to show that, historically, market efficiency on the JSE has exhibited cyclicality, as defined by the level of which a model can determine prices (or equivalently, returns). While popular tests have emerged in determining whether markets are either weak, semi-strong or strong form efficient (these are discussed in Chapter 2), there is no guideline as to how one determines market efficiency. Similarly, when the AMH was proposed by Lo (2004, 2005), no framework as provided on how cyclical efficiency can be evaluated or tested. Therefore, this thesis offers a possible framework for evaluating efficiency, whether cyclical or not, in a financial market (here restricted to the South African equities market).

1.5 Chapter Outline

The following chapters will be presented in this thesis. Chapter Two provides a literature review for this study, beginning with an examination of market efficiency from different perspectives, continuing with an overview of time series econometrics and ending with developments in the field of finance. Chapter Three details the data and methodology of the thesis. An exposition of the dataset, tests of market efficiency and models used to determine the returns generating process are presented. Chapter Four provides the results obtained from the study and a discussion thereof, while Chapter Five provides concluding remarks on the thesis.

2 Literature Review

"...I think we can suspect that there is no a priori necessity for actual Board of Trade grain prices to act in accordance with specific probability models. Perhaps it is a lucky accident, a boon from Mother Nature so to speak, that so many actual price time series do behave like uncorrelated or quasi-random walk," (Samuelson, 1965, p.42)

This chapter begins with an understanding of the concept of market efficiency, its history, theoretical and practical implications before continuing to explore the area of time series econometrics. The latter is of interest as it provides a foundation and context to examine linear and non-linear time series models. Further, a framework for identifying potential risk factors is covered, ending with some developments and esoteric areas of finance literature, as it applies to market efficiency.

2.1 A qualitative view of market efficiency

The behaviour of share prices has been a long standing enigma for academics in finance. The seminal work of Fama (1970) in defining the Efficient Market Hypothesis (EMH) and in particular, the weak form of the EMH has attracted much attention in the literature as it is perhaps the most intuitive and acceptable of the three forms to comprehend and test.

"If, in January, 1926, an individual invested $1 in one-month U.S. Treasury bills—one of the safest securities in the world—and continued reinvesting the proceeds month by month until December, 1996, the original investment would have grown to $14. If, on the other hand, an individual invested $1 in the S&P 500—a much riskier investment—over the same 71-year period, this investment would have grown to $1,370, a considerably larger sum. Now suppose that, each month, an individual were able to divine [*sic*] in advance which of these two investments would yield a higher return for that month and took advantage of this information by switching the running total of his initial $1 investment into the higher-yielding asset. What would a $1 investment in such a ''perfect foresight'' investment strategy become by December 1996? The startling answer is $2,296,183,456, more than two billion dollars!" (Farmer and Lo, 1999, p1.)

If prediction of financial markets were possible, the individuals in question would have reaped large rewards in return. It is, however, often justified to question if financial markets can be predicted at all, as the power of prediction may not always be fully accurate over time. The success of a particular trading rule over a particular time period may not always hold when adapted to other time periods. Thus, the profitability of following technical analysis waxes and wanes as time progresses, leading the trader to switch rules to ensure profitability is not diminished. Indeed, one of the implications of the AMH is that the profitability of following a particular investment strategy produces cyclical profits or losses over time. Similarly, the second implication of the AMH allows one to view market efficiency as changing over time.

Beginning with the EMH, the title of the work of Samuelson (1965) clearly indicates the author's position towards efficiency – "Proof that properly anticipated prices fluctuate randomly". In an informationally efficient market, price changes must be random since information changes randomly even if information is properly anticipated. Upon closer inspection, this assumption provides a contradictory view of efficiency – the more efficient the market, the more random the price changes. This implies that there are degrees of randomness, in which a fully random process may not be predicted, but a partially random process may be predicted. This statement lends credibility to the idea of examining market efficiency as a relative concept rather than an absolute one. The market can be efficient or inefficient at any point in time, rather than efficient or inefficient across the examined sample period. This stance is adopted in this thesis and attempts are made to provide a framework to examine market efficiency across time.

Further, the outcome of market participants trying to profit from available information eliminates the profit opportunity over time. If the profits were to be eroded instantaneously, an ideal assumption of frictionless (costless) markets is required, where no profits can be made from analysing available information. Over time, the nature and role of the market participant diminishes to the extent where they are non-existent, as the lack of profits and lack of any analysis of securities provides no benefits. At this point, traditional finance theory allows for the market to collapse altogether, as there are no participants along with any incentive to trading.

A departure to this notion is captured by the Adaptive Market Hypothesis. Consider again the case where prices are stagnant. At this point, it is reasonable to assume that arbitrage opportunities may exist, causing inactive participants to become active and strive for profits via analysis of information. Indeed, while several studies have concluded that price changes are random, no definitive conclusion has been reached (Lo and MacKinlay, 1988). In other words, while a particular study may show that prices changes are random or not random, the sheer volume of studies with either conclusion does not imply a consensus on the hypothesis itself. While academics disagree on whether this implies a violation of the EMH, the sustained profits of traders and other market participants are a clear indication of the lack of *complete* efficiency in the market. One reason for the disagreement perhaps stems from the lack of a testable hypothesis in the EMH. One needs to specify additional criteria, such as investor preference and information structure, to test the EMH. However, these criteria then make testing the EMH a test of the auxiliary hypotheses, which in themselves, cannot be generalised to other markets. In this thesis, market efficiency is tested by examining the return generating process, without making assumptions on the nature of the individual or the manner in which information is reflected in prices.

A new avenue is to then treat the EMH as a reference point as per Farmer and Lo (1999), in that one questions the relative efficiency of markets against the EMH. Conceptually, the EMH can be considered a final state model that is fixed whereas the AMH is considered a dynamic model that reaches the fixed state of the EMH. If non-linear modelling of the market cycle is to be attempted, it first begs the question of whether changes in prices (returns) can be described by a continuous linear or non-linear function. Thus, from a first principles approach, it is necessary to determine if share prices do indeed follow a random walk or whether they are a deterministic (perhaps chaotic) process. Before embarking on understanding the AMH, it is necessary to observe how the concept of market efficiency has evolved since inception.

2.1.1 The history of market efficiency

The notion of market efficiency, indeed all subsequent developments in finance, can be traced back to the 19th century. A French stockbroker, Jules Regnault, attempted to eliminate

a theoretical gap that existed in understanding stock markets. He observed that the price deviation of a security is directly proportional to the square root of time (Regnault, 1863). The longer one looks towards the future price of a security, the greater the volatility present in its future price. An alternative interpretation is that one can predict the price band of a security in the future, using historical price movements. The observation on the behaviour of price deviation was assisted through the use of statistical analysis and provided the foundation of how the current finance community analyses price movements.

Developments in the field of physics led to the concepts of random walks and Brownian motion being introduced in finance by Pearson (1905). According to Pearson (1905), a random walk is a mathematical description of a path of successive random steps. A Gaussian random walk is one where the successive steps have an underlying Gaussian distribution – an enhancement of the discovery by Regnault (1863). This has become quite popular in finance theory (for example, in the use of the Black-Scholes option pricing model).

With the emergence of interest and development in the field of economics, the earliest definition of a market is provided by Gibson (1889) in that "when shares become publicly known in an open market, the value which they acquire may be regarded as the judgement of the best intelligence concerning them". Keynes (1923) stated that investors in financial markets are rewarded not for knowing better than the market what the future has in store, but rather for risk bearing. A typical investor will be compensated at a level commensurate with the level of risk taken – a concept that would later be developed into mean-variance optimisation. One can interpret this statement to also imply that any predictive analysis on share prices would not yield superior results; a resurgence of the definition of a market by Gibson (1889). At this point in time, two seemingly unrelated areas of finance have developed, namely the quantitative aspect of modelling share prices, and the notion of a market which consists of the aggregate of all securities. In the preceding 20th century, these two areas began to merge.

According to Mandelbrot (1963), Mitchell (1915) was the first to note that price distributions are dissimilar to Gaussian population samples. Specifically, Mitchell (1915) noted that price

data since the early 1900s failed to follow a Gaussian distribution as there were too many observations near the mean and the tails (in other words, leptokurtosis was present). This is an important discovery in finance theory as it stands in contrast to subsequent theory built on the assumption of a Gaussian distribution of prices. Indeed, the Mandelbrot hypothesis holds that price distributions follow a power law (Pareto) distribution – a distribution that has leptokurtosis. After results from other authors, combined with the Wall Street Crash of October 1929, there was increasing evidence in favour of leptokurtosis in share price distributions and less belief in market efficiency. However, there were those that were convinced that share price changes were indeed random. Thus, the debate on the prediction and statistical distribution of share prices began.

Cowles (1933) analysed the performance of investment professionals in the United States. In an attempt to test whether investment professionals can forecast future stock prices or select superior stocks to invest in, the author analysed news publications written by investment professionals. He concluded that stock market forecasters lack the ability to forecast perfectly. Indeed, out of the four groups studied, the recommendations from two groups produced below average returns, one above average and the other on par with average returns. In light of the above research, it became important to determine if stock price changes are random in nature. Cowles and Jones (1937) were one of the first authors to show that serial correlation in averaged time series price indices was significant. Analysing the frequency distribution of stock prices over varying levels of time (ranging from daily to yearly), the authors find serial correlation to be present in the higher frequency data more than in the lower frequency data. Roberts (1959) demonstrated that a random walk model was strikingly similar to an actual share price series. In considering weekly price levels and price changes, the author demonstrates, graphically, that there is an equal chance of obtaining a positive change as there is of obtaining a negative change. While the Runs test (described later) is used in the analysis, no other statistical proof is offered to conclude that share price behaviour follows that of a random walk. Osborne (1959) simultaneously showed that the logarithm of share prices follows Brownian motion[2]. In analysing price data, Osborne (1959) shows that the expected gain of investing in a share is zero, implying that the investor should be indifferent in picking which share he wishes to invest in. As the academic community

[2] Also known as a Wiener process in time series, Brownian motion describes the random movement of particles.

began to generate interest in the idea of examining stock prices, so too did theories on their behaviour begin to emerge. Alexander (1961) concluded that while a random walk model best fits the data tested, there was presence of leptokurtosis in the distribution of returns. Mandelbrot (1963) first proposed that the tails of the distribution of returns follow a power law. In subsequent literature, this would become known as Mandelbrot's stable Paretian hypothesis - the hypothesis that stock returns follow a Pareto (power law) distribution in that they exhibit leptokurtosis. Meanwhile Granger and Morgenstern (1963) performed spectral analysis on market prices and found that short-run movements of the series obey the simple random walk hypothesis, but that long-run movements do not and that business cycles were of little or no importance. The authors demonstrate the applicability of a (then) new technique in statistics, that of spectral analysis, to analysing stock prices and returns. While they find minor evidence on the importance of business cycles, the results provide a basis for future studies in examining the seasonal effects of stock price behaviour. The authors further state that their results show that a short term investor (an investor with an investment horizon of less than one year), participates in a fair gamble in that his chances of earning superior returns is left to chance and not his stock picking ability; whereas an investor who chooses a longer term horizon may benefit from analysis of the business cycle. It is interesting to note that the results of Granger and Morgenstern (1963) present room to examine market efficiency over data of differing frequencies, an idea utilised in this thesis. Fama (1963) tested Mandelbrot's stable Paretian hypothesis and concluded that the tested market data conforms to the distribution. The author acknowledges that returns may well fit a power law distribution better than a normal distribution but cautions (rather strongly) against full acceptance of this new hypothesis. Thus, debates around the statistical properties of returns data began in earnest.

Subsequent refinements of random walks, martingales and Brownian motion led Samuelson (1965) to provide the first formal definition of efficient markets in terms of a martingale[3]. The author is particular cognisant of defining the martingale property of stock prices in light of the then ensuing debate on market efficiency. The work quite aptly describes that a less

[3] A martingale is a stochastic process where the next value in the sequence is equal to the present observed value given all prior observed values. Martingale strategies can be traced to early application of probability theory to gambling which resulted strategies which were aimed at producing zero profits; betting at the "fair game" stake

restrictive stochastic model (that of a martingale) of stock price movements is preferred over stricter definitions and that if stock prices do follow a random walk, it not refutable proof that they will *always* follow a random walk, especially if different frequency data is examined. While the terminology of a martingale and random walk may appear similar, there is a distinction between them, namely that in a random walk, the next observation is independent of previous observations, whereas in a martingale, the next observation is a function of previous observations. Fama (1965b) defined an efficient market (described in detail later) somewhat differently and from extensive empirical analysis, concluded that stock prices follow a random walk - the next stock price is a function of the previous stock price and a randomly generated error term. Mandelbrot (1966) proved that in competitive markets with rational risk-neutral investors, successive price changes are dependent on historical prices; they follow a martingale. The author introduced two further components for an efficient market, that of investor rationality and risk appetite. Both have roots in utility theory and the subsequent field of behavioural finance. At this point, an efficient market was one in which prices changes followed a Gaussian distribution and future prices were unpredictable. However, based on the plethora of research discussed previously, there was no single, holistic definition.

While the Efficient Market Hypothesis was publicised by Fama (1970), the term itself was first introduced in the literature by Roberts (1967) in an unpublished manuscript. This definition, as well as the three forms of market efficiency, was later used in the definitive work on the EMH in a series of three articles beginning with Fama (1970). He defines an efficient market as “a market in which prices always 'fully reflect' [relevant] available information”. Following the introduction of the EMH, the debate on whether markets are truly efficient gathered momentum. However, based on the definition by Fama (1970), the two additional elements by Mandelbrot (1966), that of rationality and risk appetite, were included in subsequent research.

2.1.2 Defining the Efficient Market Hypothesis

Following the question of whether past stock prices can be used in predicting future prices, the literature shows two diverse answers to the question. Studies cited previously are divided

amongst chartists (those who purport that there are patterns inherent in stock returns) and random walk theorists, who purport that future stock price changes are independent of past stock price changes. In examining the behaviour of stock market prices, Fama (1965b) explores the theory on random walks and martingales, before empirically testing the hypothesis that stock price changes are random. The results presented form the establishment of the author's later work in defining a theory of efficient markets. Fama (1970) reviews empirical and theoretical work combining, amongst others, the work of Fama (1965b) and Roberts (1967), in developing a theory of efficient markets.

The EMH requires that agents have rational expectations (that is, on average, the population of agents are correct, even when no single agent is) and that these agents update their expectations whenever new information arises. The EMH requires that investors' reactions follow a Gaussian distribution so that no abnormal profits can be realised and that price changes follow a random walk. Recall that under a random walk, successive price changes are independent of each other. Thus, in aggregate, these successive changes are more than likely to be normally (Gaussian) distributed. It is important to note that these two criteria must be satisfied jointly. Thus, any test on market efficiency requires that one test both the normality assumption and the independence assumption. Each of the forms of efficiency, as described by Fama (1965b) requires a differing set of additional requirements to hold true.

The weak form efficiency states that future prices cannot be predicted by analysing past prices. In the long run, investment strategies will not earn excess returns after costs. More specifically, strategies focused on technical analysis will not be able to consistently produce excess returns whereas strategies focused on fundamental analysis may still provide excess returns. Statistically, share prices do not exhibit serial correlation, in other words, there is no dependence on successive price changes, implying that they follow a random walk.

The semi-strong form of market efficiency provides that share prices adjust quickly to public information. Neither a fundamental nor a technical analysis-based strategy will earn returns in excess of the market average. However, those that have access to private information may be able to obtain superior returns.

Lastly, under strong form efficiency, share prices fully reflect both public and private information. Thus, no sustainable superior returns, after costs, can be achieved in the long run.

Tests of the above three forms of market efficiency are varied in approach and conclusion. This thesis approaches the problem of testing for market efficiency in an inductive manner. By observing the returns generating process, it is first sought to determine whether this process follows a random walk (or martingale). If it is found that the process is non-random, the next question is to provide a possible model that fits this return process, with the possible aid of exogenous variables. Given the description of the EMH above, the theoretical foundations of the EMH are now explored.

2.1.3 An axiomatic approach to informationally efficient markets

Samuelson (1965) provided the first theoretically rigorous foundation of the now Efficient Market Hypothesis of Fama (1970). He argues that the unpredictability of future price changes is not a valid basis for tests of information efficiency (recall that information efficiency refers to the instantaneous reflection of new information in the stock price). If one finds that future price changes cannot be predicted, one must be confident that the model used is robust to generalise that conclusion. Given the finding that future price changes cannot be predicted, one cannot assume that this implies that there is efficiency in market participants analysing information and reflecting this in the current price. Rather, Samuelson (1965) viewed the market as efficient where prices were equal to fundamental values when there is perfect competition and all participants have free access to relevant information. Further, Samuelson (1965) also states that actual markets may have such characteristics by chance. It is still arguably difficult to test whether prices are equal to fundamental values without a universally accepted equilibrium price model and means of ascertaining whether there is perfect competition and free access to relevant information in the market one is testing in. However, it is important to note that Fama's (1970) definition of the EMH relates only to price and not fundamental value. A market is efficient if the *price* reflects all available information; this does not imply that the fundamental value has to equate to the price. As such, the framework provided in this thesis relates to modelling the data generating process

with respect to prices (returns) and makes no assumption on the underlying fundamental value of the firm.

From Samuelson's (1965) model of price changes, the fair game theorem emerged in determining future prices. Briefly, this theorem stated that the expected price change, based on available information, is either nil or the market average, implying that the value of the information is either worthwhile and reflected in the market average, or worthless and not expected to change the price. Thus, current market prices will reflect all available information relevant to said prices. Any investor can therefore not profit from any analysis of past prices as this information is already incorporated into the current price. Fama (1965b) attempted to interpret the EMH as an empirically-based, falsifiable theory that could explain the behaviour of share prices. His motivation was different from Samuelson (1965) who strove to show that share prices do not follow a Gaussian distribution and instead follow a Paretian distribution. Samuelson (1965) offered a large-scale view of price behaviour, making minimal assumptions on the investor whereas Fama (1965b) assumed that the individual efficiency of a particular stock price aggregates to create an overall market efficiency. His assumption was that prices at which individual transactions are made are elements of the distribution whose price changes were independent and identically distributed. Fama (1965b) sought for characteristics of markets to support the assumption of Samuelson (1965) but was aware of two contradictory characteristics. First, that there are individuals who are considered leaders by others and are followed by other market participants; and second, there is inertia in the process of information dissemination - positive news is followed more often by further positive news (and *vice versa*). These two contradictions have been pronounced in the behavioural finance literature. The former is now known as herding behaviour - the tendency of investors to follow the group decision as opposed to their own (whether rational or irrational). In South Africa, Seetharam and Britten (2013) outline the literature on herding behaviour studies done in developed and emerging markets and document the herding effect on the JSE. The authors find that investors tend to herd more preceding a severe downturn in the ALSI index than preceding an upturn in the ALSI index. The authors postulate that this is due to investors becoming fearful in times of a recession and greedy in times of an expansion. The latter contradictory characteristic is the momentum effect, first documented by Jegadeesh and Titman (1993). Thus, Fama (1965b) created the image of the sophisticated trader, one where their impact on the market is significant to the extent where they reduce the dispersion

of the distribution of share prices to their expected values. Thus, as the sophisticated trader increases in expertise and in population, their approximations of the share price will converge, in line with the Samuelson (1965) hypothesis.

The works of Fama (1965b) and Samuelson (1965) show two divergent approaches to market efficiency - one which shows efficiency as a state (an axiomatic approach) and the other which shows efficiency as a process (an empirical approach). Thus, Samuelson (1965) defined efficiency as a state in which the conditions of perfect competition, zero transaction costs and complete, freely available information is met. In contrast, Fama (1965b) defined efficiency as the output produced by sophisticated traders. Both methods can be criticised in that Samuelson (1965) did not investigate the reality of his assumptions whereas Fama (1965b) did not analyse if sophisticated traders are necessary or the only influence on price convergence. It is worth noting that no research was conducted by Fama (1965b) into the nature and behaviour of these market agents, but it was rather assumed that their existence can be observed (and confirmed) by examining share price data. Indeed, in a later article by Mandelbrot (1971), it was shown that martingales alone cannot account for the variability of price changes in markets. There was thus a need to examine the agents themselves, as their behaviour is what ultimately drives stock price changes. In an attempt to reconcile the concept of market efficiency with anomalies documented in behavioural finance, Lo (2004) argues that one should view efficiency as the interaction of market participants. As their interaction increases, there is a distinct possibility that information is processed more efficiently, leading prices to eventually reflect all available information. The term "eventually" expands the current thinking on market efficiency, in that there are varying levels of efficiency. Attention is now given to the market participant and his perspective on efficiency.

2.2 A market participant's view on efficiency

There exists a large body of literature on testing the efficacy of a trading rule in earning above average returns after costs. Fama (1965b) followed this approach by stating that if no trading rule can beat the market in the long run, then the market is considered efficient. While this is a practical approach to testing efficiency, it lacks theoretical rigour. Indeed, Campbell,

Lo, MacKinlay and Whitelaw (1998) described these tests as having captured the interest of the financial community due to their practical application. It is due to this practicality that the notion of market efficiency is hardly considered an issue in the lives of financial professionals. From the viewpoint of a financial professional, say a trader, the question of whether markets are efficient or not is irrelevant. The primary drive for the trader is to "survive" by ensuring that his trades generate positive profits. Given that not all traders aim to merely survive, those that wish to thrive should be earning returns above that of the market. Taking into consideration the approach by Fama (1965b), it follows that a trader aiming to thrive would not consistently use the exact same trading rule. This can be rationalised as follows. If a particular trader is successful by following a particular trading rule, other traders would inevitably discover the high profit potential from this rule and adopt it. With many traders adopting the same rule, the finite profit pool is now shared amongst more traders, leading to at least one trader considering an innovation to the rule. If successful, the trader now earns superior profits to his peers, causing the cycle to repeat.

In the short run, there are arguably as many rules (or investors) that earn above average returns as those that earn below average returns. In other words, as the chances of having a positive return are as likely as a negative return, can one argue that the market is efficient in the short run? As one looks at a longer time period, the expectation is to find fewer rules that beat the market consistently after risk and costs have been taken into consideration. Following Alexander (1961), finding a rule that consistently beats the market is evidence in favour of market inefficiency (and *vice versa*). Realistically, one cannot define the set of all possible trading rules in existence, thus one cannot calculate the significance of the success of one rule over another. The outcome would be left to pure chance. Further, if the investment horizon is extended, it may become an impracticality in itself - a rule may have beaten the market over a longer time span than the lifespan of an average investor (or one can assume extremely low discount factors or unlimited intergenerational altruism). The discovery of any winning trading rule (which implicitly uses historical data) suffers from the benefit of hindsight. There is no plausible reason to assume that that specific trading rule will also be a winning rule in the future. As the conditions in the market, such as the number of participants, regulation and innovation, change, so too does the chances of earning abnormal profits. It is thus left to the trader to either adapt a new trading rule or adopt a previously winning rule if he wishes to thrive or merely survive in the "new" market environment. These

changing market conditions have drawn the attention of some academics as one can examine the impact of new information to price changes.

2.2.1 Tests of the speed of adjustment of prices to new information

Anomalies such as the January effect, weekend effect[4] and momentum effect have been widely found as evidence against the EMH. Proponents of the EMH argue that these anomalies are due to selection bias (the desire to focus on interesting subject matter) while other random variable distributions remain outside the attention of EMH critics. Similar arguments can be made for value and growth shares. Indeed, Lakonishok, Shleifer and Vishny (1994) term growth shares glamour shares, in that they are currently popular stock picks for investors. However, as more investors pick these glamour shares, the return gained by each investor diminishes, leading the investor to remove funds from the share. Jung and Shiller (2005) also note that Samuelson's dictum, the phenomenon of a single share following the predictions of the EMH closer than market aggregates, may indeed be prevalent in many research articles. They argue that Samuelson's dictum is more plausible if there is more information about each stock's (firm's) earnings, dividends or cash flow changes than there are of the aggregated market. Given the diversity in the above information for each firm and assuming that half of the firms have positive information with the other half having negative information, a simple aggregation of them may well signal that the aggregate market has not received any new information. Under Jensen's (1978) definition of the EMH (that the EMH is an extension of a zero profit competitive equilibrium condition bridging the gap between the certainty world of classical price theory to the uncertain world of dynamic market behaviour), it is likely that the returns from value shares are due to their more risky nature as well as the higher transaction costs of information acquisition. The dichotomy between price and fundamental value is now explored.

2.2.2 Discrepancies between price and fundamental value

Shiller (2003) provides a survey of tests of this nature, but shows that there is no consensus in their results. While it is shown that the distribution of price changes does not always follow a

[4] The tendency of share returns to be abnormally high in January or on a Monday, respectively.

Gaussian distribution, there is no agreement as to how much deviation nor whether how long the deviations persist, in order to refer to a market as inefficient. While variance bounds tests (discussed in the Methodology chapter) may assist, their underlying martingale distributions do not allow for high variability in price changes as well as providing little insight into the convergence (divergence) process of prices towards (away from) fundamentals. This may lead to a market observer interpreting a divergence as evidence of inefficiency and another as evidence of efficiency after the market has adjusted to it. Thus, a market observer's views are independent of the statistical results of, say, a variance bounds tests. Any test of efficiency would therefore have to isolate the effect of changing fundamentals from the effect of excess volatility (noise). As such, a theory of arbitrage pricing has emerged where one determines the effect of a number of factors (either fundamental or not) on the returns of an asset. This theory is discussed in detail in Section 2.6.

A further definition of efficient markets was provided by Black (1986) where prices in an efficient market would never decrease below 50% or increase above 200% of the fundamental value.

> "Still, the further the price of a stock moves away from value, the faster it will tend to move back. This limits the degree to which it is likely to move away from value. All estimates of value are noisy, so we can never know how far away price is from value. However, we might define an efficient market as one in which price is within a factor of 2 of value, i.e. the price is more than half of value and less than twice value. The factor of 2 is arbitrary, of course. Intuitively, though, it seems reasonable to me, in the light of sources of uncertainty about value and the strength of the forces tending to cause price to return to value. By this definition, I think almost all markets are efficient almost all of the time. Almost all means at least 90%." (Black, 1986, p.533).

The definition of Black (1986), while not quite theoretically grounded, offers a practical insight into market efficiency. However, the author does not account for the historical evolution of variance, a problem linked to an observation by Mandelbrot (1977) that common models of price changes do not allow for the variability, discontinuity and concentration of

price changes in markets. This has dual implications. First, one needs a model of price determination with serial dependence as markets are not able to respond instantaneously to news. Second, a sample from a fractal random process[5] may exhibit features in which a technical analyst may base a recommendation on. Given that the current view on market efficiency is still a binary one, Samuelson's (1965) deduction of the concept of efficiency as unpredictability still holds, implying that the EMH is not falsifiable as it stands. Apart from conceptual grounds of testing market efficiency, it is important to be aware of the practical limitations of financial markets in deducing whether said markets are efficient or not.

From an empirical viewpoint, one can only test what one observes. Therefore, if price is the result of an interaction between a buyer and seller, this is the only observable data one can gather to test market efficiency - one cannot observe the true value of a share, only what market participants pay (receive) for it. As such, the notion or concept of value diminishes in its significance to a researcher as the value of a share is arguably independent of its price (the buyer and seller will agree upon a particular monetary value for a particular share, irrespective of the share's perceived or true monetary value).

2.3 A statistical view of market efficiency

Grossman and Stiglitz (1980) criticised the work of Fama (1965b) in that the efficient market, as defined, requires the existence of sophisticated traders. Further, after sophisticated traders have achieved their goal of eliminating abnormal profit, they would disappear, thus making the market inefficient. This paradox becomes apparent if one views the "lifecycle" of a sophisticated trader. These traders cannot earn above-average returns in an efficient market, thus they would have no incentive to trade. In the long run, it thus becomes relatively more profitable to hold the market portfolio of securities, which is acquired at minimum cost and freely available information. Yet, by holding the market portfolio, there is no incentive for a sophisticated trader to invest in acquiring information in the first place. One comes to the conclusion that the ideal of market efficiency, as described by Fama (1965b) is unattainable as there is no incentive for the emergence of sophisticated traders. In other words, if markets

[5] A random process that appears to have a fractal (or similar) pattern.

were in a state of equilibrium, the end of information acquisition would cause the market to move from equilibrium. The argument of Grossman and Stiglitz (1980) revealed that the assumptions made on the EMH hide complex issues in the actual functioning of markets as well as the number, motives and behaviour of the sophisticated traders themselves. The definition of an efficient market by Jensen (1978) postulates that risk adjusted returns should be compared net of transaction costs, as the cost-benefit analysis of information has an important role in incentivising market participants in acquiring information. There is also the further dichotomy of a statistical and market participant view of market efficiency, the latter of which was discussed above.

Most of the extant literature on market efficiency investigated whether prices in an efficient market followed a random walk and were further normally distributed. While many tests were performed, the argument of Samuelson (1965) was often overlooked; that the unpredictability of price changes is not sufficient to test market efficiency as a rejection of market efficiency may well be due to an inappropriate equilibrium price model being used. The random walk models used formed part of a wider group of models known as martingales, which impose lesser restrictions on the price change distribution. According to Mandelbrot (1977), these restrictions, while more lax, were still too rigid to account for actual price behaviour. After two decades following Samuelson (1965), Lo and MacKinlay (1988) developed a test for variance bounds, which is more appropriate for martingales with heteroscedastic errors. The authors propose examining the variance of stock prices as this will provide more information about the time varying nature of stocks. Using weekly data from 1962 to 1985, the authors develop a specification test (later known as the variance ratio test) to reject the random walk hypothesis. However, the authors state that a rejection of the random walk hypothesis does not imply a rejection of market efficiency. Recall that the EMH provides three forms of market efficiency, implying that if the lowest (weakest) form does not hold, then higher (stronger) forms of market efficiency can still hold. Further developments, which focused on the individual's ability to process information were in the form of rational expectation models of LeRoy (1989). In this model, the author allowed for serial correlations of price changes in an efficient market when risk preferences shifted, implying that the presence of serial correlation cannot refute the EMH. This notion can be considered plausible when viewed over a period of time. Serial correlation implies that patterns exist in the data. Thus, a rational investor would want to capitalise on those patterns

by trading. Over time, the profits generated from these trades would be eliminated as more rational investors identify the same pattern and act upon it. Thus, in the long run, the profits are eliminated, resulting in no abnormal profits from being made by any market participant. This preserves the definition of market efficiency according to the EMH.

Fama (1976) stated that the EMH is not an empirically testable (falsifiable) hypothesis as a refutation of the EMH can either point to the EMH not holding or the equilibrium price model not holding. In summary, four critiques of the EMH hamper the ability to test the hypothesis. First, inappropriate models of price changes are most often used; second, the joint hypothesis problem of testing both market efficiency and the equilibrium price model; third, the theoretical possibility of serial correlation in an efficient market and fourth, the lack of an emergence of sophisticated traders. The first two point to varied statistical treatments in solving for a test of market efficiency, while the latter two point towards a philosophical approach in defining efficiency. In an attempt to overcome these criticisms, the use of a neural network, along with the framework of arbitrage pricing theory (Ross, 1976) is used in selecting the input variables. This theory provides an intuitive and encompassing view of selecting factors which can influence price changes. While there is some form of bias in that the initial dataset contains variables selected by the researcher, the use of both an unspecified *a priori* model as well as data sampling governed the APT framework, should overcome the joint hypothesis problem commonly faced when using the CAPM. Further, the neural network does not provide a list of significant variables as one only knows that from the input variables used, they are combined in some manner to provide the most accurate output. Thus, a rejection of market efficiency does not involve rejection of the APT framework used in this study. The possibility of serial correlation in an efficient market is discussed in Section 2.5 below, whilst the emergence of sophisticated traders is given indirectly by testing the cyclical nature of market efficiency. If one has a series of conclusions of whether markets are efficient or not at each point in time, then these differing conclusions can arguably be due to the emergence (divergence) of sophisticated traders.

The critiques raised above arose primarily due to Fama (1965b) attempting to create a framework around actual financial markets without diverging from the axioms set out by Samuelson (1965). This led to a range of tests being conducted on the EMH, when in reality,

the EMH could never truly be falsified, without first providing a universally accepted equilibrium pricing model. Such models are now explored below.

2.3.1 Models of market efficiency

Fama (1970) provides a review of the (then) literature on market efficiency. He states that much of the theoretical development came after empirical results were found. As the primary statement of the EMH is that prices fully reflect all available information, one must show that the expected return of a security is a function of risk based on some set of information. This describes the fair game model.

2.3.1.1 Fair game model

The fair game or Martingale model states that a stochastic process with the condition of an information set is a fair game (where the expectation of a variable is equal to the actual value of the variable) given by

$$E(x_{t+1}|I_t) = 0 \qquad \{1\}$$

Fama (1970) incorporated this model into the EMH. It is expressed as follows

$$x_{j,t+1} = p_{j,t+1} - E(p_{j,t+1}|I_t) \qquad \{2\}$$

with

$$E(x_{t+1}|I_t) = E(p_{j,t+1} - E(p_{j,t+1}|I_t)) \qquad \{3\}$$

where $x_{j,t+1}$ is the excess market return of security *j* at time *t+1*, $p_{j,t+1}$ is the actual price of security *j* at time *t+1* and $E(p_{j,t+1})$ is the expected price of security *j* at time *t+1* given the information set I_t.

According to the Fair Game model, the excess market return of a security should be zero, implying that once all information is incorporated into the current price, one cannot earn returns above that of the market.

2.3.1.2 The submartingale model

The submartingale model makes a small adjustment to the Fair Game model - the expected return can also be positive. Recall that Samuelson's (1965) hypothesis implies that the value of the information is either worthwhile and reflected in the market average, or worthless and not expected to change price. This adjustment implies that the price of a security is expected to increase over time, perhaps also due to the increased level of risk inherent in the security. The submartingale model is represented as

$$E\left(\frac{r_{i-1}}{I_t}\right) \geq P_{j,t} \quad \{4\}$$

$$E\left(\frac{r_{i-1}}{I_t}\right) = \frac{E\left(\frac{r_{i-1}}{I_t}\right)}{P_{j,t}} \geq 0 \quad \{5\}$$

The model states that the expected return of the security follows a submartingale, conditional on the information set I_t. The information set itself holds no value in forecasting security prices, except that the expected return can be equal to or greater than zero. This implies that no trading rule based only on the information set can achieve greater expected returns than a buy and hold strategy during the future period in question (Fama, 1970). Given the choice between analysis of price patterns and of the financial statements of a company, Fama (1970) would argue that in a weak form efficient market, analysing price patterns holds no value in earning abnormal returns.

2.3.1.3 The Random Walk model

The intrinsic value of a share is measured by the sum of future discounted cash flows accruable to investors. Any new information that can be expected to change a company's future performance must be immediately reflected in the share price as delays in this diffusion can be exploited by certain individuals to forecast future profitability. Thus, prices should only be able to respond to new information. Since this information arrives randomly, prices must fluctuate unpredictably. The Random Walk model of share prices is represented as follows.

$$P_{t+1} = P_t + \varepsilon_{t+1} \quad \{6\}$$

Where P_{t+1} is the price of a security at time $t+1$ and ε_{t+1} is a random error term with zero mean and finite variance.

The equation above indicates that the future price of a security is based on the arrival of new and unpredictable information. This implies that price changes are independent of past price changes. Fama (1970) argues that the random walk model is an extension of the fair game model in that the latter indicates conditions of the market equilibrium that can be stated in terms of the return generating process of the former model. Tests of the weak form of the EMH consider the above three models in their hypothesis as the determination of whether the market is weak form efficient or not is a function of both the return generating process and of the tests employed.

2.3.2 Weak form tests

The tests of the weak form of the EMH are synonymous with testing the random walk hypothesis; the notion that stock price changes are random and thus unpredictable. As with most literature on testing the EMH, tests done on the weak form of the EMH show conflicting results. Tests of the weak form are based on examining the interrelationship between current and past prices. Practically, the runs test, tests for autocorrelation and the

variance ratio test have been used to test for weak form efficiency.[6] Given the array of literature on the topic, a select few works will be discussed here.

Sharma and Kennedy (1977) employ the runs test to test for weak form efficiency on the Bombay, London and New York Stock Exchanges. Using monthly observations over an 11 year period, they find that shares on the Bombay Stock Exchange do follow a random walk. A combination of approaches is adopted by Dickinson and Muragu (1994) who also find that share prices on the Nairobi Stock Exchange follow a random walk. Their study examines weekly and monthly data for a sample of 30 of the most liquid shares on the exchange. Employing correlation and runs tests, the authors find that the majority of share prices examined follow a random walk. While they then generalise this result to conclude that the overall Nairobi stock market follows a random walk, the authors are careful to place their results in the context of literature on the EMH. They explicitly state that while their results show evidence in favour of a random walk, they are cautious to imply that the Nairobi stock market is weak form efficient. A possible reason for this hesitation is that one cannot easily generalise a result of a particular sample period and data frequency to time periods and data frequencies not used in the study. Further, the methodology used needs to be robust enough to provide comprehensive evidence that can hold across out-of-sample data. Other researchers, such as Seddighi and Nian (2004) use spectral analysis and ARCH tests for detecting if the Chinese market is weak form efficient. The authors conclude on a particularly small sample of daily share returns from the Shanghai Stock Exchange that the Chinese stock market is weak form efficient. The frequency of the time series under observation has also been investigated to determine if a result that holds for a particular frequency will hold at other frequencies. For example, Groenewold (1997) uses daily, weekly and monthly data to determine if the Australian and New Zealand Stock Exchanges are weak and semi-strong efficient. The author employs the popular tests of autocorrelation, runs and cointegration on a 17 year sample period. The results however are mixed - the returns appear to have some predictability according to the autocorrelation coefficient but are stationary in the long run. This could possibly imply that if one uses higher frequency data to test market efficiency, one might find a short term "memory" of the series, which dissipates over lower frequency data. Thus, to examine this notion, daily, weekly and monthly data are used in this thesis to

[6] Each of these tests are described in the Methodology chapter.

examine market efficiency. While much literature exists on tests of the weak form of the EMH, it becomes redundant to mention them as there was no conclusive evidence of whether emerging or developed markets are weak form efficient.

2.3.2.1 Calendar effects

The finding of abnormal returns in markets during particular time periods has prompted researchers to examine if there are any peculiarities in markets which can possibly explain these anomalies. A branch of literature which focuses on calendar effects explaining abnormal returns is vast, with various reasons provided as to why abnormal returns occur. For example, Sullivan, Timmerman and White (2001) provide a short overview of the calendar effects literature, citing observations from particular days of the week that affect returns, to weeks of the year that affect returns. From the survey of the literature, it is common to see these empirical observations being preceeded by a theoretical model explaining their existence. Thus, the authors question whether these effects are due to the researcher's ability to data mine. Here, the notion of data mining refers to testing hypotheses that are not independent of the data – in other words, the data drives the hypothesis, as opposed to some theoretical basis that drives the model. The authors mention that if multiple models (or tests) are applied to a data set, it is likely that some of them are likely to provide a positive outcome. However, if the tests are well motivated by their theoretical foundations, along with some consensus of results, then it is likely that the overall outcome is not a feature of data mining.

Thus far, common tests for weak form efficiency include: the runs test, examining autocorrelation coefficients and the ADF test for stationarity. Poterba and Summers (1988) and Lo and MacKinlay (1988) provided the foundation for the variance ratio (VR) test of the random walk hypothesis. This test compares the variance of the stock return series against stationary alternatives, under the assumption that the variance of random walk increments will be linear across the sample. The VR test can be used to test secondary hypotheses of the random walk, specifically whether stock prices mean revert. While the concept of the test is straightforward, it is often difficult to implement in practice as the test relies on overlapping data in computing the variance of long term horizon returns. Lo and MacKinlay (1988)

suggest this approach as it can improve the statistical power of the test and suggest that an asymptotic distribution be used instead of the exact distribution of the test. However, while other tests have been developed to remedy the shortcomings of the VR test, the VR test still remains popular in literature. Given an array of tests to use in examining the return generating process of share returns, one should also not assume that share returns follow certain pre-specified conditions. Three such assumptions are now discussed.

2.3.3 Assumptions underlying share returns

In testing the weak form of the EMH, three assumptions need also be considered, namely that of normality, independence and stationarity of returns.

2.3.3.1 The normality assumption

Prior to an alternative by Mandelbrot (1963), the assumption of normality in share returns was scarcely questioned. Mandelbrot (1963) conducted investigations into both the excess kurtosis and skewness of return distributions and thus developed an alternative Power Law hypothesis of distributions based on his findings. His position was later reiterated as he noted that "Bachelier's assumption, that the marginal distribution of $L(t,T)$ (returns) is Gaussian with vanishing expectation, might be convenient, but virtually every student of the distribution of prices has commented on their leptokurtic (i.e., very long-tailed) character." (Mandelbrot, 1966, p.396). Thus, while the normality assumption is required for the EMH, many practical tests of the EMH show that this assumption is violated. However, Fama (1965b) does not see this violation as evidence that the EMH does not hold.

Fama (1965b) studies the statistical properties of returns using shares on the Dow Jones Industrial Average (DJIA). He finds that a greater proportion of observations are centred around the mean as well as in the tails of the distribution. Further, when examining extreme tail observations (those that are beyond five standard deviations from the mean), he finds that they are almost 2000 times greater than that implied by a normal distribution. These findings indicate leptokurtic behaviour of the returns and Fama (1965b) concludes that a normal

distribution is ill fitting to the data. Praetz (1972) uses traditional goodness of fit measures on the Sydney Stock Exchange and finds similar leptokurtic behaviour. He further offers an alternative distribution, based on Brownian motion, which is claimed to fit the data better than that of the Pareto or Gaussian distribution. Officer (1972) has similar findings over a longer time period (1926 to 1968) on data from the Centre for Research in Security Prices (CRSP) database. He finds that the distributions are reasonably stable across time, yet not across the sample of stocks used. Further, under differing frequencies, the stability of the distribution changes - daily returns produce a stable distribution only up to 20 days, whereas monthly distributions are stable up to 5 months. These results point towards examining efficiency using differing frequencies of data as the results may not always hold or be generalised if only one particular sampling frequency is used. This approach is adopted in this study as daily, weekly and monthly return data are examined.

The effect of non-normality can also been seen in event studies. Brown and Warner (1985) quantify the level of kurtosis in shares in the CRSP database. The authors find that the kurtosis detected is more than double that in a normal distribution and that the frequency of data plays a significant role in the conclusion of non-normality. Specifically, daily returns will exhibit greater departures from normality relative to monthly returns. Indeed, the presence of leptokurtosis is found in numerous studies and pointed out by Engle and Patton (2001). The latter authors show that the range of kurtosis is normally between 4 and 50 times that required in a normal distribution.

From a risk based perspective, Arditti (1967) is among the first set of authors to show that leptokurtosis is also accompanied by asymmetry in return distributions. The author hypothesises that a risk averse investor will be unwilling to invest if the investment will potentially yield a higher loss relative to its gain. This asymmetry can be captured by skewness – where the given outcome is more likely overall, but the skewed distributions affect the likelihood of this outcome. Using cross-sectional analysis, Arditti (1967) established factors that affect returns of firms during 1946 to 1963. Skewness was found to be significantly negatively related, implying that investors prefer positive skewness (as positive skewness implies that there is a higher likelihood of an observation being greater than the mean).

2.3.3.2 The statistical independence assumption

Durbin and Watson (1950) describe independence as the serial correlation function of returns that should decay to zero. It is given by

$$C(t) = corr\big(r(t), r(t + Dt)\big) = \rho \qquad \{7\}$$

Where *C(t)* is the serial correlation coefficient of order *t*, *r(t)* is the return of a given series at time *t* and *Dt* is the time scale. A market is thus described as efficient in the absence of linear serial correlation. Further, if serial correlation is present, then the anomaly is short lived. The assumption of independence can be viewed either from a statistical perspective or from an investor's perspective. If an investor finds that returns are not independent, then investors can theoretically use knowledge of past returns to increase future profits (Fama, 1965b).

Kendall (1953) studies the properties of returns and finds that the pattern of events in a price series is less systematic than what is generally accepted. He concludes that these price changes follow a random walk and are thus independent. Further, the author argues that it is generally difficult (at least at the time) to distinguish between a true random series and one where the systematic element is particularly weak. This implies that when testing any hypothesis, one should take caution to the results and model(s) used. Lastly, the author states that given his results on a lack of serial correlation in the sample of stock prices, he argues that it is near impossible to predict values, in their case one week ahead, without any additional information. While Fama (1965b) states that it is difficult to find a series that conforms to the independence assumption, statistical independence holds even if some level of dependence is present. Further, the simplest explanation for the assumption of independence is due to the arrival of new information, which does not follow any consistent pattern. After testing returns on the DJIA, he finds that most follow the independence assumption with the remainder being serially correlated but with the serial correlation decreasing at higher orders. When correlation is statistically significant, they are low enough to ignore any statistical or practical implications. Noting that the empirical evidence for market efficiency was publicised before the theory, one questions whether the results of Fama (1965b) were taken into consideration in developing higher hurdles for the EMH. Over the

long run, Campbell *et al.* (1998) show that the independence assumption is violated. They test returns of shares on the CRSP value and equally weighted index and find that there is significant first order serial correlation in weekly and monthly returns. Further, the serial correlation decayed slower on the equally weighted index than the value weighted index. This implies that market capitalisation plays a role in efficiency. To test this hypothesis, the authors employ VR tests and find that indeed, market capitalisation plays a role in determining whether the aggregate stock market, comprising of individual stocks, is efficient.

2.3.3.3 The stationarity assumption

The third assumption of returns is that of stationarity. According to Mandelbrot (1966), stationarity implies that the statistical moments of the distribution do not change from one sample to another. Giannopoulos (2000) argues that while the evidence on the stationarity of return distributions is inconclusive, the non-stationarity of return variances is widely recognised. Further, Cont (2001) mentions that seasonal effects (such as the January effect or weekend effect) may confound the tests of stationarity.

Gibbons and Hess (1981) show that the distribution of returns is not identical over all days of the week and provide evidence for the Monday effect; where returns on a Monday have a higher first and second statistical moment. Tests were run on daily data of the S&P 500 index, the CRSP database and shares on the DJIA from 1962 to 1978. They find that the returns distribution is not equal across time, yet returns on Mondays are lower than expected. It is concluded that there exists seasonality in the daily returns and that it is most likely caused by a persistently negative mean return on a Monday.

Taylor (1997) focused on the time varying property of variance in his study of share returns. The author shows that the absolute and square transformations of U.S. share returns are good proxies for volatility and exhibit high levels of first order serial correlation. Further, this correlation over an extended period can imply that there is a time varying structure in variances over the sample period of 1966 to 1976. Following each assumption, literature has progressed to examining the time varying nature of share returns, with particular emphasis on the second moment of the return distribution, volatility.

2.3.4 The behaviour of share returns

Volatility, as defined by Poon (2005) is the spread of all likely outcomes of an uncertain variable. While volatility is one of several factors in a share return distribution, it has an important role in portfolio management, derivative pricing and risk management. Similar to share returns, one needs to check for the presence of volatility clustering, persistence, leveraged effects and mean reversion to determine if the chosen price model is appropriate under the data set used.

2.3.4.1 Volatility clustering

Mandelbrot (1963) defines volatility clustering as large changes in price that tend to be followed by further large changes. This implies that a return series can experience times of stability and instability and that the periods of instability can be persistent (Poon, 2005). Thus, in testing the decay of the autocorrelation function on stock returns, if volatility clustering is present, this decay will be prolonged. Given that stock returns are driven by the interaction of market participants, arbitrageurs who observe a linear trend in returns will exploit this trend through a particular investment strategy. This further drives the persistence of the trend. The detection of volatility clustering is fairly straightforward in that a plot of the time series can easily expose areas where there is clustering (Engle, 2001). Jacobsen and Dannenburg (2003) show that volatility clustering can be statistically observed by Ljung-Box statistics on returns of six international markets. Their results show that all markets examined have volatility clustering at daily and weekly frequencies at all lags. These findings hold for bi-weekly observations but not for monthly observations again highlighting the importance of investigating differing sampling frequencies in the investigation of market efficiency.

2.3.4.2 Volatility persistence

Volatility persistence is closely related to clustering in that volatility clustering implies volatility persistence if extended periods of time are characterised by greater variability in returns. This suggests that the variability must have a degree of persistence to be identified initially. Engle (2001), *inter alia*, refers to this as long term memory, where a single shock will have a noticeable and persistent impact on future volatility. McMillan and Ruiz (2009)

describe the standard approach to detecting volatility persistence through a sample serial correlation function for a non-linear transformation of the returns. If an extended period of time lapses until the serial correlation declines to zero, then the series is said to have a long term memory. Particular time series models, such as Vector AutoRegression (VAR) models[7], and some types of neural networks have been introduced to model this long term memory. These are discussed in Section 2.5. Ding, Granger and Engle (1993) examine the long term memory property of the S&P 500 over the period 1928 to 1991. They find that while first order serial correlation is statistically present, it is of small magnitude and short lived. However, when the transformation of returns is examined, the serial correlation is statistically significant over longer time periods.

2.3.4.3 The Leverage effect

While the original definition of the leverage effect is according to Black (1976), in that there is a negative relationship between share prices and the debt-to-equity ratio of a firm; the later definition instead focuses on the relationship between share prices and volatility (Engle and Patton, 2001). The leverage effect implies that the relationship between returns and volatility is asymmetric, in that negative shocks will have a greater effect on both variables than positive shocks.

Haugen, Talmor and Torous (1991) find evidence of a negative relationship between volatility and returns on the DJIA between 1897 and 1988. Their results show that, following an increase in volatility, there is a decrease in average returns for a four week period. Indeed, this asymmetric response can be indicative of a non-linear risk aversion function of investors, an idea linked to the loss function of Kahneman and Tversky (1979). With the considerations of volatility and assumptions underlying share returns in mind, the next section provides a first principles approach to modelling share prices (returns).

[7] A time series model that captures linear interdependencies amongst multiple time series.

2.4 Modelling share prices

At the heart of the EMH lies the question of share price predictability. While a vast amount of research has been conducted in this area, no definitive answer has been reached. To forecast the returns of a share (or any asset), the returns must necessarily be correlated across time (Skaradzinski, 2003).

However, the mere notion of forecasting share prices goes against the random walk model of Fama (1965b) – where a security's returns are independent and normally distributed. Studies by Conrad and Kaul (1998) and Lo and MacKinlay (1988) rejected the notion of the random walk as described by Fama (1965b) due to the existence of time-varying parameters and sampling dependence, respectively. However, as a lack of correlation does not necessarily imply independence, studies have used more sophisticated methods of testing which rely on higher-order statistical moments of the distribution of returns. If a higher-order dependence is found, then the underlying data exhibits non-linear behaviour. The independent variables used in forecasting may be past values of the asset's returns, micro-economic or macro-economic in nature. According to the EMH, prices can be modelled as a function of noise and the past share price.

$$S_{t+1} = S_t + \varepsilon_t \qquad \{8\}$$

where ε_t is a white noise error term. Thus, the best estimate of the future share price is the current share price

$$\widehat{S_{t+1}} = S_t \qquad \{9\}$$

If a series of share prices is thus truly a random walk, then the best estimate of the future share price is the current share price. If we now assume that share prices can be predicted, the equation is modified somewhat as follows:

$$S_{t+1} = S_t + f(S_t + S_{t-1} + S_{t-2} + \ldots + S_{t-n+1}) + \varepsilon_t \qquad \{10\}$$

where, in addition to the above specifications, f is a (possibly non-linear) function of past share prices. Similarly, the best estimate of the future share price is given by:

$$\widehat{S_{t+1}} = S_t + f(S_t + S_{t-1} + S_{t-2} + \ldots + S_{t-n+1}) + \varepsilon_t \quad \{11\}$$

If the time series in question were to have a trend, then prediction using the above equation results in highly inaccurate outputs. Granger and Newbold (1986) propose that the series be transformed to represent first order differences, commonly referred to as returns if share price data is used.

$$R_{t+1} = f(R_t, R_{t-1}, R_{t-2}, \ldots, R_{t-n+1}) + \epsilon_t \quad \{12\}$$

Where

$$R_{t+1} \triangleq S_{t+1} - S_t \quad \{13\}$$

and ϵ_t is white noise. In this specification, the best estimate of R_{t+1} is given by

$$\widehat{R_{t+1}} = f(R_t, R_{t-1}, R_{t-2}, \ldots, R_{t-n+1}) \quad \{14\}$$

In other words, assuming an element of predictability of share prices, the best estimate of a future share price is some function of past share prices. With the aid of time series econometrics, many models have been proposed to model share prices. These are discussed below.

2.5 Time Series Methods

A set of data points measured at uniform periods of time is referred to as a time series. To model a time series, one needs to be aware of the varying types of seasonality, stationarity and determinism (level of randomness) present in the series as the presence of each can point towards a different model. Often, in analysing a time series, one can mistake the presence of chaos in the series as randomness. Chaos can be defined as the irregular behaviour of solutions to deterministic equations of motion (Casdagli, 1991). The necessary requirement is

that the system of equations be non-linear in order to generate chaotic solutions as a linear system will necessarily generate a trend in its output. These outputs are often mistaken as random time series and are only accurate for a length of time governed by the errors of the initial conditions and the Lyapunov exponent[8] of the system. The following sub-sections discuss the various methods of modelling and explaining time series.

2.5.1 Exponential Smoothing

Exponential smoothing (ES) methods were first developed by Holt (1958). These methods were widely used for business and industrial applications but were often considered a collection of *ad hoc* techniques by academics. Pegels (1969) provided a means of classifying a time series by its trend and seasonal patterns. Both can be linear (additive), non-linear (multiplicative) or neither, giving rise to nine different stochastic models. By graphical illustration of the time series, the classification by Pegels (1969) assists with choosing the best forecasting model to use. Box and Jenkins (1970), *inter alia*, showed that some linear ES forecasts were special cases of ARIMA models. Indeed, the simple ES model can be classified as an ARIMA (0,1,1) model (refer to Definitions page) with no constant term. Snyder (1985) showed that simple ES methods can be considered to originate from an innovation state space model (a model with a single error source). This work prompted later research into state space models and ES methods.

The classification hierarchy by Hyndman, Koehler, Snyder and Grose (2002) describes the various ES methods. Each ES method can consist of one of five types of trend (none, additive, damped additive, multiplicative and damped multiplicative)[9] and one of three types of seasonality (none, additive or multiplicative). This gives rise to 15 different methods, the most common being that of Simple Exponential Smoothing (which has no trend and no seasonality in the data). Further, the authors provide a theoretical framework which maps ES methods to a state space, showing that they are in the same taxonomy as ARIMA models.

[8] The Lyapunov exponent describes the exponential divergence of the output vectors in a chaotic system.

[9] Where a damped additive trend refers to a time series that has an additive trend that decays over time and a damped multiplicative trend refers to a time series that has a multiplicative trend that decays over time.

2.5.2 Prediction intervals

A criticism of ES was that it could not produce prediction intervals for its forecasts. The first analytical approach to this problem by Brown (1963) was to assume that the time series were deterministic functions of time and white noise (refer to Definitions page). If this held true, then a regression model could be used instead of an ES method. This assumption was heavily criticised by Newbold and Bos (1989). The authors note that under the assumption that the time series were deterministic functions of time and white noise, one would: overestimate false signals (Type 1 error), misestimate the probability of the forecast value, misjudge appropriate starting values for the ES method and incorrectly assume that the forecast errors are serially correlated. Other authors since attempted to obtain prediction intervals by examining the equivalence of ES methods and statistical models. In a follow up study, Hyndman, Koehler, Ord and Snyder (2005) used state space models to derive analytical prediction intervals for 15 ES methods, providing a comprehensive algebraic approach to handling the prediction distribution problem (that an ES model would provide estimates, but not a distribution of forecasts). Given the exploration into ES methods, their more general forms, that of ARIMA models, are now briefly discussed.

2.5.3 ARIMA models

Early attempts to study time series in the 20th century began with the idea of a deterministic world, where a change to an initial condition did not result in a different outcome. Yule (1927) provided the first significant contribution of regarding every time series as a stochastic process, where a change in the initial state produces a different final outcome. As such, the concept of an autoregressive (AR) model and moving average (MA) model was developed. Wold's decomposition theorem[10] led Kolmogorov (1941) to formulate a solution to the problem of linear forecasting (and later the Kolmogorov Smirnov test for normality). The work of Box and Jenkins (1970) integrated the then existing knowledge on time series and has become a staple addition to any time series course. The Box-Jenkins method is widely used in first testing for stationarity and seasonality, and then proceeding to specify and evaluate the model. With the advent of the computer, autoregressive integrated moving

[10] Every covariance-stationary time series can be decomposed into the sum of one deterministic and one non-deterministic series.

average (ARIMA) models could be developed and used in forecasting discrete time series processes through their univariate forms.

2.5.4 Univariate models

During the 1960s, the selection of an ARIMA model was largely left to the researcher's judgement, as there was no algorithm available to specify the model correctly. Since then, information criterion techniques have been developed, such as the Akaike Information Criterion (AIC) and the Bayes Information Criterion (BIC). Often, it becomes a task of minimising these criteria that would result in the best model fit as one would prefer to have estimates as close as possible to actual values to show that the model best describes the data.

There are a number of methods for estimating the parameters of an ARIMA model, yet they are prone to error when there are large differences in the finite sample properties. Newbold, Agiakloglou and Miller (1994) showed that this difference is significant across the then available software packages and can result in inaccurate forecasts. As a means to overcome the problem, the authors suggest the use of full maximum likelihood estimation[11] to ensure the parameters are statistically consistent. If a time series is known to follow a univariate ARIMA model, forecasts using disaggregated observations[12] are as good as using aggregated observations under the MSE criterion.

As an alternate to the univariate ARIMA model, Parzen (1982) proposes an ARARMA methodology where the time series is transformed from a long term memory AR filter to a short term memory filter. Using data for airline passengers, Parzen (1982) shows that the ARARMA model is a better fit than other more traditional time series models. Meade and Smith (1985) are part of the few authors who test the ARARMA methodology and show that it achieves a significantly low Mean Absolute Percentage Error (MAPE) for longer forecast

[11] A method of estimating the parameters of a statistical model, maximum likelihood estimation provides estimates of the mean and variance of a distribution given sample information.

[12] (Dis)Aggregated observations - Observations combined (removed) from several measurements.

horizons. While software is available for implementation, these methods are often opaque in that the researcher cannot fully describe the model (it is considered a black box). While there are guidelines for the choice of automatic forecasting methods, Me'Lard and Pasteels (2000) suggest the use of an Expert System[13] as the expert system can more optimally configure the parameters of the model, speeding up the time taken to produce results and quite possibly producing more accurate results.

2.5.5 Non-linear models

Compared to the study of linear time series, the development of non-linear time series is still in its infancy (De Gooijer and Hyndman, 2006).The first work in this area is by Volterra (1930) who showed that any continuous non-linear function can be approximated by a finite series with a memory property, later known as a Volterra series. While the probabilistic properties of these models have been studied, little exists in the problem of parameter estimation, model fitting and forecasting. Poskitt and Tremayne (1986) attribute this to the lack of computational power at the time as well as the complexity of the Wiener model itself. While linearity in itself can solve many practical applications, it is often restricted by the existence of complex real world problems. One hindrance of the forecasting ability of non-linear models was pointed out by De Gooijer and Kumar (1992) in that the models made it difficult to obtain analytical expressions for closed-form multi-step ahead forecasts. The models could not be applied (non-analytical) by a researcher to obtain a finite valued solution (the solution was not closed-form). In principal, the Chapman-Kolmogorov relationship (the mapping of joint probability distributions to a stochastic process) can be used to obtain exact least squares multi-step ahead forecasts through integration techniques and currently, these forecasts have been obtained through Monte Carlo simulation or bootstrapping approaches. The latter approach is preferred as it requires no assumptions about the distribution of the error process. Indeed, Clements, Franses and Swanson (2004) concluded that "... the day is still long off when simple, reliable, and easy to use non-linear model specification, estimation, and forecasting procedures will be readily available." Four such non-linear models are presented below.

[13] A program that mimics the decision-capability of a human being – discussed in Chapter 3.

2.5.5.1 ARCH and GARCH models

A feature of financial time series is that there are periods of high and low volatility which are often clustered together. This volatility clustering is ideally suited to be modelled by autoregressive conditional heteroscedastic (ARCH) models of Engle (1982). These models describe the conditional variance as a deterministic (quadratic) function of past returns. As the variance is known at time *t-1*, one step and multi-step ahead forecasts can be made. The more general form of ARCH model is given by GARCH models where there are additional dependencies on the lag of the condition variance. These models are fairly similar to ARIMA models and thus share many statistical properties. Sabbatini and Linton (1998) test a simple GARCH (1,1) model on daily returns of the Swiss market index and find that the out-of-sample forecasts were not accurate. Engle and Ng (1993) point out that asymmetric volatility is often present in financial returns and their conditional variances. Negative (positive) returns are generally associated with an upward (downward) revision of the conditional volatility. As such, researchers have developed GARCH type models to account for this asymmetric volatility.

2.5.5.2 Long term memory models

When the integration parameter, *d*, in the ARIMA process is fractional and greater than zero, the process is said to have a long term memory. This implies that the observations that are a long time span apart have some sort of dependence between them. Stationary long term memory models or fractionally integrated models (ARFIMA) models have also been developed to allow real values (as opposed to integer values) of the integration parameter. These are thus more apt to modelling long term dependence as the integration parameter can now take on more values. Souza and Smith (2002) investigated the effect of different frequencies of data on ARFIMA models. They find that the bias in the fractional parameter of a non-aggregated series is influenced by the short run autoregressive and moving average parameters.

2.5.5.3 SETAR Models[14]

One of the initial applications of non-linear models to business cycles was shown by Hamilton (1989)'s application of Markov Switching techniques. These non-linear models assumed that changes to market phases were governed by an unobserved Markov chain (a process where the next state depends only on the current state). This assumption meant that the exact times a regime (market phase) change occurred were unknown (the *unobserved* part of the assumption), and could only be estimated using probabilities (Hamilton, 1989). Another property of Markov models is that the change (or switch) between regimes is abrupt. In financial markets, it is often difficult to justify this assumption. Further, the changes between an expansionary and contractionary phase of the market cycle need not necessarily be symmetric. It can therefore be inferred that modelling the changes between these regimes of the business cycle can be problematic as they can be either be symmetric or asymmetric and is an issue that STAR models are aptly suited towards. Investors have heterogeneous beliefs, different time horizons and learning speeds (see Harrison and Kreps, 1978 and Bernatzi and Thaler, 1995). These all point to a gradual change in markets as opposed to a more abrupt one. Thus, a new family of models were developed, namely Transition Autoregressive (TAR) models, where they address the issue of a change between regimes. In TAR models, movements are governed by an observed variable and are referred to as Self Exciting TAR (SETAR) models when the observed variable is a lag of the dependent variable. Tong (1983) provided an extensive discussion of Self-Exciting Threshold AR (SETAR) models. These models are piecewise linear models that "partition" the non-linear time series into linear pieces, making estimation of the overall model quicker and less computationally expensive. Other modifications to these models include Threshold VAR (TVAR) models and continuous threshold AR (CTAR) models. While the CTAR models provide highly accurate estimates, they are often impractical due to the higher dimensional integration involved in parameter estimation.

Through examination of the literature on modelling techniques, it emerges that as the accuracy of the model increases, the ability to estimate its parameters and interpret the model itself decreases. This phenomenon is seen in the case of neural networks discussed below. The author states that a small amount of outliers in a time series can often mask the simplicity

[14] Portions of this sub-section are taken from Seetharam and Britten (2015).

of the series itself. This brings into question the use of macro-economic and micro-economic variables in addition to lagged dependent variables. McMillan (2005) provides international evidence in favour of non-linear modelling of financial markets. An interesting avenue of research is explored by the author linking non-linear behaviour of share prices to performance of noise traders, in spirit of the behavioural finance literature of asset prices being dictated by the interaction of noise traders and sophisticated traders. The results show that a non-linear model is able to capture the effect of noise traders on share prices as well as providing significant gain in forecasting prices out-of-sample for Asian-Pacific economies.

Alagidede and Panagiotidis (2009) provide evidence on testing time series models in several African countries. The authors test each country index for the presence of non-linearity and then proceed to model returns appropriately. Using daily closing prices, the authors find non-linearity in most of their sample, with the exception of South Africa. Bonga-Bonga and Makakabule (2010) use a Smooth Transition Regressive (STR) model is used to investigate the relationship between macroeconomic variables and stock returns. The difference between the STAR and STR models is that the former uses lagged values of the independent variable. Whilst the modelling approach is similar, direct comparison of results between STAR and STR models is inappropriate. van Gysen, Huang and Kruger (2013) conduct a comprehensive study of linear and non-linear modelling techniques in forecasting returns on the JSE. The authors find that non-linear methods are favoured over their linear counterparts, but less so during turbulent market conditions, such as the financial crisis between 2007 and 2009. Specifically, Markov switching models provide the most accuracy from the family of non-linear models considered.

2.5.5.4 Neural Networks

While ANNs are adept at forecasting non-linear time series, some have questioned their accuracy. For example, Tkacz (2001) shows that the forecasts of an ANN are outperformed by a naive random walk model. Some attention has also been given to define the border between ANNs and traditional techniques. Balkin and Ord (2000) show that ANNs can work better for high frequency data and also stress the importance of a large dataset to obtain more accurate training and forecasts of the ANN. An observation is made by Qi (2001) in that an

ANN is more likely to outperform other methods when the input data is as current as possible and using a recursive modelling procedure. Swanson and White (1997) show that a simple feed-forward ANN with a single hidden layer offers a highly useful and flexible alternative to a linear model, particularly in multi-step ahead forecasts, as the linear model needs to be specified in advance whenever new information becomes available over an extended time period. A comparison between ANNs and an ARIMA model is given by Ghiassi, Saidane and Zimbra (2005). They find their dynamic ANN performs significantly better than a traditional ARIMA model based on MSE statistics and the Morgan-Granger-Newbold test for autocorrelation between the positive and negative sum of the error terms.

Given the array of time-series based models to choose from, the researcher must also be cognisant of the inputs to the model. Indeed, choosing appropriate inputs are as important as choosing the correct functional model form. The next section of the literature discusses a prevalent theory in finance in choosing appropriate factors that influence stock prices and returns.

2.6 Considerations in asset pricing

The Arbitrage Pricing Theory (APT) of Ross (1976) is seen as a (perhaps superior) alternative to the Capital Asset Pricing Model (CAPM) of Sharpe (1964) and Lintner (1964). The weaknesses of the CAPM are mainly in its often unrealistic assumptions and empirical shortcomings. Tests of the CAPM usually display poor explanatory power in overestimating the risk free rate and underestimating the market risk premium. This therefore limits its practicality, particularly in the use of betas to predict a share's return.

The APT has the potential to overcome these weaknesses by providing a model that generates asset returns via multiple factors and it's explanatory power can thus be theoretically better than the CAPM. Despite this, the APT has failed to replace the CAPM mainly due to its weakness to explain variation in asset returns by a given, limited number of easily identifiable factors. Indeed, most empirical tests of the APT begin with the selection of

candidate variables into the model. While this may prove useful in explaining returns, the generalisation ability of the model is poor. For example, Chen, Roll and Ross (1986) attempt to provide candidate macroeconomic variables that are felt to influence asset returns. While this selection is not based on a rigorous, theoretical identification, the approach nevertheless uses factors that are intuitive and justifiable. The aim of their study was to identify a common set of factors that will be robust across time and dataset-specific characteristics.

The APT model assumes that the return to security *i*, given by r_i, is generated by a multi-factor model.

$$r_i = b_{i,0} + b_{i,1}F_1 + b_{i,2}F_2 + \cdots + b_{i,J}F_J + e_i \quad i = 1,2,3,\ldots,N \qquad \{15\}$$

where F_i are the factors *(j = 1,2,...,J)*, $b_{i,j}$ are the factor loadings and e_i is a random variable in a universe of *N* assets. Assuming that in equilibrium all arbitrage opportunities are exhausted, the model implies that the relationship between expected return of asset *i* is given by:

$$E(r_i) = r_f + (\delta_1 - r_f)b_{i,1} + (\delta_2 - r_f)b_{i,2} + \cdots + (\delta_J - r_f)b_{i,J} \qquad \{16\}$$

where the existence of a risk free asset with return r_f is assumed and δ_j is the expected return to the portfolio with a unit sensitivity to factor *j* and a zero sensitivity to other factors. A special case of the APT, where *j = 1* and $F_1 = r_m$ is given by the CAPM equation.

To test the APT, one first needs to estimate the factor loadings for each asset and then regress the sample mean returns on the factor loadings in a cross-sectional regression. However, it is up to the researcher to determine the value of *J* (the number of factors) as well as to identify those factors. Literature makes use of either principle component analysis or factor analysis to identify the factors and estimate the factor loadings. These are then used to explain mean asset returns in the manner described above (Roll and Ross, 1980).

Alternatively, one can identify the factors *a priori* based on justifiable reasons for their inclusion. Chen, Roll and Ross (1986) used this method to identify several candidate macroeconomic factors that could affect asset returns. While the results were not compared to other models, Cheng (1996) compares a macroeconomic APT model to alternatives using canonical correlations (a method for finding the highest correlation pairs between multiple variables). He finds that the canonical correlations method works reasonably well and can successfully identify factors of economic risk in the APT context.

The estimation procedure of the APT suffers from the errors-in-variables problem, where the independent variables are often measured incorrectly, leading to a spurious regression. Gibbons (1982) suggests that a multivariate regression approach can be used to overcome this problem. Given that δ_j is the expected return to the portfolio with unit sensitivity to factor j and a null sensitivity to other factors, it implies that

$$\delta_j = b_{j,0}^p + E(F_j) \qquad \{17\}$$

where $b_{j,0}^p$ is a constant and $E(F_j)$ is the expectation of factor j. Equation (18) can be rewritten as

$$E(r_i) = r_f + \left(b_{1,0}^p + E(F_1) - r_f\right)b_{i,1} + \left(b_{2,0}^p + E(F_2) - r_f\right)b_{i,2} + \cdots + (b_{J,0}^p + E(F_J) - r_f)b_{i,J} \qquad \{18\}$$

Taking the expectation of Equation (18):

$$E(r_i) = b_{i,0} + b_{i,1}E(F_1) + b_{i,2}(F_2) + \cdots + b_{i,J}(F_J) + e_i \quad i = 1,2,3,\dots,N \qquad \{19\}$$

and combining the above two equations

$$r_f - b_{i,0} = \left(b_{1,0}^p - r_f\right)b_{i,1} + \left(b_{2,0}^p - r_f\right)b_{i,2} + \cdots + \left(b_{J,0}^p - r_f\right)b_{i,J} \quad i = 1,2,3,\dots,N \qquad \{20\}$$

which provides N restrictions on the parameters of a system of N equations. In total, there are $J+1$ parameters (the constant terms), such that there are $N-J-1$ restrictions where $J < N$. The procedure above can be estimated using non-linear least squares and is suggested by McElroy, Burmeister and Wall (1985). Having established the general equation for the APT model, attention is now focussed on particular candidate factors.

2.6.1.1 The APT model with macroeconomic factors

Chen (1983) suggested that a time series of statistically derived factors be correlated with a time series of identified macroeconomic factors. In establishing a strong statistical relationship between the two series, one can infer that the macroeconomic factors are suggestive of systematic risk. Chan, Chen and Hsieh (1985) provided a motivation for returns being sensitive to changes in the macroeconomic environment which can theoretically be hedged by investors. The authors identify the monthly growth rate in industrial production, the unanticipated inflation rate, changes in expected inflation, changes in the term structure of interest rates, the default spread and changes in the business cycle as being systematic factors that affect returns. While innovations in these factors appear a more intuitive variable to use, the authors caution that the generation of these innovations through pre-whitening[15] may lead to a loss of information. All of their chosen factors show some level of statistically significant correlation with each other and the returns on an equally weighted New York Stock Exchange (NYSE) index. In typical studies of the macroeconomic APT model, the presence of correlation between returns and candidate variables is cited as a justification for their inclusion into the return generating process (van Rensburg, 2000). The authors follow the typical APT approach outlined previously and find that the default spread, growth rate in industrial production and unanticipated inflation are significant variables in the models in explaining approximately 35% of the cross-sectional variation in expected returns. Barr (1990) applies this procedure to the JSE and finds that measures of real economic activity are appropriate factors to use in an APT model, a finding that is similar to international studies.

To further the work of Chan *et al.* (1985) and to explain expected returns of an asset by systematic risk factors, a macroeconomic APT model was found by Chen, Roll and Ross

[15] The removal of a signal at a particular frequency from a data set.

(1986). The authors utilise the framework of the APT given by Ross (1976) and focus on the influence of unanticipated events (especially in systematic or macroeconomic risk factors) on asset returns. The existence of systematic risk factors can be observed by the co-movement of share returns and in the diversification strategies of investors. This implies that factors that are associated with the economy have an impact on the prices of market indices. Chen *et al.* (1986) begin with an investigation into the dividend discount model to aid in the identification of risk factors within the APT framework. If any systematic risk factor influences either the expected cash flows or the discount rate, then it will also affect the asset's price. If one assumes that all current information is already incorporated into prices, it follows that only unanticipated information needs to be account for. To circumvent the errors-in-variables problem, the authors use a rate of change methodology to represent innovations. Using the monthly and annual industrial production growth rates, the change in expected and unexpected inflation, the change in the default spread, the term structure, consumption growth and changes in the oil price, the authors attempt to explain returns on both an equally weighted and value weighted NYSE index over the period 1953 to 1983. Amenc and Le Sourd (2005) term the final set of factors identified as the "classic" factors of a macroeconomic APT model. It can be represented as:

$$r_i = \alpha + \beta_1 MP + \beta_2 DEI + \beta_3 UI + \beta_4 UPR + \beta_5 UTS + e_i \qquad \{21\}$$

where r_i is the return on security i, α is the constant term and the factor loadings are represented by β_1 to β_5. They are respectively, the monthly industrial production growth rate (*MP*), the change in expected inflation (*DEI*), unexpected inflation (*UI*), the change in the default spread (*UPR*) and the change in the term structure (*UTS*). Chen *et al.* (1986) state that the above equation shows that returns can be modelled by innovations in multiple macroeconomic risk factors. Further, they also note that the above equation is not the only representation of a macroeconomic APT model - other candidate factors can easily be used and found to be (more) significant than the current specification.

Indeed, Hamao (1988) investigates the robustness of the Chen *et al.* (1986) model by testing it on the Japanese market. Under the APT framework, the author chooses industrial production, inflation, the equity risk premium, interest rates, the Japanese Yen/ United States

Dollar exchange rate and oil prices as factors into the APT model. Apart from finding the same factors to be significant, the change in oil prices, the unexpected change in the foreign exchange rate and changes in the terms of trade are also found to be significant. Further, the test against the CAPM beta factor shows that the CAPM does not capture any additional risks not already captured in the macroeconomic factors chosen. In other words, the inclusion of additional risk factors assists in describing the returns generating process. Apart from macroeconomic factors being considered, some authors have looked towards the arrival and assimilation of information as descriptors of stock returns.

2.6.1.2 From unanticipated factors to agent expectations

Connor and Korajczyk (1988) use principal component analysis to estimate an APT model and find that the model outperforms the CAPM. Using similar macroeconomic data by Chen *et al.* (1986), the authors apply their new technique of principal component analysis to the issue of identifying relevant factors in an APT framework. They find that the factors identified are robust to firm size and to equal or value weighted methods. Given that some statistical techniques offer little economic intuition in interpreting the estimated risk premia, attention has been given to pre-specifying observed macroeconomic and financial factors as candidates for systematic risk (most famously, Chen *et al.*, 1986). These tests rely on the assumption that prices react to news regarding macroeconomic and financial variables and that this news is unanticipated. Consequently, agents form expectations around these factors. In tests of the APT, it is therefore necessary to generate an expectations formulation process in order to examine the unanticipated components of the news. While the APT does not mention how agents form their expectations about observed factors, one possible condition that can be enforced is that the expectations produce a mean-zero and serially uncorrelated white noise process which follows a random walk.

Some authors address the issue of generating unexpected components for the observed factors of the APT and show that previous techniques employed in this area may well be misleading in identifying the appropriate set of risk factors. Two techniques emerge from the literature, namely the rate of change model and the autoregressive model. The former uses the first difference of the factor as the unanticipated component and assumes that the factors follow a

random walk where the expected value is the current value. The latter allows for the former as a special case and assumes that agents use autoregressive models to form expectations where the unanticipated component is the residual from these models. It is found that the rate of change methodology fails to meet the criteria that the components are serially uncorrelated; whereas the autoregressive methodology fails to provide an expectations generating process where the agents do not make systematic forecast errors. Chen *et al.* (1986) then propose a new methodology based on learning – the Kalman filter – which is discussed below.

2.6.1.3 Unanticipated factors in an APT

Define the i^{th} factor, $f_{i,t}$ as

$$F_{i,t} = (X_{i,t} - E_{t-1}(X_{i,t})) \qquad \{22\}$$

where $X_{i,t}$ is the actual value of the i^{th} observed factor at time t and $E_{t-1}(X_{i,t})$ is the expectation of factor X_i at time *t-1*. In this definition, the assumption that $E(F_{i,t}) = 0$ must be satisfied and the expectation forming process of $X_{i,t}$ must be considered.

If agents are rational according to the definition of Muth (1961)[16], then the unanticipated component should be a zero mean, serially uncorrelated innovation that is orthogonal (independent and uncorrelated) to the information set. This leads one to generate unanticipated factors using rational expectations models (Priestly, 1996). Indeed, Priestly (1996) shows that an APT model with Kalman filter innovations outperforms the two other methods of the rate of change and autoregressive models.

[16] According to Muth (1961), the expectations of agents tend to be distributed for the same information set about the objective probability distribution of outcomes. Thus, rational agents do not waste information, they form expectations based on the structure of the relevant system describing the economy and public opinion has no substantial effect on an agent's expectation.

Chen *et al.* (1986) generate unanticipated factors by simply differencing the candidate variables whereas Clare and Thomas (1994) use autoregressive models. The rate of change approach assumes that the unanticipated component is the first difference of the variable and that all information is included in the most recent observation. If this is true, then agents do not make use of past information when it is relevant and the unanticipated component will not be white noise. While the unanticipated components allow for the use of past information, they also assume that the parameters are stable. The author theorises that any econometric model based on the optimal decision making of economic agents over time will not capture "arbitrary", unanticipated information in its output. The alternative means of generating unanticipated components is given by Friedman (1979). It is assumed that agents use a simple linear model with time-varying parameters that will approximate the true model. Thus, agents learn and update their expectations recursively each period as more information is available such that the problem of estimating an expectations series and generating the unanticipated component becomes, in the simplest scenario, one of signal extraction through a Kalman filter.

One can represent an unanticipated shock as follows. Assume that X_t is the variable of interest, X_t^* is the expectation of X_t, shocks to X_t and X_t^* are statistically independent and that changes to X_t^* are time varying with parameter γ_t which evolves as a random walk, the model is written as

$$X_t = X_t^* + u_t \quad \{23\}$$

$$X_t^* = X_{t-1}^* + \gamma_t + \zeta_t \; where \; \gamma_t = \gamma_{t-1} + \omega_t \quad \{24\}$$

where u_t, ζ_t and γ_t are white noise processes. The first equation is known as the measurement equation and the second as the transition equation which determines the evolution of the expectation of X_t^*.

The model fits well within the framework of Ross (1989) who assumes that information evolves as a random process. The equations then describe a stochastic environment where agents form expectations based on a stochastic trend model. The residuals of the model, if

they are serially uncorrelated, are used as the expected components in the APT model. When they are serially correlated, a more general structure is allowed for by specifying an autoregressive model with time varying parameters as the expectations generating process. In the latter scenario, the measurement and transition equations will take the form

$$X_t = \delta_{i,t} X_{t-i} + \epsilon_t \quad \{25\}$$

$$\delta_{i,t} = \delta_{i,t-1} + \omega_{i,t} \quad \{26\}$$

Antoniou, Garrett and Priestley (1998) examine the APT model using data from the London Stock Exchange. The authors, using the above returns generating process, find three common factors that can explain stock returns - unexpected inflation, money supply and excess returns on the market portfolio. An interesting result emerges that two particular companies in their sample are greatly affected by an additional two factors. The authors argue that these two factors contribute marginally to the overall performance of the APT model in explaining the cross section of returns across all firms studied. Thus, the exclusion of those two companies affected by the two additional factors still provides an APT model with substantial explanatory power. Thus far, attention in the APT framework has been given to domestic factors. However, as financial markets continue to become integrated, an investor's exposure to currency risk and other international risk factors need to be taken into account.

2.6.2 International asset risk

While the macroeconomic factors of Chen *et al.* (1986) offer valid interpretations of risk in developed markets, it is also necessary to establish any additional risk factors that may be present in emerging markets. Clare and Priestley (1998) examine returns on the Malaysian market and include domestic and international factors in their APT model. Their results show that unexpected changes in the risk free rate, the term structure of interest rates, unexpected inflation and changes in expected inflation are statistically significant. The macroeconomic model is then modified with the inclusion of a domestic market index and is also found to be significant. The authors also incorporate the MSCI World Index as a proxy for international risk. Once again, this new factor is found to be significant, indicating that international risk (or at least some form of international influence) is significant in explaining returns on the

Malaysian share market. A comparison between the CAPM, the domestic APT and international APT indicates that the international APT is superior to the others, implying that an international risk factor contributes towards the pricing of assets in the Malaysian stock market.

To price assets in an international context, one needs to make assumptions on the utility functions of individuals, sources of uncertainty and market structures. Further, to model these prices requires an appropriate stochastic model specification which is robust to the arrival of new information and the possible correlation of market index movements both between countries and across time. The International APT (IAPT) (Solnik, 1983) provides a framework for finding and evaluating international factors that may influence stock returns. The author shows that that the IAPT allows investors to value the returns of the same asset differently, analogous to a domestic setting where the investors have heterogeneous beliefs. However, due to the issues around currency translation, a globally defined market portfolio and the *ex ante* specification of international factors, the IAPT has not received much attention in the more recent literature. Indeed, the author states that if international markets are segmented, the power of his own theory diminishes significantly. He advises that the APT model itself offers a better alternative to pricing international assets than the traditional CAPM and suggests that one use a combination of international factors common to all asset classes along with factors common only to domestic asset classes in deriving an international pricing model. Naturally, one needs to be aware of limiting the number of factors to ensure the resulting model can be tested and interpreted correctly. This process is discussed below.

2.6.3 Variable selection in the APT

Chen (1983) suggests that candidate variables be chosen based on which factors are justified to influence asset pricing. Thus, any factor that influences the expected cash flow or discount rate can be included in an APT model. Berry, Burmeister and McElroy (1988) set out additional criteria that must be met by candidate variables to be included in an APT model. First, each factor must have a pervasive influence on the asset's returns. Second, each factor must be unpredictable at the beginning of each period and third, relevant factors must influence expected returns. The first criterion implies that firm specific factors are not

candidate risk factors as they can be diversified away. The second criterion suggests that, at the start of every period, the expected value of the given factor is null and the series of observations is uncorrelated. Under this criterion, unexpected changes in a factor can be included as they would meet the statistical properties by definition. The final criterion can only be investigated through trial and error. The authors state that while there is no correct set of factors, an extended number of factors may produce equivalent results. While the framework by Berry *et al.* (1988) is useful in selecting factors for an APT model, it is not the focus of this thesis.

Sharpe (1963) advocates the use of a market index in the CAPM (which is a special case of the APT). From a practitioner's viewpoint, the market index is considered a benchmark by most investors when making investment decisions, suggesting that prices reflect the movements of a market index. Hamao (1988) offers that market indices capture unexpected shocks to macroeconomic factors. As prices respond quickly to public information, returns on the market index should be related to innovations in macroeconomic factors. This argument can be extended to international indices, where they are able to capture their respective domestic innovations in macroeconomic factors. Indeed, van Rensburg (1996) shows that returns on the DJIA influence the ALSI. Using variables, such as the unexpected movements in gold returns, the returns on the DJIA Industrial index, the term structure of interest rates, inflation expectations and the residual market returns factor, the author shows that all factors except gold returns are significant in the APT model constructed over the period 1980 to 1989.

Clare and Thomas (1994) suggest that changes in the expected rate of inflation can affect expected cash flows and discount rates. A higher inflation rate will imply a higher short term interest rate and a higher discount rate as investors expect to be compensated more for bearing additional risk in investing in equity. Fama (1981) argues that the negative relationship between inflation and returns is likely a result of the proxy effect (a mirror of the relationship between real activity and inflation). When using both real variables and measurements of expected and unexpected inflation, the stock return and inflation relationship disappears. The author finds evidence that real stock returns are positively related to real activity measures, such as capital expenditure, the average real rate of return on

capital and output. Further, there is a consistent negative relationship between inflation and real activity, implying that agents rationally set current prices on forecasts of relevant future real variables. Termed "the proxy effect hypothesis", Fama (1981) shows through regressions using monthly, quarterly and yearly data that the growth rates of money and real activity are effective at eliminating any perceived relationship between real stock returns and expected inflation rates. This explanation is however contested by Wei and Wong (1992) who suggest that the proxy effect can explain the spurious negative relationship between returns and expected inflation, but not between returns and unexpected inflation. From a microeconomic point of view, the authors argue that the transfer of wealth between creditors and debtors set off unanticipated inflation (as an increase in inflation helps debtors and hurts creditors). If one views the debtor and creditor firms as having listed stocks, one can expect a positive relationship between debtor firm's stocks and unanticipated inflation. Wei and Wong (1992) examine this hypothesis using common stocks from different industry groups over the period 1926 to 1985. They find that the proxy hypothesis is fully supported for natural resource stocks and partially supported for other sector stocks. By including a measure of future real activity, they find that the spurious negative relationship between stock returns and *expected* inflation is eliminated. However, the relationship between stock returns and *unexpected* inflation is still significant.

The growth rate in industrial production and the growth in GDP are often used as competing measures of real activity (Fama, 1990). Past literature shows that variables that measure time-varying expected returns shocks capture 30% of the variance on the annual real return of the value weighted NYSE index. Similarly, future growth rates of industrial production captures 43% of the variance. However, as production growth rates, expected returns and shocks to expected returns are correlated to the business cycle, the combined explanatory power of these variables is only 58%. In testing the variables that influence stock returns, Fama (1990) finds that the goodness of fit from monthly returns regressions on future production growth rates understates the information about production. This information is captured better over longer term returns.

Chen *et al.* (1986) use the default spread as one of their macroeconomic proxies. Measured as the difference between low grade corporate bonds and long term government bonds, the

authors hypothesise that the variable should have a zero mean in a risk neutral world. Thus, the default spread is a direct measure of risk aversion in pricing securities. Further, the default spread can capture a leverage effect, with more levered firms having a lower corporate bond rating. The default spread is also able to influence share returns. The spread between yields on corporate bonds and yields on government bonds increases during adverse economic conditions and would thus decrease equity returns. Thus, the default spread can have either a positive or negative effect depending on the business cycle.

Oil prices can theoretically influence the performance of the stock market. An upward movement in oil prices can arguably increase the level of uncertainty in the market, inducing a fear amongst participants and a decrease in stock prices. Studies such as Hamilton (1996) examine the influence of oil prices on industrial production, inflation and stock market returns. It is found that higher oil prices lead to a higher cost of production, a lower production rate and lower expected earnings. Huang, Masulis and Stoll (1996) examine this relationship on the U.S. stock market and find that there is a lead-lag effect between future oil prices and oil company stock returns. However, in the overall market, oil prices do not have any significant explanatory power, in line with Chen *et al.* (1986). The conclusion of whether oil prices should be considered in an APT model is still open to discussion. More recently, Miller and Ratti (2009) find that over the 1971 to 2008 period, stock markets react negatively to oil price changes in the long run. The authors use a Vector Error Correction Model (VECM) approach and find that the impact tends to reduce to zero for years after 1999. They justify this finding by stating that the oil and stock market bubble of 2000 could have influenced their results.

Given that some local industries compete with their international counterparts, the depreciation of the foreign currency foretells a loss of sales and profit from local suppliers (as consumers can now purchase the same item for a cheaper price elsewhere). Griffin and Stulz (2001) explore the impact of exchange rate fluctuations on stock prices across similar industries internationally. Using data from various countries over 1975 to 1997, Griffin and Stulz (2001) show that the impact of exchange rates varies across industries. Intuitively, exporting industries will be adversely affected by domestic currency appreciation while importing industries will benefit from it. Exchange rate fluctuations are hypothesised to affect

firm value by affecting the demand for its products and thus expected future cash flows. However, given their results, the authors state that exchange rate fluctuations alone only account for 1.5% of the variation in stock prices. By including industry-specific effects, they can account for an additional 3.8% of variation, implying that exchange rate fluctuations alone do not explain stock variation across countries.

Cutler, Poterba and Summers (1989) seek to determine the underlying causes for movements in stock prices. The (then) viewpoint was that these movements were caused solely by changes in fundamentals of the stock itself. Given various forms of event studies, literature has shown that stock prices react to announcements about a change in fundamentals. The authors then attempt to test whether the only factor in moving stock prices is the arrival of news, irrespective of whether it is fundamentally related or not. While not finding any definitive evidence on news alone influencing stock prices, their results suggest that changes in the money supply, short and long term interest rates act as proxies for economic news. Changes in the interest rate will lead to changes in the discount rate in the manner described previously. Thorbecke (1997) suggests that monetary policy tightening will decrease a firm's value and limit its ability to borrow. The lower investment expenditure will deter investors as the expected future cash flows will be lower and result in a decline in the firm's share price.

Since the Asian financial crisis of 1997, the role of gold as a hedge against economic uncertainty has again risen to considerable heights for investors. Chan and Faff (1998) examine a number of potential factors that can theoretically influence the returns of gold firms' stocks. Using the market index, gold price, interest rate and foreign exchange rate (Australian Dollar/United States Dollar), their findings suggest that changes in the gold price be used as an additional factor in an APT model as gold prices influence returns by their impact on interest rates. Davidson, Faff and Hillier (2003) test the inclusion of the gold price in an international asset pricing setting. They find that 22 global industries show sensitivities to changes in the gold price, over and above normal market fluctuations. Upon inspection of the industries, there is no discernible characteristic that defines them, implying that industries across that stock market can be influenced by the gold price. The authors conclude that the gold price is a significant factor to include in an asset pricing model. In South Africa, unpublished research by Bodington (2014) investigated the hedging properties of gold for the

South African investor investing in either local bonds, local equity or international equity. Her findings suggest that a South African investor will hedge local bonds and local equity against gold in times of a market downturn, but not against international equity. This implies that gold is not correlated to equity prices and could be included as a risk factor in an APT model. Having considered the models behind returns generating processes, one now investigates the issues around using said models to test market efficiency.

2.7 Considerations in testing market efficiency

Given a model that can explain the variation in stock returns, one then questions whether this explanation can turn into a prediction. This is a contentious issue and a simple reasoning is offered here to avoid the trap of prediction and explanation. Any empirical work conducted would rely on data gathered over a particular sample period. While the results of this empirical work can well vary, it is often difficult to simply generalise the results over the examined sample period to hold in the universal scenario. Thus, while a model might be excellent at explaining returns, it does not imply that it is also good at predicting returns. Therefore, the stance adopted in this thesis is one of providing explanatory evidence as opposed to predictive evidence.

Evidence that returns can be explained (modelled) range from literature on the aggregate, macroeconomic level (Fama and French, 1989) to the microeconomic level (Fama and French, 1992). Fama and French (1989) offer evidence on the variation of expected returns of stocks and bonds through time and across the business cycle. Their results show that the expected excess returns on bonds and stocks are correlated. The default and maturity spreads accounts for much of the variation in stock returns and are further related to long term business cycle phases. In particular, the dividend yield and default spread had a higher weighting when the business cycle is at a trough, and a lower weighting when the cycle is at its peak. The term spread is more correlated with short term changes in the business cycle and is also low around peaks and high around troughs. The authors then question whether their findings are indicative of a rational assessment of expected returns. While their results appear intuitive and in line with traditional economic theory (in that all variables relate to changing business conditions and their variation can be explained using the monetary theory of demand

and supply), the authors are hesitant to conclude that their findings imply market rationality. The work by Fama and French (1992) was a precursor to their now famed three factor APT model. Beginning with common factors identified by literature to add explanatory power to the CAPM, the authors test the impact of market equity (size of a firm) and the ratio of book equity to market equity as factors in an APT model. Offering little on the implication for market rationality, the authors simply state that their findings prove that the traditional CAPM is insufficient in explaining the risk-return relationship of stock returns. Rather, they posit that equity risk is multidimensional, influenced by the additional factors of size and book-to-market ratios.

A simple model of stock prices describes the current stock price as a function of the present value of rationally expected or optimally forecasted future dividends, discounted by a constant (or fairly stable) discount rate. This implies that movements in the returns of stocks should be attributed to movements in the forecasted dividend stream. However, some argue that stock return series are often too volatile to be explained by any new, objective information incorporated into the new dividend forecast. Shiller (1981) points out that price volatility cannot solely be explained by changes in dividends. In an attempt to reconcile the data with the efficient markets model (here assumed to be a dividend discount model described previously), the author uses time series approaches to describe the trend like behaviour of dividends and reconcile it with the chosen model. He concludes that the movements in the detrended price over the sample period can be seen as a rational response to new information about movements in the detrended dividend series, if and only if these future movements were larger than those actually observed over the data period. In other words, there are necessarily additional factors that can assist in explaining price volatility, not all of which are rational. Researchers are required to judge whether these additional factors are more consistent with rational behaviour or irrational mispricing. Some authors argue that there is a third possible explanation, that of parameter uncertainty. When investors have imperfect information about expected returns or cash flows, they must learn about the unknown process using information that is available, which can be modelled using Bayesian analysis. Parameter uncertainty will necessarily affect prices at a given point in time through its impact on investors' beliefs as well as the evolution of prices over time as investors learn more about the economy.

Lewellen and Shanken (2002) provide an example of parameter uncertainty. Suppose that dividends are independent and identically distributed over time with unknown mean, *d*, and known variance, *s*. Thus, dividends are serially uncorrelated and have constant volatility and any test of this series will reveal these properties. From a rational investor's perspective, the mean of the dividend process is random and represented by a posterior belief about *d*. Realised dividends provide information about future dividends and the perceived volatility declines as the investor learns. The empirical properties of the series are clearly different to that perceived by the investor. The authors show that for this reason, asset pricing tests can find patterns in returns that are neither part of the subjective distribution nor caused by irrationality.

2.7.1 Risk aversion, uncertainty and market efficiency

Rationality in markets implies that investors correctly use all available information in forming security prices. A consequence of this definition is that to investigate how returns are generated behoves the consideration of how market participants determine and assimilate relevant data in their decision making. The EMH assumes that investors learn to make correct inferences about the impact of new information on the probability distribution of returns, thereby forming rational expectations about the future. What the traditional definitions of rationality do not imply, however, is the speed at which security prices react to information surprises. For example, when an event clearly conveys good or bad news about a firm's future prospects, the full extent of this impact may well be uncertain. Thus, with incomplete information, the best an investor can do is estimate the parameters of a conditional probability distribution summarising various possible outcomes.

Lewellen and Shanken (2002) present an alternate market hypothesis, the Uncertain Information Hypothesis (UIH), where the price setting behaviour of investors before a dramatic financial event are known. The UIH then predicts that after new information is processed, the risk and expected return of the security in question increase in a systematic fashion; in addition to a noisy piece of favourable or unfavourable news that immediately causes a market of risk averse investors to set their prices significantly below their conditional expected values. As the uncertainty over the eventual outcome is resolved,

subsequent price changes tend to be positive on average, irrespective of the nature of the causal event. Further, if investors exhibit decreasing absolute risk aversion (the rate of change of curvature of the utility function decreases as wealth increases), then the average price change will be larger following bad news than good news. In contrast, the AMH offers no similar definition to the behaviour of investors, other than their inherent ability to learn and survive when conditions in the marketplace change. Where the UIH investigates how agents assimilate information from a microeconomic viewpoint, the AMH aggregates each agent's ability to assimilate information to create a market where the adaptive agent survives.

While no study has tested the implications of the UIH, literature indirectly does exist to provide indications of its implications. Prior to the UIH being publicised, French, Schwert and Stambaugh (1987) show that the *ex ante* risk premium on common shares is positively related to the expected volatility of returns. A positive relationship between the expected risk premium and predictable level of volatility on common stocks is found over the period 1928 to 1984. This positive relationship implies that a positive unexpected change in volatility would increase future expected risk premiums, thereby lowering current stock prices. The magnitude of this strong relationship is found to not be solely due to the leverage effect, implying that a positive relationship exists between expected risk premiums and *ex ante* volatility, in line with the implication of the UIH that there is a systematic increase in both expected risk and return.

In a pioneering study on investor behaviour, De Bondt and Thaler (1985) attempt to reconcile market behaviour and the psychology of individual decision making. They study the concept of overreaction – the tendency of prices to move past their "true" values. Conceptually, if one is willing to accept that there can be overreaction in the stock market, then it follows that some level of reaction is deemed acceptable. A means of classifying this acceptable reaction is through Bayes' rule of updating probability beliefs. However, Kahneman and Tversky (1979) show that this rule does not match the reality of how investors perceive new information and they instead use heuristics. De Bondt and Thaler (1985) demonstrate that investors tend to overreact to information and must therefore consistently revise their original forecasts. While empirically observing this phenomenon in the returns of prior high and low return portfolios, the authors provide many more questions to answer. For example, if

overreaction is consistently observed as long as (in their study) three years after portfolio formation, it begs the question of what prompts these portfolio returns to move back to their "normalised" values. One such theoretical explanation is offered by the UIH.

2.7.1.1 Developing the UIH

Assume that: investors are rational according to the standard utility axioms of von Neumann-Morgenstern; they are risk averse; the market incorporates all available information into prices quickly and that major surprises can be identified as either good or bad news, but the full extent of these surprises is uncertain. The last assumption implies that investors can form conditional probability distributions of returns given that the news is either good or bad.

Given these assumptions, Lewellen and Shanken (2002) now proceed to prove that rational investors' reactions to unfavourable news will produce a short run price pattern similar to overreaction. Conversely, the reaction to favourable news will produce a price pattern similar to underreaction. These are shown in Figure 1. Panel A shows the adjustment of prices to bad news. The arrival of bad news on the event day drives the pre-event value of the security, *P*, to *PB* and there is no response after the event. Thus, the present value of the certainty equivalent of risky cash flows is reduced to *PB* because the event discloses a certain decrease in the share's expected future cash flows.

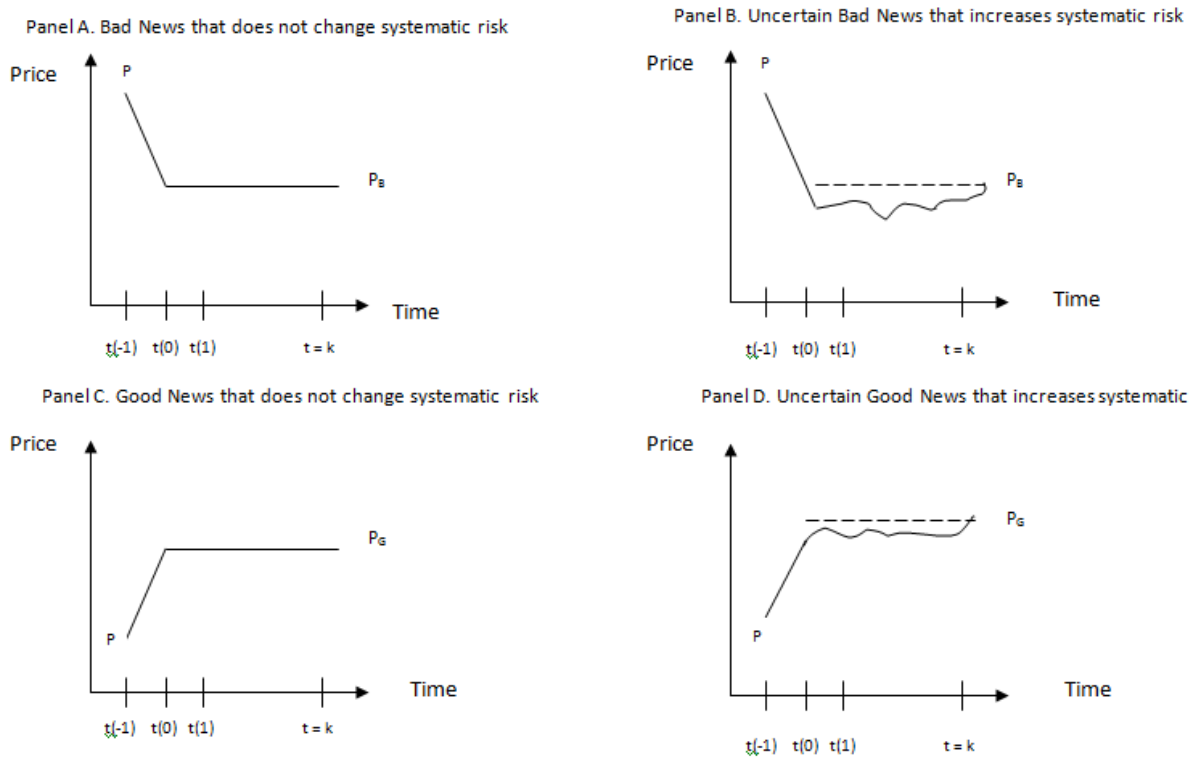


Figure 1 – Share price changes in response to favourable and unfavourable information (Lewellen and Shanken, 2002)

In contrast, Panel B shows the pattern of price changes that would be caused by unfavourable surprises that decrease the expected cash flows of the share and increase its systematic risk. With the additional uncertainty, the present value of the certainty equivalents, *PB**, is strictly less than *PB* in a market of risk averse investors. After the uncertainty of the event dissipates on day *k*, the price increases from *PB** to *PB*.

The impact of a favourable suprise is shown in Panels C and D. When the full extent of the good news is certain, the price increase from *P* to *PG*. The adjustment is instant and there is no abnormal response after the event. However, when the good news increases the systematic risk as well as the expected value of future cash flows, the price rises from *P* to *PG**. Similarly to the previous case, when the uncertainty surrounding the event dissipates on day *k*, the price further increases from *PG** to *PG*.

The above interpretation can easily be generalised to encompass marketwide surprises that affect the price of stock indices. The UIH claims that major favourable and unfavourable

surprises about the economy will increase the risk of stocks in general. Thus, an index is expected to behave to those shown in Panels B and D. Moreover, when investors experience decreasing absolute risk aversion and hold a broad, diversified portfolio of equity, the price reaction to unfavourable marketwide surprises will be more pronounced than the reaction to an equivalent favourable marketwide surprise. This implies that in both situations, the portfolio is rationally priced and there are no *ex ante* arbitrage opportunities. Further, when empirical tests are done on samples of only bad or only good news, it may create the impression of investors consistently overreacting to bad news and underreacting to good news. Thus, one needs to be aware of both data mining and parameter uncertainty. While the former may produce statistically significant results that are actually random, the latter is arguably of more concern as it might produce patterns that are statistically significant, but of no importance to the investment decision.

2.7.1.2 Testable implications of the UIH

The UIH has the following testable implications. Firstly, share return variability will increase following the announcement of any major unanticipated news. Secondly, the average price response following negative events will be positive and *vice versa*. Thirdly, on average, post event price changes will be larger for a sample of unfavourable events than favourable events if investors experience decreasing absolute risk aversion.

The second and third implication together implies that while the average reaction across good and bad news may not be the same, the reaction should not be negative. Further, when considering the first and second implication, the UIH predicts that following the arrival of unanticipated information, investors can expect to be compensated for bearing higher risk. This is line with satisfying the risk averse investor. The UIH can also be extended toward individual firm events with minor adjustments to the wealth and utility function of the investor.

While the literature on psychological biases in markets is growing, it is often misinterpreted as evidence against market efficiency. Studies in behavioural finance document anomalies

ranging from overconfidence (Barber and Odean, 2001) to herding behaviour (Seetharam and Britten, 2013). These anomalies are, in light of the argument presented above, descriptive elements of how an actual financial market functions. Shefrin (2002) mentions that profit opportunities associated with these behavioural anomalies are often associated with higher levels of risk, preserving Jensen's (1978) definition of efficiency. It is thus difficult to interpret which part of the anomaly is associated with higher levels of risk and the remaining "anomaly". Litvinova and Ou-Yang (2003) introduced the assumption that in choosing an optimal level of effort in acquiring information, agents are cognisant of the existence of other like-minded agents in the market. This creates competition amongst agents decreasing the marginal benefit of a single agent acquiring information. As the marginal benefit decreases to a level that eliminates the desire to obtain more information, sophisticated traders bear higher risk and higher costs in trading. While the number of agents may increase (or decrease), this does not necessarily lead to market efficiency and no equilibrium in their model.

Therefore, it is imperative to account for the population of traders in a market, segmented by their psychological characteristics (such as the need for competition), in obtaining costly information, as well as the speed and costs of learning such information. As such, the description of the AMH captures the behaviour of agents, albeit at an aggregated level. It provides a framework in which the individual agents are autonomous in their search for valuable information, thereby always creating an ever changing equilibrium where profits can be made from time to time. The UIH can be seen as a microeconomic view of market efficiency. Another alternate theory explores market efficiency from a macroeconomic point of view - indirectly bringing a robust understanding of the AMH.

2.7.2 The Market Fraction Hypothesis

Brock and Hommes (1998) develop an asset pricing model where the agents have heterogeneous beliefs. In this model, agents can update their beliefs of the future price of a risky asset based on a fitness measure of past realised profits. The agents can thus be

clustered according to their beliefs. Using dynamical systems analysis[17], they provide analytical evidence for circumstances in which chaos exists causing the fractions of the clusters to change. They then extend this analysis to show that chaos can exist under the two and four belief types[18]. Based on these observations, Chen (2008) and Chen, Chang and Du (2012) suggest a Market Fraction Hypothesis (MFH) which describes the constant variability among the fraction types as being driven by the types of trading strategies of the agents.

The MFH is characterised by three statements. First, in the short run, the fraction of different clusters of strategies changes over time, implying a short dominance duration for any one cluster. Second, in the long run, different clusters are equally attractive and their market fractions are equal. Third, the size of each type of trading strategy is positively correlated to its earnings performance. These statements imply that it is not possible for a single strategy to dominate the market by attracting an overwhelming fraction of market participants for many consecutive periods. Further, if the market has two trading strategies, their fraction should keep changing over time such that in the long run, they have the same market share. There is also a positive (albeit counterintuitive) correlation between survivorship of a strategy and the profits obtained from a strategy.

Agent based modelling can be categorised in a binary fashion. The first group would examine a financial market by allowing the agents to choose between different types of portfolio strategies. The agents in this model are presented with, say, three types of strategies at each point in time and are required to choose one; leading the researcher to examine the fractions of different strategies that are chosen over time. A shortcoming of this approach is that the strategies are predefined and do not change over the lifetime of the simulated experiment. As such, the second type of agent based modelling examines rather the evolution of the agents themselves, where the agent is allowed to create his own strategy at each point in time. While this may be a more realistic simulation, it focuses on the evolution of price rather than the fraction of different strategies chosen over time. Kampouridis, Chen and Tsang (2012)

[17] A geometric analysis that aims to reliably compute objects of dynamic significance, such as the swing of a pendulum.

[18] A fundamentalist or contrarian investor, coupled with a bullish or bearish investor.

combine these two approaches to test the Market Fraction Hypothesis. The authors use artificial intelligence techniques (specifically genetic programming (GP) and Self Organising Maps (SOMs), both of which are described in Chapter 3) for testing the MFH. The first technique, genetic programming, was used to evolve the agents and their behaviour over time; while the second, Self Organising Maps, was used to cluster agents with similar characteristics (thereby creating in effect fractions of agents with a similar portfolio strategy). They find that their GP algorithms produce robust results across 10 international markets, implying that in the long run, these markets tend to favour five to six types of agents to capture the behaviour of 95% of market participants. However, in using SOMs to identify clusters of agents, their results do not support the MFH particularly well. The explanation of this finding is left to their future research. Deviating from the relationship between risk aversion and parameter uncertainty, literature also provides alternative forms of market efficiency based on the rationality of the investor.

2.7.3 Rationality and Market Efficiency

Perhaps the first authors to offer an alternative to the EMH, Daniel and Titman (1999) defined a new form of market efficiency, that of adaptive efficiency. A market is considered to be adaptive efficient when profit opportunities apparent in historical data are dissipated as soon as they appear. The authors argue that a rational arbitrageur will take time to understand the trading strategies and possible irrationality of other traders. Once this knowledge is gained, patterns that were caused by trading of irrational traders can be removed over time. Assuming a risk-averse arbitrageur with limited capital, it follows that price patterns cannot be instantaneously removed, as suggested by the EMH. Thus, adaptive efficiency can be considered a weaker form of market efficiency than that suggested by the EMH. In tests of the U.S equity market, Daniel and Titman (1999) reject the adaptive efficient hypothesis as the profits discovered from a zero cost trading strategy appeared to remain persistent over the time period 1963 to 1997. The long-short strategy is based on market capitalisation, momentum and book-to-market ratios. It has negative betas each year; consistent, positive profits and an extremely high Sharpe ratio. In effect, the authors test what can now be considered the second implication of the AMH; that the profitability from following a particular strategy over time will be cyclical. While their findings are contrary to this

implication, the thinking around alternate efficient market hypotheses began from the work of Daniel and Titman (1999).

The traditional economic paradigm of rational individuals implies that these individuals can make optimal decisions based on available information. Grossman and Stiglitz (1980), *inter alia*, show that this implies that asset prices reflect all available information, such that abnormal profits can only be achieved through the use of private information. Further, the typical investor, who can reasonably be assumed to not have access to private information, would never earn abnormal profits according to this view. Daniel and Titman (1999) develop their notion of adaptive market efficiency based on the behavioural bias of overconfidence. The overconfidence bias is chosen in particular as the authors believe that it is the most established, most likely evident in security valuation and most likely to arise through evolutionary selection. The last reason, that of evolutionary selection, can be explained as follows. If a behavioural bias distorts an investment decision with no offsetting benefit, it follows that that bias would likely be eliminated through natural selection. Given the individual's ability to learn, the acknowledgment or discovery of the bias would lead the investor to determine if the bias assists in earning abnormal profits. If no abnormal profits are gained from this bias, the astute investor will discard the bias in search of another (barring any finer points on whether a character trait can easily be discarded).

Traditional economics is of the view that irrational investors have a minor effect on prices. Thus, rational investors can change prices to a point where the profit opportunity is eliminated, implying that prices are in effect determined by mostly rational investors. Behavioural-based models gained favour since the work of DeBondt and Thaler (1985) in showing that there exists overreaction in stock prices. While traditional economics criticises behavioural theories, in that the array of irrational behaviours in a given setting is unlimited, no singular theory can explain a multitude of financial anomalies out of sample. Daniel, Hirshleifer and Subrahmanyam (1998) show that the evidence is more consistent with particular behavioural biases than the standard rational model. The authors develop a theory based on investor overconfidence and biased self-attribution. In the context of financial markets, individuals will overestimate their abilities to analyse information and underestimate their error in making forecasts. Thus, an overconfident investor is defined as an agent who

attributes more confidence to his own assessments about a private signal to the market than a public signal; overweighting this information and causing the price to overreact. The behavioural model of Daniel *et al.* (1998) is based on the premise of overreaction of prices to private information and underreaction to public information. Examining the investor in the above model, the authors argue that their agents are quasi-rational; they are Bayesian optimisers for the most part, except when analysing private information.

2.7.4 The response from the rational investor

The hypothetical investor in Daniel and Titman (1999) is assumed to have shifted his capital towards strategies that have performed well in the past. The magnitude of this shift would determine the magnitude of his profits. However, without the benefit of perfect hindsight, the investor would have cautiously shifted capital, earning moderate returns over time. In the presence of irrational investors, there is no accepted metric to determine how much capital a rational investor will shift towards strategies that have performed well in the past. With both rational and irrational investors present, that learn from past price movements, non-stationarity in the data is the root cause of the problem. Whilst the hypothetical investor discovers price patterns, other rational investors would most likely discover the same price patterns at the same point in time. If all rational investors acted on their discovery, the profits from the strategy would be eliminated. The investor should ideally have a theory of inefficient markets that assists in understanding irrational behaviour as well as the extent to which these patterns are detected by other rational investors (assuming that some of the overconfident investors have access to private information). The belief of market efficiency then becomes a non-trivial question. If the investor believes that he is the only one to conduct such an analysis on the market, then he would most likely shift more capital towards the chosen strategy. If however, the investor believes that the inefficiencies detected are corrected by other investors, then he may well decide to not shift any capital to the chosen strategy. If most investors act in a like manner, then the profit opportunity may well persist. Alternatively, if they are contrarian in nature, they may well shift capital into an opposing strategy, reversing the pricing anomaly. The traditional paradigm of efficient markets, as described by Fama (1965b), implies that an investor during those times, with access to high-level technology and costless processes, would be able to earn above-average returns. The

evidence presented in Daniel and Titman (1999) firmly rejects this notion, in favour of an alternative form of adaptive efficiency.

2.7.5 The Adaptive Market Hypothesis

As per Simon (1982), human beings are boundedly rational. We cannot compute complex calculations mentally in any feasible amount of time. This is just one of the many "shortfalls" of *homo economicus*. If the EMH is the cornerstone of traditional finance theory, and if *homo economicus* is a component of the ideal investor, then the stark reality is that *homo sapiens* are not ideal investors. As such, we do not live up to the impossibly high standards of theory. This leads us to the debate of whether markets are truly efficient. From the arguments presented previously, the time horizon under investigation is extremely important. It is plausible that no practitioner may beat the market in the long term. It is also equally possible that an analyst might consistently beat the market in the short term (in independent trades). This thesis defines efficiency for markets to be such that the market is efficient at every point in time.

Lo (2004, 2005) describes a new form of market theory – the Adaptive Market Hypothesis (AMH). This approach utilises concepts from finance and the principles of evolution. It is simply stated as follows: "Prices reflect as much information as dictated by the combination of environmental conditions and the number and nature of 'species' in the economy" (Lo, 2005, p. 19). Species refer to market participants (asset managers, hedge funds, traders, *inter alia*). Thus, the efficiency of the market at any point in time is related to the factors of evolution and competition present.

The AMH is built upon Wilson's (1975) concept of socio-biology and Simon's (1982) concept of bounded rationality. As decision makers learn through trial and error, the feedback from these actions determines their survival capability. As market conditions change, participants develop new heuristics to replace the old, inappropriate ones and adjust their investment strategy accordingly. Research by Hunt and Ellis (1999) shows that emotion affects people's memory and judgement. Indeed, when making investment decisions under

uncertain conditions, it is reasonable to expect that the investor will deviate from full rationality and modify his investment strategy based on past mistakes and successes. A form of Darwinian evolution ensues in that only the investor who adapts is able to survive by making profits under changing conditions. Conceptually, the EMH can be considered a final state model which is fixed whereas the AMH is considered a dynamic model that reaches the fixed state of the EMH. It presents a simple, philosophical and pleasantly intuitive view of market efficiency. Market efficiency can be seen as cyclical. There are times of inefficiency and efficiency. For a market to become efficient, it must first be inefficient and *vice versa*. The influence of market participants (through trading or financial product innovation) influences this efficiency, sometimes in a disruptive way. To date, no formal methodology has been published on testing the AMH. However, authors have nonetheless proposed and tested methods.

The AMH of Lo (2004, 2005) captures the characteristics of the changing psychology of different investor groups. It applies evolutionary principles to financial markets, and attempts to explain investor "irrationality" as a rational reaction to a changing environment. Further, the AMH implies that market efficiency is relative to time. In other words, markets can be both efficient and inefficient over a sample period as efficiency is measured at each time interval. The AMH incorporates elements of asymmetric information of Grossman and Stiglitz (1980) and the "noise" trader[19] of Black (1986). In ecological terms, they are the prey of sophisticated traders and suffer from psychological biases as described by behavioural finance. Statistically, their behaviour may cause a serial dependence in price changes unless eliminated quickly by sophisticated traders. In addition to these two investor groups, there is a feedback loop between them. According to Shefrin (2002), sophisticated traders may be aware of an increased number of noise traders through large and sudden price changes, which would increase their perception of risk and decrease their enthusiasm for trade. This would in turn lead to larger price fluctuations, changing the equilibrium between noise and sophisticated traders over time. One can thus expect to find larger price fluctuations during periods dominated by noise traders.

[19] An arguable definition, a noise trader is an investor who makes buy and sell decisions with no regard to fundamental analysis.

If one had more precise information on the number and behaviour of noise traders, one could model an ecological environment and possibly the complexity of competition for resources. From the suggestions of Farmer and Lo (1999), this would enable insight into the emergence and extinction of certain investor groups and behaviour. The inclusion of heterogeneous beliefs in an adaptive learning model assists in depicting financial market dynamics, which according to Lo (2004, 2005), provides an opportunity to: explain changes in risk premia, explain changes in risk attitude and, explain changes in winning investment strategies; ultimately aiding in understanding the process of market efficiency.

For example, Todea, Ulici and Silaghi (2009) test the profitability of a moving average trading rule in six Asian markets. By examining a technical analysis trading rule, the authors are in effect testing the weak form of market efficiency. The authors state that an acceptance (failure to reject) the hypothesis implies informational efficiency, but a rejection of the hypothesis does not imply informational inefficiency due to the joint hypothesis problem described earlier. Examining the evolution of profits from a trading strategy, they find that the profits generated vary through time and postulate that this cyclicality is similar to that described by the AMH. While this may seem like weak evidence in favour of the *AMH*, recall that one cannot readily test the *EMH* without first specifying an appropriate equilibrium price model. Thus, one can only offer indirect evidence of the implication of the AMH.

Neely, Weller and Ulrich (2009) test for stability of returns over time in the foreign exchange market. The authors find that trading rules used during the 1970s and 1980s provided statistically significant profits whereas those used during the 1990s did not. They infer that the lack of consistent profits over time implies a cyclical form of efficiency, favouring the AMH, similar to Todea *et al.* (2009). Butler and Kazakov (2012) test two implications (cyclical profitability and cyclical efficiency) of the AMH using computational intelligence techniques. To test the former, the authors use a popular trading rule, Bollinger Bands[20], to determine profits from following that rule over time. The rule is adapted using a particular AI

[20] A technical analysis trading rule that measures the high and low values of a share's price relative to previous trades.

technique, that of Particle Swarm Optimisation (PSO)[21] to chose optimal parameters in the Bollinger Bands. They find that this particular rule is able to outperform the market index 35% of the time, implying that the profits from such a strategy vary over time in line with the implications of the AMH. Further, to test the latter implication of cyclical efficiency, the authors examine the returns generating process. Assuming returns to be generated from a GARCH(1,1) model, the authors divide the output into a sample exhibiting random walk behaviour and another exhibiting deterministic (non-linear) behaviour. This division is done using the Hinich Portmanteau bi-correlation test. Thereafter, using several AI techniques to determine predictability of the two samples, the authors find that a Support Vector Machine[22] or a decision tree[23] are equally effective in forecasting future values of deterministic share returns. Using time-series econometrics, the authors demonstrate that non-linear dependence, if detected, can provide more reliable forecasts.

Cajueiro and Tabak (2004) test for market efficiency across 13 different countries. The authors use a rolling window approach to view efficiency over time rather than across the entire sample period in comparison to most other studies. Using a Hurst exponent and Rescaled Variance (R/S) method, the authors calculate and rank relative efficiency across the 13 market indices. They find that most Asian indices studied exhibit long run dependencies compared to South American indices. Further, in line with previous studies, there is no evidence of dependence in return observations within the developed market indices studied. The authors present a practical framework for testing serial dependence, which is adopted in this thesis. These tests are apt for detecting long term patterns that may be apparent in the data. Some authors argue that the existence of non-linear serial dependence is a challenge to the unpredictability criterion of market efficiency. In principle, when one examines the autocorrelation between return observations, one should also examine higher order (non-linear) autocorrelation that may be present.

[21] A method of optimisation that iteratively improves the candidate solution allowing a population of candidate solutions to converge on a particular option.

[22] A supervised learning model that can analyse data and recognise patterns in either classification problems or regression analysis.

[23] A support tool that represents decisions and possible outcomes in a tree-like graph.

Johnson, Jefferies and Hui (2003) demonstrate that a simple return generating process may exhibit higher order correlation which appears random when one examines the (linear) autocorrelation function. Specifically, the function produces a value of zero, indicating that the returns are not correlated, and thus implying evidence in favour of the EMH. Lim (2007) takes the above into consideration in testing for relative efficiency across market indices. Adopting the approach from Cajueiro and Tabak (2004), the author uses a Hinich Portmanteau test with a rolling window approach to capture higher order correlations across time. He finds that market efficiency is not a static measure as previously assumed by the literature. Instead, there is evidence of non-linear dependence of stock returns that evolves over time.

Indeed, a branch of literature exists which uses theories and techniques from the physical science discipline in solving problems in economics. Labelled *econophysics*, the fledgling field can provide means of solving some of the long standing questions in finance. Zunino, Zanin, Tabak, Prez and Rosso (2009) test whether market efficiency is cyclical in developed and emerging markets using techniques from the physical sciences. They introduce the concept of entropy, a measure of disorder or chaos, to rank stock market efficiency. This measure does not rely on any particular pricing model, but does rely on the probability distribution of prices. If stock prices followed a random walk, this entropy measure would be maximised. A further measure, the number of forbidden patterns, is also used. These forbidden patterns capture the existence of missing sequences in a time series and was proposed by Amigo, Kocarev and Szczepanski (2006) as a distinguishing factor between a random and deterministic process. Zunino *et al.* (2009) find that both measures have lower values in developed markets (indicating greater efficiency) and higher values in emerging markets. The results of these three particular studies demonstrate that the paradigm of considering efficiency as a binary state is changing, along with the techniques used to solve the long standing debate on market efficiency. The methodology used in this thesis adopts some of these approaches, discussed in detail in Chapter 4.

2.7.6 The behavioural view of market efficiency

The field of economics can be seen as having foundations in biology, sociology and more recently psychology. While current economic thought is now dominated by equilibrium-based models, some authors have nonetheless proposed models based on biological processes. Miller (1986) argues that these models place limited demands on the abilities of the economic agent. Borrowing from Simon (1982), agents can be considered boundedly rational. In other words, they can process as much information as is humanly possible. This sets a far less restrictive assumption on the capability of the agent compared to optimisation models traditionally used. Further, these biological models are dynamic in nature, making them well-equipped to handle disequilibrium conditions. In contrast to equilibrium models, these evolutionary models can model a large amount of economic behaviour which (arguably) occurs in disequilibrium states.

While the biological-based model is appealing, it is not without its disadvantages. These models often lack analytical solutions that traditional models provide. Their dynamic nature often requires simulation to achieve a high level of accuracy. The results from these evolutionary models can provide insight into the conditions that an optimal state is reached, prompting further research using equilibrium models. Miller (1986) creates a biological-based model on genetic programming. The model has sufficient theoretical structure and works well in explaining many economic concepts. Further, Miller (1986) suggests that the model has strong optimisation abilities, inferring that the optimisation and adaptive approach may not be mutually exclusive, as previously thought. Indeed the role of Artificial Intelligence, specifically neural networks as approximators of functions, has aided to solve many complex issues in economics.

2.7.6.1 Neural Networks

Any form of explanatory analysis on share returns makes the implicit assumption that publicly available information has a relationship to future share returns. Such information could range from economic variables, fundamental (accounting-based) variables to rumour and speculation. This assumption clearly violates the EMH which states that it is impossible

to forecast future prices as all relevant information is already accounted for in current market prices. When new information enters the market, prices will adjust instantaneously in a random manner according to the random walk hypothesis. This line of reasoning implies that the best forecast of future share prices is the current share price, thus resulting in a random walk model. A major caveat of studies that show the contrary, which is exposed by proponents of the EMH, is that the evidence presented relies on a linear dependence between the share price and the independent variables. Practically, it is reasonable to infer that non-linear relationships do exist between economic and financial variables. Given this inference, one can then proceed to model these relationships. However, this model-driven approach requires that the model first be specified before estimation of the parameters can commence. Neural networks have thus been introduced to model financial problems precisely because of the reason outlined above. They are capable of non-linear modelling without any *a priori* knowledge about the relationship between the input and output variables. Desai and Bharati (1998) test the predictability of four asset class returns using a neural network. If a neural network is mistakenly applied to linear data, the network will either be relatively computationally expensive to train compared to simpler linear models or will overfit the data and learn the noise in the series. Thus, to avoid the latter, one should first investigate the series for neglected non-linearity before attempting to use a neural network to predict any future values. Using two popular tests of the sort, Desai and Bharati (1998) test the return series for large stocks, small stocks, corporate bonds and government bonds. They find that non-linearities do exist in large stocks and corporate bonds and attempt to fit a neural network to predict future values of these two asset class returns. The neural network outperformed both a linear regression and GARCH model, showing that over the sample period covered, neural networks are more suitable for modelling non-linear behaviour of asset classes.

Notwithstanding their ability to perform non-linear modelling, the accuracy of results from a neural network is heavily biased towards the ability of the researcher. In other words, a neural network is only as successful at predicting future prices based on the inputs received (which are selected by the researcher). Often, no justification is given for the selection criteria of input variables. It is apparent that the inclusion (or exclusion) of (irrelevant) relevant input variables can be detrimental to the success of the network. Given some background on neural networks, one proceeds to be informed of their development since inception.

Hardin (2002) suggests that neural networks were developed in part due to the ordinal revolution in economics and decision theory. As the choices of economic agents have a social and interactive context, one needs to construct a means of mapping all potential and actual responses from the interaction of these agents. As our choices have social and interactive elements, it becomes near impossible to theoretically describe all potential paths from these responses and interactions. Thus, Hardin (2002) argues that these models exhibit a fundamental indeterminacy. It is impossible to practically describe all possible interactions and responses. Assuming the rational individual understands the product of their interactions with others, it follows that the reactions of those other participants may not necessarily be similar or unique – similar to the "prisoner's dilemma". When the element of time is added to these models, the agents may react quite differently to what was assumed by the other agents. These time based models would be dependent upon some initial condition, which would cause a chaotic series of actions to emanate from each change in condition. Further assuming that a complex model can be constructed and empirically tested, the problem of aggregation arises, where information contained in the individual data are lost due to aggregation. This can be observed through application of Arrow's Impossibility Theorem (Arrow, 1950). Aggregation of preferences into a general choice rule makes it impossible to determine the optimal allocation of resources in the face of disagreement[24].

ANNs assume ambiguity[25] in the ability of the researcher – the researcher does not know that he is incapable of conceiving, designing or constructing a complicated, interactive model of human behaviour. Thus, the alternative would be to learn from past observations, without imposing a determinate principle on it (Krippendorf, 2002). In economics and finance, this does not necessarily pose a problem as initial conditions are dependent on future expectations - the price of a stock today does not necessarily depend solely on its previous price, but also on the forces of supply and demand for it. Applying an ANN to an economic or financial problem, the focus would be to detect and test for non-linear relationships as they are more likely to be present than linear relationships according to Granger (1991). While neural

[24] The theorem is most commonly described in an example of an election. Assume a finite set of candidates for the election, a finite set of voters and their individual preferences for outcomes. These preferences are unrestricted - they are independent of other influences. The theorem states that it is not possible to derive a complete and consistent social choice rule exclusively from the individual preferences, except in dictatorships.

[25] Ambiguity is defined as unknown outcomes with an unknown distribution.

networks are capable of processing voluminous amounts of data, they lack insightful imagination (Weiss and Kulikowski, 1991). In other words, while they are capable of processing voluminous data and performing calculations beyond the natural ability of humans, the results of ANNs are essentially taken to be true, provided the data and the network itself is adequate.

2.7.6.2 Advantages and Disadvantages of ANNs

The use of ANNs over conventional statistical methods presents many useful advantages. First, ANNs have the ability to analyse complex patterns in the data with a high degree of accuracy. Second, there are no assumptions made as to the underlying distribution of the data. They thus provide an unbiased analysis, especially when the relationships between variables do not fit an assumed model. Maasoumi, Khotanzed and Abaye (1994) stress that since time-series data is dynamic, it is necessary to have non-linear tools to discover any relationships among the data. They conclude that ANNs are the best at discovering such relationships. Given that not all data sets are complete, ANNs can perform well with missing or incomplete data. The ANN can readjust its connection weights to account for the new data presented to it, enabling a dynamic updating of the node thresholds and providing a more accurate forecast. In comparison to an econometric model, it is easier for an ANN to forecast data over short intervals, given that the argument of anomalous characteristics disappearing when data is aggregated. If the ANN is used for solving an economic or financial problem, this advantage is quite appealing. In an attempt to circumvent data aggregation, data of differing frequencies is thus used in this study to provide robustness.

Given the complex nature of economic and financial systems, it is difficult (if not impossible) to develop a model which accounts for all possible reactions and counter-reactions. If one tries to account for all possible outcomes and dynamic interactions, the resulting model becomes both overly complex and impractical to test. Thus, using principles such as profit or utility maximisation produce inaccurate results. Recall that the most important maxim in the AMH is that of survival - not necessarily of profit or utility maximisation. ANNs, while not attempting to provide a complete model of the system, attempt to emulate it. ANNs can handle the indeterminacy of the system by either utilising probability and statistics; or by

using fuzzy logic on the input and output data. The activation function can thus be adjusted accordingly. While the ANN does not solve the indeterminacy problem, it provides a means of reducing it; thereby allowing forecasts and predictions to be carried out with some degree of accuracy (often higher than traditional econometric methods). This is indirectly tested by using differing frequency data in this study as well comparing the result of the neural network to more traditional econometric models.

As much as an ANN solves many problems, there are also flaws in utilising them. Firstly, ANN development is often left to the researcher in that there is no structured methodology available for constructing an ANN. Further, the output quality may be unpredictable regardless of the architecture of the network. The researcher may have followed each reasonable heuristic in designing an optimal network, but the output may nonetheless be poor. An ANN is also considered to be a "black box" (a system that cannot be fully described despite it accurately predicting output data). As such, it is impossible to determine the relationship between nodes in the hidden layer without further additions (Li, 1994). One such method that has emerged in the recent literature is the use of a Deterministic Finite State Automaton (DFA). These DFAs produce a symbolic representation of how input data is processed and transformed into output data, however, this method is still in its infancy as of time of writing.

ANNs can be thought of as autopoietic systems – they produce their own patterns from a set of inputs that are, in turn, used to operate the future production of outputs which are emulative (provide empirical evidence) instead of theoretical (provide theoretical evidence) (Krippendorf, 2002). In contrast, a regression model is usually built around first principles in statistics and physics. Thus, a regression model provides a higher level of structure and explanatory power compared to an ANN. As such, it is important to understand the different types of ANN architecture to determine if the advantages and disadvantages of each are acceptable to the researcher. ANNs usually have long training times that require the researcher to perform multiple iterations to enhance confidence in the predictive ability of the network. Another disadvantage is that neural networks are data-dependent. The success of the ANN depends on the input data. In solving financial problems, it is crucial to test the ANN

on out-of-sample data as the input data may be inherently different (similar) to the out-of-sample data.

Kanas (2001) developed an ANN for the Dow Jones and Financial Times indices. Tests from both the ANN and a linear model revealed little in predicting directional changes in the indices, however, the non-linearity in share prices was confirmed. An ANN may easily over-fit or under-fit the data, an implication from an indeterminate system. Therefore, an ANN does not contain explicit causal relationships nor is built on first principles.

It should be noted that many of the disadvantages highlighted above can be solved using pre-processed data (for example, using returns instead of raw price levels). As per Schwartz (1995), using a few well-chosen variables will result in a better result than using every known economic variable as inputs. Often, the ES can be used to eliminate either insignificant or highly correlated variables – speeding up the training time and enhancing accuracy. This also adds an element of indeterminacy – the choice of the ES differs each time based on the iteration and choice of expert. More such applications are now discussed, with particular reference to the use of neural networks as opposed to other AI techniques.

2.7.6.3 Application of ANNs to finance

Swales and Yoon (1992) test whether an ANN is better at forecasting than multiple discriminant analysis. Given the popularity of the former technique, the limitations of the technique suggest that a non-linear approach may better assist analysts and investors in making investment decisions. The authors show that an ANN is superior at predicting share prices compared to the discriminant analysis method, based on analysing information content in news alerts from select Fortune 500 companies.

In the insurance arena, Brockett, Cooper, Golden and Pitaktong (1994) construct an early warning system to predict insolvency on insured clients. The authors use a feed-forward neural network with the backpropagation learning algorithm and compare its performance

against the more traditional measures in the field to predict insolvency, namely discriminant analysis and publically reported insurance regulator ratings. They find that the neural network shows a high level of predictability and generalisation for predicting insolvency two years after the end of their sample period. While it is now known that a feedforward network is not the best architecture to use for time series data, the results of the authors show the power of artificial intelligence techniques in solving relevant issues in finance (or at least in insurance).

In the pricing of derivatives, the most common practice is to use the Black-Scholes option pricing framework. However, this approach rests on the parametric specification of the dynamics of the underlying asset's price. If there is a misspecification in this stochastic return generating process, then it follows that the price derived from the framework will be error prone. In effect, the success of establishing a true price of the derivative rests on correctly specifying the stochastic process of the underlying asset price. Hutchinson, Lo and Poggio (1994) propose a non-parametric approach for pricing derivatives. By selecting those factors believed to influence the derivative's price, the authors compare the error terms from three different models, a radial basis function network, a multilayer perceptron and a projection pursuit regression (PPR) (a technique unrelated to artificial intelligence, PPR is a means of analysing high dimensional datasets by examining their lower dimension projections). While the authors do report that the networks are better than other methods, they are hesitant to generalise their findings given the short data period (three years) and single derivative instrument used.

While the ANN is better at prediction, it does not imply that the ANN is a determinate system. Some authors, such as Hill, Marquez, O'Connor and Remus (1994) find that the ANN is comparable to traditional statistical methods. Indeed, the ANN performs as well as the classical regression model at forecasting yearly prices, but better in forecasting monthly and quarterly prices. When non-linearity is present in the data, the ANN can necessarily outperform regressions in modelling human behaviour. Kuo and Reitsch (1996) test regression and ANN methods at forecasting data. They use two datasets, one with a dependent variable and a number of explanatory variables (a cross sectional dataset) and the other with a single dependent variable measured across time (a time series dataset). Further

employing exponential smoothing techniques to the time series data, the authors find the neural network models generated the most accurate forecasts in both datasets.

Kuan and Liu (1995) investigate the out-of-sample forecasting ability of neural networks in predicting exchange rates. As foreign exchange rates are integrated of order one and their changes are uncorrelated over time, these changes are not linearly predictable. Thus, one needs to employ non-linear methods to forecast them. Utilising a two step procedure to estimate and select the appropriate feed-forward and recurrent network, the results from their study are mixed. Out of six daily exchange rates studied (the U.S Dollar, British Pound, Canadian Dollar, Deutsche Mark, Japanese Yen and Swiss Franc) over 1980 to 1985, only two networks offer either significant market timing ability (predicting the correct direction of the future exchange rate) or a lower out of sample error. While their results are not overall in favour of using neural networks to forecast exchange rates, the authors do propose an easily implemented procedure in selecting the best network for use in the modelling exercise. The procedure allows for a family of networks to be estimated that produce the best predictive ability. Thereafter, statistically better estimates for these networks are derived using non-linear least squares regressions. The authors test this procedure and find that it performs well in determining the optimal ANN.

Shachmurove and Witkowska (2001) investigate the dynamic relationships between major world stock markets using neural networks. Using daily data from seven major indices (six country indices and one world index), the authors propose that the daily return on a particular index is a function (contemporaneous and lagged) of other indices. They first apply ordinary least squares regression methods to determine which variables are significant to be input to the neural network, a multilayer perceptron. They find that the neural network predicts daily stock returns better than the more traditional methods of ordinary least squares and general linear regression models. Further, there are different network architectures that exist for each index. The results of their study point towards a simple, yet powerful application of neural networks in predicting stock returns. In the case of the authors, their objective was to determine if there exist interrelations between global stock indices and to determine if a non-parametric model provided superior forecasting ability. Indeed, asset managers and investment banks such as Goldman Sachs and J.P. Morgan utilise ANNs (Shachmurove and

Witkowska, 2001). The authors describe how a unit trust by Fidelity Investments bases its portfolio allocation on the recommendations of its ANN. The increased usage of ANNs in business indicates the usefulness of the ANN in solving financial problems and can be considered a pioneering field in the realm of empirical finance.

2.8 Frontiers in finance

There has been a recent emergence of non-traditional fields in finance - most notably that of behavioural finance, evolutionary finance and neurofinance (Tseng, 2006). The extent of market efficiency and indeed its participants' rationality is a matter of perspective. While the empirical evidence remains unchanged, the unique characteristics of each individual that views this information can possibly lead to differing conclusions. This is more apparent in the professional realms of trading and investing, where market participants have a variety of backgrounds, experience and heuristics for analysing identical information.

Rational behaviour theories either prescribe how people should behave in order to achieve certain goals under certain conditions, or they describe how people actually do behave. When risk, uncertainty or incomplete information is introduced, it is well documented that people behave differently from the strict and often abstract definition of rationality. Simon (1997) defines an alternative (and more realistic) form of rationality, which he calls bounded rationality.

> The term 'bounded rationality' is used to designate rational choice that takes into account the cognitive limitations of the decision-maker, limitations of both knowledge and computational capacity. Bounded rationality is a central theme in the behavioural approach to economics, which is deeply concerned with the ways in which the actual decision-making process influences the decisions that are reached. The theory of subjective expected utility (SEU theory) underlying neo-classical economics postulates that choices are made: (1) among a given, fixed set of alternatives; (2) with (subjectively) known probability distributions of outcomes for each; and (3) in such a way as to maximize the expected value of a given utility function (Savage, 1954). These are convenient assumptions, providing the basis for a very rich and elegant body of theory, but

> they are assumptions that may not fit empirically the situations of economic choice in which we are interested. Simon (1997, p. 291)

The concept of bounded rationality has been firmly rooted in many theories in behavioural finance and is preferred over its stricter counterpart. Indeed, Shleifer (2000) argues that the attitude of investors towards risk, their sensitivity to the framing of problems and their non-Bayesian expectation formations bias investors toward deviating from rationality. The first assumption of the EMH can thus be modified in terms of bounded rationality or minimal rationality (Rubinstein, 2001).

Rubinstein (2001) describes the debate between himself and a famed behaviouralist, Richard Thaler, on market rationality. The first assumption is that markets are maximally rational if all investors are rational. This implies that investors would not trade much and rather invest in the market or an index fund. In practice, the author argues that this is hardly believed to be true by many investors. The second assumption is that asset prices are determined as if all investors are rational. Again, in practice, it can be said that not all investors are actually rational, for if they were, then, for example, fund managers would correct their own and their client's irrational investment choices. Therefore, in a rational, but not maximally rational market, investors can either trade too much or too little. If markets are not rational, it does not imply that profit opportunities exist. In such a case, Rubinstein (2001) refers to this as a minimally rational market, where prices are not set as if all investors are rational, but there are no abnormal profit opportunities for the rational investor. Further, Shleifer (2000) shows that real world arbitrage opportunities are risky and limited. This would imply that if an opportunity arose, it may not necessarily be eliminated through trading as that action is dependent upon the risk attitude of the investor willing to undertake the arbitrage. The final assumption is also considered unrealistic by Simon (1982). It is assumed that an investor will comprehensively and accurately analyse information available to decide upon choices in the present and the future. In reality, this assumption cannot be met by market participants, due to the limitations of their mental capacities.

As much as the empirical evidence can criticise the EMH, supporters of the EMH show that the methodology used in testing market efficiency are the very cause of such anomalies as the methods are designed to detect an anomaly they create (Fama, 1998). Conceptually, this point is valid, yet in a similar manner, there has been no correct methodology used to show that the market has been efficient at all times. An interesting viewpoint was raised by Constantinides (2002) in that "several examples of apparent deviation from rationality may be reconciled with rational economic paradigms, once we recognise that rational investors have incomplete knowledge of the fundamental structure of the economy and engage in learning". This finding was taken into consideration in developing the AMH where agents are boundedly rational and engage in learning in order to generate profits and survive over changing market conditions.

2.8.1 Bounded Rationality

According to Simon (1982), the assumptions of SEUT are that: (1) the utility function is well defined and cardinal, (2) there is a well defined set of alternatives, (3) a joint probability distribution can be assigned to all future sets of events and (4) the decision maker is a utility maximiser. The theory of bounded rationality aims to relax the assumptions of SEUT. It is important to note that rationality cannot be considered in binary form – there are varying degrees of rationality. Thus, bounded rationality does not imply irrationality on the part of the investor but merely a less strict form of perfect rationality (as discussed above).

Relaxing the second assumption, one can assume that the alternatives will follow a generating process. This generating process can be considered complex and difficult to analyse in a given amount of time, as outside factors may affect asset prices which will in turn cause the process to change again. Therefore, it is unlikely that the given set of alternatives assumed by SEUT will be complete. As investors find alternatives, evaluate them and decide which to follow, the given time period in which the alternatives remain fixed is too small to provide an accurate assessment. Modern cognitive psychology shows that in these situations, humans will follow some heuristic in finding satisfactory answers instead of perfect answers.

Similarly, the third assumption requires the investor to have *a priori* knowledge of all future events – an impossible task. Instead, the investor will rely on estimates of joint probability distributions given that future events are uncertain. If both the second and third assumptions are relaxed, the decision maker is unlikely to have a well defined utility function (Tseng, 2006). It follows that the last assumption will also be relaxed as a utility function that is not well defined cannot be maximised. Thus, the decision maker will have no alternative but to settle for a type of satisficing strategy (Simon, 1982). In other words, the decision maker will settle for a decision regardless of whether it is the most optimal one.

Conlisk (1996) provides four reasons for incorporating bounded rationality into traditional finance and economic theory. First, bounded rationality provides empirical insights.

> There is a mountain of experiments in which people: display intransitivity; misunderstand statistical independence; mistake random data for patterned data and vice versa; fail to appreciate law of large number effects; fail to recognize statistical dominance; make errors in updating probabilities on the basis of new information; understate the significance of given sample size; fail to understand covariation for even the simplest 2x2 contingency tables; make false inferences about causality; ignore relevant information; use irrelevant information (as in sunk cost fallacies); exaggerate the importance of vivid over pallid evidence; exaggerate the importance of fallible predictors; exaggerate the ex ante probability of a random event which has already occurred; display overconfidence in judgment over evidence; exaggerate confirming over disconfirming evidence relative to initial beliefs; give answers that are highly sensitive to logically irrelevant changes in questions; do redundant and ambiguous tests to confirm a hypothesis at the expense of decisive tests to disconfirm; make frequent errors in deductive reasoning tasks such as syllogisms; place higher value on an opportunity if an experimenter rigs it to be the "status quo" opportunity; fail to discount the future consistently; fail to adjust repeated choices to accommodate intertemporal connections; and more. (Conlisk, 1996, p. 670).

Second, there are economic and financial models that already incorporate bounded rationality, which were subsequently proven to be more useful than their counterparts. Anufriev, Hommes and Philipse (2013) examine the influence of expectations on market

prices as past literature shows that heterogeneous agent models and non-fundamental expectations can result in price bubbles. As the expectations of market participants in real markets are not easily observable, a controlled experiment is particularly difficult to construct. Thus, the authors fit a Heuristics Switching Model to determine if agents can learn to forecast. In this model, agents can switch between heuristics by learning which heuristic performed better in the past, causing the impact of different heuristics on price to change over time. Their results showed that the participants relied on simple first order forecasting heuristics and anchored their expectations to past prices and extrapolated past trends. Third, there may be a case where the environmental conditions favour either bounded or unbounded (maximal) rationality. Last, based on a foundation of economics, limitations on cognitive abilities can be considered a scarce resource.

Gabaix and Laibson (2000) develop a boundedly rational decision algorithm which makes quantitative predictions on heuristics. Based on algorithms that simplify decision trees, the authors show that when cognitive efforts are costly, the agent will rely on a simplification of the decision tree. The model proposed makes quantitative behavioural predictions, offering an alternative to rational models that is psychologically plausible. The decision algorithms are widely used and documented in the psychological literature, showing that the model is empirically testable. The authors find that the model fits well and reject the notion of a fully rational model.

Given a choice of receiving R120 today and receiving 12 monthly payments of R10, it is easy to choose the optimal option rationally. The choice above requires both knowledge of the problem (financial mathematics in this example) and computation of the answers in both choices. Thus, rationality requires the decision maker to have knowledge and computational skills.

As a contrasting example, suppose that a salesman needs to visit two customers that are situated at opposite ends of the neighbourhood. In choosing which customer to visit first, the salesman needs to have knowledge of the costs, time and distance to travel. Further assumptions will also need to be made, such as the availability of each customer. With only

two customers, the optimal solution is relatively straightforward given all the required information. However, if the number of customers increases linearly, the time taken in solving the problem increases exponentially. This is known as the "travelling salesman" problem in operations research and computer science. While heuristics exist to solve the problem, they are computationally difficult to implement. The problem is classified as non-deterministic polynomial-time hard (NP-hard). In other words, the time required to find the optimal solution grows exponentially as the number of customers increases linearly (Tsang, 2008). Again, while a solution exists, it may not be feasible for the salesman to compute within a fixed amount of time. In such an event, the salesman would settle for the second best option found within the allotted time frame as well as the lowest cost (where knowledge gain and computational time are costly). This provides a backdrop for the definition of bounded rationality of Simon (1982).

From a scientific viewpoint, one can study the effect of relaxing the assumptions of SEUT using algorithms and heuristics (Tsang, 2008). In a financial market, often the most feasible option available to the trader or investor is to rely on heuristics to find the optimal solution to a problem. The search for such a solution not only provides an interesting study in itself, but also has implications for market efficiency. If one defines perfect rationality as being able to find the optimal solution in a given situation, then the level of optimality settled upon by constrained resources determines the level of rationality. As technology advances and general living conditions (in particular, education levels) increase, the decision maker is able to implement what was a once computationally infeasible solution to current problems. Thus, a theory where computational intelligence determines effective rationality (CIDER) is introduced by Tsang (2008). Rubinstein (2001) states that real world agents do not necessarily attempt to find the optimal decision. If it is assumed that an agent's task is to pick that decision from a finite set of options that satisfies all given constraints, then an agent's actions will be impacted by the actions of others in say, an organisation or a market. In such an environment, finding that decision which satisfies all constraints is often easier than finding the optimal one (Tsang, 2008). This problem of constraint satisfaction is also bounded by the problem's computational complexity and available resources of the decision maker. A discussion of investor rationality is not complete without considering other psychological factors that affect the investor. One such factor, emotion, is considered below.

2.8.2 Emotion

Elster (1998) defines emotions based on six characteristics: cognitive antecedents (beliefs), intentional objects, physiological arousal (changes in the hormone levels of the nervous system), physiological expressions (body expressions), valence (a ranking of emotion on the pleasure-pain scale) and action tendencies (impulses that lead to a response). Emotions and their precedent beliefs form the difference between human beings and other animals. While animals experience, say pain or hunger, they do not form beliefs on these experiences. Thus, a human being will form an intentional object due to the emotion felt for that object – it is based on cognitive antecedent. Biologically, emotions are caused by hormonal changes as well changes to the autonomic nervous system which results in physiological expression. They have some sort of measurable scale and lead to actions if the emotion is powerful enough. In a perfectly rational world, it is considered that any action or belief formed by emotion has no place. However, emotions can help to maximise utility by the act of rational decision making - they force the agent to make a decision and sometimes make the most optimal decision. Elster (1998) shows that emotions play a dual role in decision making for choice and reward. Given a set amount of time, emotions assist the agent to limit the information received and analysed, forcing the agent to make a decision based on the options available. The ideal of maximal rationality assumes that there are no surprises, misunderstandings or irresolvable conflicts but this maxim cannot guide actions that are available in a given amount of time. Elster (1998) argues that bounded rationality forces a decision to be made and avoids an "addiction to reason" in which the agent will always procrastinate for the arrival of new information.

Visceral factors, according to Loewenstein (2000) refer to a range of negative emotions that motivate the agent to engage in a specific behaviour. Tseng (2006) argues that contrary to popular belief, visceral factors are systematic instead of erratic and unpredictable. However, the cognitive deliberations of these visceral factors are unpredictable. These factors result in long-lasting and significant consequences that affect behaviour. Thus, they play an important role in decision making under uncertainty as they force the agent to make a decision and not procrastinate. Visceral factors have been usually left out of traditional economic and financial modelling as they have been seen as too unpredictable. However, a rational assessment of a risk with a corresponding choice of action will often differ from the emotional reaction to that

risk. For example, the lack of a response to an emergency is normally caused by a heightened emotional reaction in the agent. The emotional arousal serves as a functional equivalent for the rational faculties it has temporarily suspended, by inducing a behaviour that is rationally required that would have been reasoned out by the agent if more time was available. While emotions can assist in solving problems quicker than rational processing, the capacity for emotions to enhance rationality at times would not exist if the same emotions also undermine it at times.

Emotional states can be categorised into "hot" and "cold" (Tseng, 2006); where an emotion in a hot state can be fear or greed and an emotion in a cold state can be rational calmness. The difference between behaviour in these two states is known as the empathy gap. As seen by real world evidence, the behaviour of investors during hot states biases those investors towards making mistakes (and a subsequent loss). Tseng (2006) argues that traders need to close this empathy gap to earn long-term financial returns and satisfaction. Indeed, since emotions encompass bounded rationality, behavioural finance and neuro-finance, they have a profound impact on the decision making process. In practice, these emotions are seen through the decisions of investors which are inexplicable according to traditional finance theory.

2.8.3 Neuro-finance

Using scientific methods from other fields, finance has made significant advances in its theory. With the development of behavioural finance, there arose a need to empirically test some of the assertions about investors and their behaviour. Neuro-finance has emerged as a front-runner in this regard. This field analyses financial markets by applying neuro-technology to observe and understand the trading behaviour of market participants (Tseng, 2006). The underlying assumption in neuro-finance is that market participants have different psycho-physiological traits which affect their decision making ability. Behavioural finance investigates the actions of investors during the act of trading and decision making and evaluates these against the backdrop of established psychological theory. In contrast, neuro-finance examines why and how these behaviours occur based on the biological profile of the investor (through hormonal changes and brain activity). Neuro-finance is closely related to

neuro-economics with the main emphasis being on financial market activity instead of all economic behaviour.

Tseng (2006) also refers to neuro-finance as being medical finance in that the biological profile of the investor can be explained through knowledge from the medical field. For example, damage to a particular area of the brain (the orbital frontal cortex) may result in abnormal financial decision making; melancholic depression may cause excessive sleepiness and chronic risk aversion; anxiety can be characterised by excessive risk perception and may lead to panic selling, impulsive overtrading or avoidance of financial markets. Results from experimental studies show that several medications can change the risk-return perception of participants. Further, investors may need psychological support to avoid common behavioural biases. Investors will have the tendency to minimise denial, disappointment and anger when they have made the wrong financial decision. Lo and Repin (2002) conduct controlled experiments on investors and traders using positron emission tomography (PET) and magnetic resonance imaging (fMRI) to understand brain activity and psycho-physiological characteristics when making financial decisions. They find that emotional responses are a significant factor in real-time processing of financial risk amongst professional traders. Kuhnen and Knutson (2005) use similar techniques to examine deviations from the decisions of a rational agent (one where risk-seeking or risk-aversion mistakes are not warranted). Their results show that when people anticipate physical pain, adverse visual stimuli, risky choices or anxiety, the part of the brain that handles cognitive functioning (the anterior insula) is activated. In contrast, when people anticipate monetary gain, the emotional centre (also the centre of addiction) of the brain is activated.

2.9 Summary

This chapter outlined the literature on market efficiency, beginning with a qualitative exposition on how the concept of market efficiency emerged in finance academia. Simply, a market is considered efficient if one cannot use any means available to consistently earn abnormal returns, through the prediction of future stock prices. Market efficiency is not a new concept in the literature as the term has been used since the late 19th century. However, the concept became popularised by Fama (1970) in defining the Efficient Market Hypothesis,

which stated that no abnormal profits may be made over time as prices reflect both private and public information.

Since its popularisation, the EMH has generated a multitude of both empirical evidence and more recently, alternate market theories. Concerned with whether asset prices reflect all available information, researchers who conducted tests on the EMH were regularly faced with problems in the form of either unrealistic assumptions, or one of a joint hypothesis - testing both an asset pricing model and the EMH simultaneously. While some may have overcome this obstacle, the work of Fama (1970) has certainly fostered a greater understanding of financial markets. From the viewpoint of a market participant, studies have attempted to analyse the speed of adjustment of prices to new information; while others have taken the statistical definition of the EMH (that share prices follow a random walk) and have attempted to test the hypothesis. However, irrespective of the viewpoint chosen, there is no consensus on whether markets are efficient according to the EMH.

A digression to time series methods was thereafter discussed, to provide a foundation for the econometric and neural network used in this thesis. Often, in analysing a time series, one can mistake the presence of chaos in the series as randomness. The necessary requirement is that the system of equations be non-linear in order to generate chaotic solutions as a linear system will necessarily generate a trend in its output. These outputs are often mistaken as random time series and are only accurate for a length of time governed by the errors of the initial conditions and the Lyapunov exponent of the system. Various time series models, ranging from simple to complex, were presented as an "evolution" of the field to what led to models being developed in the field of computer science. This evolution can be seen as the search for the "perfect" model. Once a model is developed and permeates into the academic community, empirical testing of it leads to robust descriptions of its appropriateness. In the event that it has a particular shortcoming, a new research question emerges in that one then tries to improve on the existing model. Thus, while one can use logic to deduce which model is more appropriate than another, this argument is limited by the universe of available models as well as the shortcomings inherent in any model. In other words, while a neural network may be more appropriate to use, it is not without its disadvantages nor is it the "best" out of the universe of models that can be chosen.

A discussion of asset pricing followed, where both considerations of investor rationality and the influence of exogenous factors were presented. Coupled with the foundation provided for time series methods, the discussion on asset pricing would then provide a background and motivation for the artificial intelligence models used, along with the inclusion of exogenous factors that could influence stock returns.

Lastly, some of the emerging (and perhaps esoteric) areas of finance research were discussed, providing a well-rounded view of how inter-disciplinary collaboration can provide solutions to long standing questions in finance. At face value, the field of finance is concerned primarily with observation and empirical testing. Any new theory introduced to the field is grounded on a set of assumptions either related to the market participant or to the applicable world at large. As such, the realm of behavioural finance, evolutionary finance and neurofinance provide alternate views on finance theory. For example, the extent of a market participant's rationality is often considered a matter of perspective as some theories rely on a participant having full or strict rationality, whereas others rely on reasonable levels of rationality (bounded rationality).

Conceptually, the arbitrage pricing framework covered in this chapter can be considered the starting point for any investigation into the inclusion of additional risk factors, however, one must first ensure that the data used is correctly processed and indeed deterministic (or at least non-random). Data collection and processing are fundamental to ensuring accurate insights are generated from the attempt to answer a research question. Further, one should also take cognisance of any nuances inherent in a method or model that is used - as is the case with neural networks.

3 Data and Methodology

Testing for cyclical efficiency requires one to first test whether returns follow a random walk or some deterministic process. If the latter is found to be true, then it implies that some form relationship over time is present in the data and this data generating process can be modelled. A first attempt at modelling this process is to use autoregressive models which show the relationship between the contemporaneous return and historic returns. If this model is found to be suitable, yet still contain significant constant or error terms, it implies that additional factors apart from historic returns influence contemporaneous returns. One then investigates this hypothesis using, in this particular case, a neural network, where the data generating process can be non-linear but unknown to the researcher *a priori*. Further, the use of neural networks also acts as a further test of random walk behaviour, adding to the library of existing methods. Using the proposed framework to examine cyclical market efficiency, one can also investigate whether the sampling frequency has any impact on the results in both an individual share level and aggregated index level. This chapter outlines the data collected and used in the study as well as those particular techniques selected for determining if market efficiency is cyclical.

3.1 Data

As alluded to above, this thesis will examine the hypothesis of cyclical market efficiency on both an individual share level as well as aggregated index level, over different sampling frequencies and over sub-samples. Closing prices for the local equity, equity indices (local and international), macroeconomic data, fundamental and behavioural related data were obtained for the period September 1997 to October 2014. Three data sources were used, ranging from McGregor BFA, Bloomberg and the South African Reserve Bank; and each variable contains total returns (inclusive of corporate actions or dividends where applicable). Given the task of ensuring returns inclusive of dividends are correctly incorporated, the simplifying assumption of using the dividend yield (converted to the appropriate frequency) was used. Thus, the total return is the sum of the share price change and the dividend yield. With respect to the indices used on a monthly basis, the total return index (TRI) of those indices were obtained and used instead of the method outlined previously. The sample period was chosen so as to ensure full daily, weekly and monthly data were available (a longer

sample period could have been used if a particular frequency or less shares were required). The number of observations ranges from 4480 for daily data, 896 for weekly data, 206 for monthly data, 68 for quarterly data and 34 for semi-annual data. While the different frequencies add an element of robustness to the study, the choice of frequencies is certainly not exhaustive. Indeed, one can extend the frequencies to include the highest (high-frequency data or tick-by-tick data), to lower ones (perhaps annual). For the purposes of this thesis, these five popular frequencies are chosen, with remaining frequencies left to future research.

Table 3 below shows the individual equity series used, along with select equity indices; whereas Table 4 below describes the candidate variables and frequency of data used for modelling purposes. Forty four local equity series were randomly selected from the top 100 shares by market capitalisation on the JSE, as of October 2014, along with six local equity indices. From the local equity series, the nine shares presented in bold below, along with the JSE Top 40, are used in the results and discussion, whereas the remaining results are displayed in the Appendix.

Table 1 - Shares and indices used

Share Code	Share Name	Industry
SAB	South African Breweries	Consumer Goods - Beverage
BIL	BHP Billiton	Basic Materials - Mining
NPN	Naspers	Consumer Services - Media
MTN	MTN Group	Telecommunications - Mobile Telecommunicatons
SOL	Sasol	Oil and Gas - Oil and Gas Producers
AGL	Anglo American	Basic Materials - Mining
FSR	Firstrand Group	Financials - Banks
SBK	Standard Bank Group	Financials - Banks
APN	Aspen Healthcare	Healthcare - Pharmaceuticals and Biotechnology
BGA	Barclay's Group Africa	Financials - Banks
RMH	RMB Holdings Ltd	Financials - Banks
MDC	Medi-Clinic Corp	Health Care - Health Care Equipment and Services
SHF	Steinhoff International Holdings	Consumer Goods - Household Goods and Home Construction
INP	Investec	Financials - Financial Services
MPC	Mr Price Group	Consumer Services - General Retailers

IMP	Impala Platinum	Basic Materials - Mining
NTC	Network Healthcare	Health Care - Health Care Equipment and Services
MMI	MMI Holdings	Financials - Life Insurance
ANG	Anglogold	Basic Materials - Mining
IPL	Imperial Holdings	Industrials - Industrial Transportation
NPK	Nampak	Industrials - General Industrials
GFI	Goldfields	Basic Materials - Mining
ASR	Assore	Basic Materials - Mining
INL	Investec Limited	Financials - Financial Services
PIK	Pik N Pay Stores	Consumer Services - Food and Drug Retailers
TFG	The Foschini Group	Consumer Services - General Retailers
SNT	Santam	Financials - Nonlife Insurance
HYP	Hyprop Investments	Financials - Real Estate Investment Trusts
SAP	Sappi	Basic Materials - Forestry and Paper
CLS	Clicks Group	Consumer Services - Food and Drug Retailers
GND	Grindrod	Industrials - Industrial Transportation
PPC	Pretoria Port Cement	Industrials - Construction and Materials
AFE	A E C I Ltd	Basic Materials - Chemicals
RCL	RCL Foods	Consumer Goods - Food Producers
SUI	Sun International	Consumer Services - Travel and Leisure
ILV	Illovo Sugar	Consumer Goods - Food Producers
RLO	Reunert	Industrials - Electronic and Electronic Equipment
FBR	Famous Brands	Consumer Goods - Travel and Leisure
MUR	Murray & Roberts	Industrials - Construction and Materials
SPG	Super Group	Industrials - Industrial Transportation
FPT	Fountainhead Property	Financials - Real Estate Investment Trusts
SAC	SA Corporate Real estate Fund	Financials - Real Estate Investment Trusts
OCE	Oceana Group	Consumer Goods - Food Producers
WBO	Wilson Bayley Holmes Ovcon	Industrials - Construction and Materials
J150	JSE Gold Mining Index	
J200	JSE Top 40	
J203	JSE All Share Index (ALSI)	
J211	JSE Industrial 25	
J213	JSE Financial and Industrial 30	
J177	JSE Mining Index	

The candidate variables below are examined before being used in any modelling procedure for stationarity, normality and correlation. These variables are defined as the first difference of the original variable dataset. Those variables that have passed this initial screening are then used as exogenous inputs into the respective econometric models presented below. Apart

from the financial crisis in 2007, the last reported trough in the market cycle was during 1999. This sample period provides a sufficient framework to examine the market cycle; especially given the global recession since 2007 (the period under investigation corresponds to a 'complete' market cycle in that it includes a peak between two troughs. From the candidate variables in Table 4 below, those that "pass" the initial screening are described in the results of Chapter 4.

Table 2 – Variables considered

	Variable	Acronym used	Frequency released
Economic	R153 bond	R153	Daily
	R157 bond	R157	Daily
	Oil Price	Oil	Daily
	Gold Price	Gold	Daily
	Prime Rate	Prime	Monthly
	PPI	PPI	Monthly
	CPI	CPI	Monthly
	GDP	GDP	Quarterly
Equity - Returns	MSCI World Index	World	Daily
	MSCI BRIC Index	BRIC	Daily
	MSCI EMEA Index	EMEA	Daily
	FTSE 100	FTSE	Daily
	S&P 500	S&P	Daily
	Hang Seng 100	Hang Seng	Daily
Equity - Fundamentals	ALSI Earnings Yield	EY	Daily
	ALSI Dividend Yield	DY	Daily
	ALSI Volume	Vol	Daily
	ALSI Price/Earnings Ratio	PE	Daily

Five frequencies of the data are used in this study - daily, weekly, monthly, quarterly and semi-annually. For those variables that did not have daily observations, they were made to follow a step-wise linear function, in that once new data is captured, the current variable maintains the same value until the next data point is captured. In the event that a daily observation was not available, an assumption was made that the release date of say, GDP data, occurs on the last day of the month. To extend the example of GDP data, if GDP data

was released on 31 January 2013, the daily value of future observations remains the same until the next GDP release data one quarter later. Thus, when converting returns to a logarithmic scale, the log returns remain at a value of 0. This is found to be intuitive in the manner an investor interprets information that is not released daily. The "latest" value of information is kept by the investor until new information is released. It should also be noted that the international indices used exclude dividends.

Further, the models are run on the full sample, non-overlapping samples and over-lapping samples to add robustness to the results. The full sample period is split evenly into 10 sub-samples that do not overlap and span 21 months of data. These sub-samples consist of 448 observations.

3.2 Methodology

In effect, there are three phases of the methodology. The first phase involves testing for random walk behaviour on the return series; the second testing for an autoregressive data generating process with no additional variables apart from the lagged dependent variable; and the third with testing the data generating process with additional (lagged) variables without pre-specifying the functional form of the model.

3.2.1 Testing for normality

Three tests for normality are presented to ensure robustness of the results. These tests cover parametric, non-parametric and graphical evidence on whether the data used exhibits normality or not.

3.2.1.1 The Jarque Bera test

The Jarque Bera (JB) provides a goodness of fit statistic of whether a sample distribution matches a normal distribution by examination of the skewness and kurtosis measures. The JB test statistic is defined as:

$$JB = \frac{n}{6}\left(S^2 + \frac{(K-3)^2}{4}\right) \quad \{37\}$$

where S is the skewness of the sample and K is the kurtosis of the sample. These are respectively given by:

$$S = \frac{\hat{\mu}_3}{\widehat{\sigma^3}} = \frac{\frac{1}{n}\sum_{i=1}^{n}(x_i - \bar{x})^3}{\left(\frac{1}{n}\sum_{i=1}^{n}(x_i - \bar{x})^2\right)^{3/2}} \quad \{38\}$$

$$K = \frac{\hat{\mu}_4}{\hat{\sigma}^4} = \frac{\frac{1}{n}\sum_{i=1}^{n}(x_i - \bar{x})^4}{\left(\frac{1}{n}\sum_{i=1}^{n}(x_i - \bar{x})^2\right)^{2}} \quad \{39\}$$

where $\hat{\mu}_3$ and $\hat{\mu}_4$ are estimates of the third and fourth moments of the distribution, $\bar{x}$ is the sample mean and $\hat{\sigma}^2$ is the sample variance. The test statistic is asymptotically distributed with a chi-squared distribution with two degrees of freedom. This is used to test the null hypothesis that both the skewness and excess kurtosis are set to zero.

3.2.1.2 The Q-Q plot

The Quantile-Quantile (Q-Q) plot is a visual aide for depicting the probability distributions of two samples (populations) against each other. The set of intervals for the quantiles are chosen from each distribution and are plotted as a pair of coordinates. A particular coordinate corresponds to one of the quantiles of a distribution plotted against the same quantile of the other distribution. Therefore, the ensuing line of coordinates form a curve across each numbered quantile. The Q-Q plot is most often used to compare a sample distribution to a normal distribution. If the two distributions are similar, then the line of coordinates would be roughly shown at a 45° angle. As such, the Q-Q plot is a non-parametric approach to determining if a distribution is normal.

3.2.1.3 The Kolmogorov-Smirnov test

The Kolmogorov-Smirnov (K-S) test is a non-parametric test for normality. It compares one continuous probability distribution to a reference probability distribution (considered a one-sample K-S test). The null hypothesis for the test is that both samples are from the same distribution, and the test statistic follows a Kolmogorov distribution. While this test is a quite popular non-parametric means of testing for normality, it is often less accurate than other tests such as the Shapiro-Wilk test or Anderson-Darling test. However, it is deemed appropriate to use this test as opposed to the others mentioned as the former does not perform well under data that has multiple identical values; while the latter does not perform well with small samples. The K-S test statistic is given by:

$$D_n = \sup_x |F_n(x) - F(x)| \qquad \{40\}$$

where $\sup_x$ is the supremum of the set of distances, $F_n(x) = \frac{1}{n}\sum_{i=1}^{n} I_{X_{i\leq x}}$ and $I_{X_{i\leq x}}$ is an indicator function that is equal to unity if $X_{i\leq x}$ and zero otherwise.

3.2.2 Testing for linearity

The BDS test can be considered a popular means of establishing whether a series is non-linear. It was originally designed to test if a distribution's observations were independent and identical for the purposes of detecting non-random chaotic behaviour. The test statistic defines a correlation integral which measures the frequency of which a temporal pattern is repeated. Consider a time series x_t for $t = 1, 2, 3, ..., T$ and define its *m-history* as $x_{mt} = (x_t, x_{t-1}, x_{t-2}, ... x_{t-m+1})$. Then, the correlation integral at point *m* can be estimated by:

$$C_{m,\epsilon} = \frac{2}{T_m(T_m - 1)} \sum_{m\leq s} \sum_{<t\leq T} I(x_t^m, x_s^m, \epsilon) \qquad \{41\}$$

where $T_m = T-m+1$ and $I(x_t^m, x_s^m, \epsilon)$ is an indicator function equal to unity if $|x_{t-i} - x_{s-i}| < \epsilon$ and zero otherwise. Thus, the correlation integral measures the probability that any two *m-dimensional* points are within a distance ϵ of each other. If the observations are independent and identically distributed, then this probability is equal to:

$$C_{1,\epsilon}^{m} = \Pr\left(|x_t - x_s|\right)^m \quad \{42\}$$

Therefore, the BDS test statistic is given by:

$$V_{m,\epsilon} = \sqrt{T}\frac{C_{m,\epsilon} - C_{1,\epsilon}^{m}}{s_{m,\epsilon}} \quad \{43\}$$

where $s_{m,\epsilon}$ is the standard deviation of $\sqrt{T}(C_{m,\epsilon} - C_{1,\epsilon}^{m})$ and the test statistic follows a normal distribution.

3.2.3 Testing for stationarity

Two tests for stationarity are presented below. These tests are commonly used in tandem to determine if a series is both stationary and has a unit root.

3.2.3.1 The Augmented Dickey Fuller test

Detection of a random walk first requires tests for autocorrelation. This would measure the relationship between the share return at the current period to a value in a previous period. The Augmented Dickey Fuller (ADF) test is a popular metric used to measure autocorrelation in a series. The version of the ADF test applied includes an intercept and trend.

$$\Delta y_t = c_0 + c_1 t + \delta y_{t-1} + \beta \sum_{i=1}^{p} \Delta y_{t-1} + \mu_t \quad \{44\}$$

where c is a constant term, $c_1 t$ is the trend term, p is the number of lags of y and u_t is white noise.

3.2.3.2 The Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test

The KPSS test examines the null hypothesis that the time series under consideration is stationary around a deterministic trend. The series is decomposed into a deterministic trend, random walk and stationary error and the test uses the Lagrange multiplier method to test the hypothesis that the random walk component has a zero variance. The KPSS test complements the ADF test in that by utilising both, one can determine if a series appears to be stationary and appears to have a unit root, for which more data points is required. The test statistic is given by:

$$KPSS = T^{-2} \frac{\sum_{t=1}^{T} \hat{S}_t^2}{\widehat{\lambda^2}} \qquad \{45\}$$

where $\hat{S}_t^2 = \sum_{j=1}^{t} \hat{u}_j$, and $\hat{u}_t$ is the residual of the regression of the deterministic component on the series itself and $\widehat{\lambda^2}$ is the estimate of long-run variance of $\hat{u}_t$. Under the null hypothesis that the series is stationary, the test statistic follows a Gaussian distribution, with the parameters of the distribution being dependent on the form of the determinstic terms in the regression.

3.2.4 Testing for random walk behaviour

Given the multitude of tests for random walk behaviour, one needs to be cognisant of which tests are used to accurately provide results. As such, four tests are considered, with each being an improvement on the prior.

3.2.4.1 Runs test

The runs test is a non-parametric test for detecting whether a series is random. If a series is random, then the observed number of runs should be close to the expected number of runs. A run is defined as a sequence of consecutive (price) changes with the same sign. Thus, there are three categories of run: upward, downward and flat. Under the null hypothesis of independence in share returns, the total expected number of runs (*m*) is estimated as:

$$m = \frac{N(N+1) - \sum_{i=1}^{3} n_i^2}{N} \quad \{46\}$$

where N is the total number of observations and n_i is the number of price changes in each of the three categories.

For a large number of observations, the sampling distribution of m is approximately normal and the standard error of m is given by

$$\sigma_m = \sqrt{\frac{\sum_{i=1}^{3} n_i^2 \left[\sum_{i=1}^{3} n_i^2 + N(N+1)\right] - 2N \sum_{i=1}^{3} n_i^3 - N^3}{N^2(N-1)}} \quad \{47\}$$

Standard normal Z statistics can be used to test whether the hypothesis of independence is rejected. A disadvantage of the Runs test shown above is that it can only detect randomness at a lag order of one only. While other versions of the Runs test have been developed, more powerful tests examine the decomposition of variance.

3.2.4.2 Variance ratio test

The variance ratio test of Lo and MacKinlay (1988) is shown by many authors to be an adequate test of the weak form of the EMH. The test assumes that the variance of increments in the random walk series is linear in the sample interval - the variance should be proportional to the sample interval. Specifically, if a series q follows a random walk, the variance of its *q-differences* should be a q multiple of the first difference.

$$Var(p_t - p_{t-q}) = qVar(p_t - p_{t-1}) \quad \{48\}$$

The Variance Ratio (VR) is then calculated as

$$VR(q) = \frac{\frac{1}{q} Var(p_t - p_{t-q})}{Var(p_t - p_{t-1})} = \frac{\sigma^2(q)}{\sigma^2(1)} \quad \{49\}$$

For a sample size of *n(q+1)* observations the formulae for computing the variances are modified as follows

$$\sigma^2(1) = \frac{\sum_{t=1}^{nq}(p_t - p_{t-q} - q\hat{u})^2}{(nq-1)} \quad \{50\}$$

$$\sigma^2(q) = \frac{\sum_{i=q}^{nq}(p_t - p_{t-q} - q\hat{u})^2}{h} \quad \{51\}$$

where $h = q(nq + 1 - q)(1 - \frac{q}{nq})$

and $\hat{u} = \frac{1}{nq}\sum_{t=1}^{nq}(p_t - p_{t-1}) = \frac{1}{nq}(p_{nq} - p_0)$

Under the assumption of either homoscedasticity or heteroscedasticity, two standard normal statistics, *Z(q)* and *Z*(q)*, can be used.

$$Z(q) = \frac{VR(q) - 1}{\sqrt{\emptyset(q)}} \quad \{52\}$$

$$Z^*(q) = \frac{VR(q) - 1}{\sqrt{\emptyset^*(q)}} \quad \{53\}$$

A shortcoming of the Lo and MacKinlay (1988) VR test is that the lag order q is required to be specified. Thus, a modified version of this test is employed, by Chow and Denning (1993) as this tests for multiple lags of order q. As both the single and multiple order VR test statistics have shortcomings in their reliance to an approximated distribution, these tests can often give rise to size distortions or low power (Wright, 2000).

Thus, the modification of Wright (2000) is used as this provides a non-parametric version of the Lo-MacKinlay test, displaying results for the variance decomposition based on ranks (*R1*,

R2 variables in the test) and sign (*S1*). Assume that $r(Y_t)$ is a rank of return Y_t among T_1, T_2, ..., T_r, then $r(Y_t)$ is the number from *1* to *T* given by

$$r_{1,t} = \frac{r(Y_t) - \frac{T+1}{2}}{\sqrt{\frac{(T-1)(T+1)}{12}}} \quad \{54\}$$

$$r_{2,t} = \Phi^{-1}\left(\frac{r(Y_t)}{T+1}\right) \quad \{55\}$$

where Φ is the standard normal cumulative distribution function and Φ^{-1} is its inverse. The series $r_{1,t}$ is a linear transformation of the ranks that is standardised to have a sample mean of 0 and a sample standard deviation of 1. The *R1* and *R2* test statistic are defined as

$$R_1 = \left(\frac{\frac{1}{Tk}\sum_{t=k}^{T}\left(r_{1,t} + r_{1,t-1} + \cdots + r_{1,t-k+1}\right)^2}{\frac{1}{T}\sum_{t=1}^{T} r_{1,t}^2}\right)\left(\frac{2(2k-1)(k-1)}{3kT}\right)^{-0.5} \quad \{56\}$$

$$R_2 = \left(\frac{\frac{1}{Tk}\sum_{t=k}^{T}\left(r_{2,t} + r_{2,t-1} + \cdots + r_{2,t-k+1}\right)^2}{\frac{1}{T}\sum_{t=1}^{T} r_{2,t}^2}\right)\left(\frac{2(2k-1)(k-1)}{3kT}\right)^{-0.5} \quad \{57\}$$

Similarly, Wright (2000) defines a sign statistic, s_t, by being equal to 0.5 if the return Y_t is positive and -0.5 otherwise. The sign based variance ratio test statistic, *S1*, is thus defined as:

$$S_1 = \left(\frac{\frac{1}{Tk}\sum_{t=k}^{T}\left(s_{1,t} + s_{1,t-1} + \cdots + s_{1,t-k+1}\right)^2}{\frac{1}{T}\sum_{t=1}^{T} s_{1,t}^2}\right)\left(\frac{2(2k-1)(k-1)}{3kT}\right)^{-0.5} \quad \{58\}$$

Therefore, the Chow and Denning as well as Wright modifications of the VR test are used as the former examines multiple variances and the latter ranks and signs. In addition, a graphical plot of the variance decomposition over time would reveal if the series follows random walk behaviour or not. If the variance decomposition is not within acceptable confidence intervals,

then it implies that the variance does not decompose over time as expected. Perhaps a more sophisticated version of the variance ratio test, the Hurst exponent provides a measure of long term memory in a time series and as such is discussed below.

3.2.4.3 Hurst Exponent

To test for non-linear dependence, the Hurst exponent is used. Zunino *et al.* (2009) argue that the exponent measures the long range dependence in stock market indices, where an existence of autocorrelation between distant observations will imply market inefficiency. The exponent provides a measure of memory and fractality of a time series. Ranging from values between 0 and 1, the Hurst exponent can identify if a time series follows a random walk or is persistent. A value of 0 indicates that the series is anti-persistent (mean-reverting); a value of 1 indicates that the value is persistent and a value of 0.5 indicates that the series is random. Further, there are various permutations of calculating the Hurst exponent, leading one to be cautious in the preference of one calculation over another. Taqqu, Teverovsky and Willinger (1995) conduct simulations of the different methods of the Hurst exponent on data of differing sample sizes to empirically determine the best method to use for a particular sample size. The authors find that for series that have between 4000 and 7000 data points, the Peng estimator should be used; for series with 700 to 1000 data points, the Whittle Estimate be used; and for series less than 700 data points, the R/S method be used. All three methods are discussed below. The first considers analysing both the mean and standard deviation of a time series; the second on detrending the time series and then analysing the variance to determine the Hurst exponent; whereas the third relies on a periodogram fit.

3.2.4.3.1 Rescaled range estimate

Hurst (1951) describes the process for running the Hurst exponent. A sliding window approach partitions the series into subsamples that exhibit random walk behaviour and non-linear dependence. In a given time series, $\{Y_t\}_1^T$, and a window of size d, an initial sub-sample is created that consists of d observations. The tests to classify the prevailing pattern in the data are run and the window is then incremented by one observation – from *2* to *d+1*.

This continues until all observations are classified (from observation *T-d* to observation *T*). To calculate the Hurst exponent, one uses rescaled range (R/S) analysis.

1. Calculate the arithmetic mean of the series (or window), *m*.

2. Calculate the mean-adjusted series, $\{X_t\}$, by subtracting m from each observation

3. Calculate $\{Z_t\}$, the cumulative deviation of $\{X_t\}$.

4. Calculate$\{R_t\}$,the range of $\{Z_t\}$

5. Calculate$\{S_t\}$, the standard deviation of $\{Y_t\}$

6. Calculate the rescaled range as $\frac{R_t}{S_t}$

7. The Hurst exponent is then estimated from $\frac{R_t}{S_t} = c * t^H$, where c is a constant.

3.2.4.3.2 Peng Estimate

Consider a noisy time series, $u(i)$, where $i = (1,2,3, \dots, N_{max})$. This time series can be integrated to obtain:

$$y(j) = \sum_{i=1}^{j} (u(i) - \bar{u}) \quad \{59\}$$

where $\bar{u} = \frac{1}{N_{max}} \sum_{j=1}^{N_{max}} u(i)$ and is divided into *n* equal partitions. In each partition, the integrated time series has a polynomial function that is fit to it, $y_{fit}(i)$, which is known as the local trend. The integrated time series, $y(i)$, is detrended by subtracting the local trend in each partition and is calculated as:

$$Y(i) = y(i) - y_{fit}(i) \quad \{60\}$$

For a given partition size, *n*, the Root Mean Square fluctuation is calculated as:

$$F(n) = \sqrt{\frac{1}{N_{max}} \sum_{i=1}^{N_{max}} [Y(i)]^2} \quad \{61\}$$

The above calculation is repeated for *n* partition sizes to provide a relationship between F(n) and *n*. A power-law distribution between *F(n)* and *n* indicates the presence of scaling, given by:

$$F(n) \sim n^{\alpha} \quad \{62\}$$

The parameter, α, is called the scaling exponent and represents the correlation properties of the time series. If α = 0.5, then there is no correlation and the time series is white noise. If α < 0.5, then the series is mean-reverting; and if α > 0.5, then the series is trending.

3.2.4.3.3 Whittle Estimate

Assume that the spectral density of a self-similar process is denoted by $f(\lambda, \theta)$ where $\theta = (\theta_1, \theta_2, \theta_3, \dots, \theta_M)$. If $\theta_1 = \sigma_\epsilon^2$ where σ_ϵ^2 is the variance of the innovation ϵ of the infinite autoregressive representation of the process, then this implies that $\int_{-\pi}^{\pi} log\{f(\lambda, \theta)\}\, d\lambda = 0$ and θ_2 represents the Hurst parameter, *H*. The Whittle estimator, $\hat{\eta}$, of $\eta = (\theta_2, \theta_3, \theta_4, \dots, \theta_M)$ minimises the quality of fit function:

$$Q(\eta) = \int_{-\pi}^{\pi} \frac{I(\lambda)}{f(\lambda; (1, \eta)} d\lambda \quad \{63\}$$

where $I(\cdot)$ denotes the periodogram of the time series of length *N* and is defined as:

$$I(\lambda) = \frac{1}{2\pi N} \left| \sum_{j=1}^{N} X_j e^{ij\lambda} \right|^2 \quad \{64\}$$

$\widehat{H}$ is given by $\hat{\theta}_2$ and the estimate of σ_ϵ^2 is given by:

$$\sigma_{\epsilon}^{2} = \int_{-\pi}^{\pi} \frac{I(\lambda)}{f(\lambda,(1,\hat{\eta}))} d\lambda \quad \{65\}$$

3.2.4.3.4 Significance levels

While the Hurst exponent (and its various methods) are considered powerful tests of random walk behaviour, the method in general suffers from a lack of distribution theory to correctly allocate confidence intervals to interpret the results. In other words, faced with an answer of 0.49 for the Hurst exponent, one does not have a clearly defined interval to determine if 0.49 is statistically close (or not) from 0.5. As such, authors have proposed three avenues to determine the significance of the Hurst exponent. The first relies on conducting the test using a variety of methods and simply choosing the consensus. The other relies on simulating data to obtain confidence intervals that can be applied in general to a sample of finite observations; and the final considers a simple case of the inverse of the number of observations in the sample (this provides a point estimate as opposed to a confidence interval). This thesis relies on the second method and uses robust estimates obtained from the literature. Weron (2002) provides equations based on simulations to estimate the confidence intervals for the Peng and Whittle estimators. These equations are as follows, where $N = \log_2 n$ and n is the sample size. Further, Weron (2002) notes that the Whittle estimator is the only known Hurst exponent method which has known asymptotic properties. In other words, one can only rely on approximate statistical results as opposed to exact statistical results, the former of which is based on the behaviour of those statistics in large samples. As a consequence, the confidence interval for the Whittle estimate is considerably larger than the other two methods used in this thesis, as the latter two do not have asymptotic properties.

Table 3 - Hurst exponent confidence intervals for the Peng estimate

Confidence interval	Lower bound	Upper bound
90%	$0.5 - e^{(-2.99\ln(N)+4.45)}$	$0.5 + e^{(-3.09\ln(N)+4.57)}$
95%	$0.5 - e^{(-2.93\ln(N)+4.45)}$	$0.5 + e^{(-3.10\ln(N)+4.77)}$
99%	$0.5 - e^{(-2.67\ln(N)+4.06)}$	$0.5 + e^{(-3.19\ln(N)+5.28)}$

Table 4 - Hurst exponent confidence intervals for the Whittle estimate

Confidence interval	Lower bound	Upper bound
90%	$0.5 - e^{(-0.71N^{\frac{2}{3}} + 1.87)}$	$0.5 + e^{(-0.68N^{\frac{2}{3}} + 1.62)}$
95%	$0.5 - e^{(-0.71N^{\frac{2}{3}} + 2.04)}$	$0.5 + e^{(-0.68N^{\frac{2}{3}} + 1.78)}$
99%	$0.5 - e^{(-0.73N^{\frac{2}{3}} + 2.45)}$	$0.5 + e^{(-0.65N^{\frac{2}{3}} + 1.92)}$

Rasheed and Qian (2004) provide a confidence interval for the traditional R/S method used in this study, also based on simulations. The authors do not provide an equation to calculate the Hurst exponent yet do provide a mean and standard deviation value. As such, the confidence intervals used in this thesis are as follows. The reader will note that the higher the confidence interval, the wider the range between the lower and upper bound. A wider (larger) confidence interval implies that the chance of the observation at hand being equal to the true population value is higher. Or equivalently, the researcher is less likely to reject the null hypothesis that the observed value is not "close enough" to the true value.

Table 5 - Hurst exponent confidence intervals

Frequency	Method	90%		95%		99%	
		Lower	Upper	Lower	Upper	Lower	Upper
Daily	Peng	0.4508	0.5432	0.4429	0.5515	0.4260	0.5685
Weekly	Whittle	0.2492	0.7241	0.2027	0.7630	0.0913	0.8471
Monthly/Quarterly/ Semi-annual	R/S	0.4656	0.6252	0.4503	0.6405	0.4205	0.6703

In summary, the preliminary tests of the dataset involve normality, linearity, stationarity and random walk behaviour. To test for normality, three tests are used - one parametric, one non-parametric and one graphical method. To test for linearity, the most powerful and popular test, the BDS test, is used. To test for stationarity and for a unit root, the ADF and KPSS tests are used. Lastly, to test for random walk behaviour, a non-parametric Runs test, variance ratio tests (one parametric and the other non-parametric), the variance decomposition plot; and the Hurst exponent is used. If the time series under consideration indicates non-random behaviour, it is possible that one can model this time series. As such, econometric models can be used.

3.2.5 Modelling the return generating process

If the returns process does not follow a random walk, one proceeds to model the process itself. Two such models are presented here, the first being a sophisticated autoregressive model (the SETAR model) and the second being a neural network (a NARX network).

3.2.5.1 SETAR Models

A SETAR model is the simplest form of a threshold AR model. The SETAR model divides a time series into a piecewise linear function over a particular threshold value. Conceptually, the SETAR model creates several linear time series models from a non-linear time series. When the piecewise function is a function of the lagged dependent variable, the model is referred to as a SETAR model as the dependent variable is dependent on lagged values of itself ("self exciting"). Let Y_t be a univariate time series and let $X_t = (Y_{t-1}, Y_{t-2}, Y_{t-3}, \dots, Y_{t-p})'$, a *k x 1* vector with *k = 1 + p*. The *SETAR(m)* model can be defined as:

$$Y_t = \alpha_1' X_{t-1} I_{1,t}(\gamma, d) + \cdots + \alpha_m' X_{t-1} I_{m,t}(\gamma, d) + e_t \qquad \{66\}$$

where $\gamma = (\gamma_1, \dots, \gamma_{m-1})$ with $\gamma_1 < \gamma_2 < \cdots < \gamma_{m-1}$ and $I_{j,t}(\gamma, d) = I(\gamma_{j-1} < Y_{t-d} \leqslant \gamma_j)$. Further, $I(\cdot)$ is the indicator function, $\gamma_j$represent the j^{th} threshold, with *j* being any integer and *d* is the delay parameter, which is usually strictly positive.

3.2.5.2 SETAR Modelling Procedure

The typical modelling procedure presented here is consistent with that of Granger and Terasvirta (1993). In modelling returns, one can rely on the grid search procedure in the add-in for the statistical program, *R*, referred as the "*tsdyn*"[26] package, to find optimal parameter values.

[26] A time series package that implements non-linear autoregressive models by Di Nurzo, Aznarte and Stigler (2009).

1) First, it is required to specify a linear autoregressive model of order *k*. The appropriate order is chosen by the lowest Akaike Information Criterion or Ljung-Box Statistics.

2) Once specified, the researcher needs to test the linearity of the above model to determine the delay parameter, *d*. Note that the equation below tests for linearity in the return series. The function "setartest" in R tests the series against 3 alternatives - a linear AR model, a SETAR model with one threshold value and a SETAR model with two threshold values.

3) Now that the use of the SETAR model has been established, the coefficients remain to be estimated. This can be done through non-linear least squares (NLS) or equivalent methods. Insignificant coefficients are dropped from the regression until only those coefficients that are significant remain. In this process, accurate estimates of γ are usually difficult and accompanied by a high standard error. A low significance of γ should not be interpreted as weak evidence of non-linearity, as large changes in γ have little impact on the value of the transition function. Thus, high significance of γ is not necessary.

4) To further evaluate the accuracy of the model, forecasting of future share returns is done. There are primarily two methods of approach. One could forecast out of sample by either assuming the parameters to be constant, or one could forecast for the next interval (the next month) and use the new observation to re-estimate the SETAR model. This thesis uses the former approach. The most common method employed for testing of forecasts, is the Root Mean Squared Error (RMSE):

$$RMSE = \left\{ \left(\frac{1}{n}\right) \sum_{\tau-t+1}^{n} (A_\tau - F_\tau)^2 \right\}^{0.5} \quad \{67\}$$

where *n* is the number of observations in each forecast, A_τ is the actual return at time *t*, F_τ is the forecast return at time *t*. While the SETAR model assists in modelling regime changes, it must still be specified in advance before it can be applied. Attention is now turned to neural network models, which do not require an *a priori* specification.

3.2.6 Building an ANN

Basheer and Hajmeer (2000) provide a process for building an ANN, as shown in Figure 2 below. While these steps are intuitive, it is nonetheless instructive to discuss them here.

Phase 1: Problem definition and formulation

Prior to conducting any research, the problem needs to be adequately understood, with particular attention to causal relationships that may be present. The authors also suggest that other techniques be explored before a final decision is made to use ANNs.

Phase 2: System design

Reliant on the abilities of the modeller, the ANN is designed. This would involve data collection, any filtering or processing of the data, statistical analysis of the data and partitioning of the data into training, test and validation samples.

Phase 3: System realisation

This phase involves training the network and the optimal selection of the parameters used (such as connection weights, learning rates and number of hidden nodes). Recently, evolutionary techniques, such as a genetic algorithm, have been shown to assist in this phase.

Phase 4: System verification

Once the optimal parameters are selected, the ANN is then tested on the validation sample.

Phase 5: System implementation

If the network has performed adequately, the ANN is then programmed using the appropriate computer hardware and software.

Phase 6: System maintenance

If there are exogenous factors that cause the characteristics of the data to change over time, the ANN would need to be retrained. This final phase involves ensuring that the minimum acceptable error is adhered to at all times.

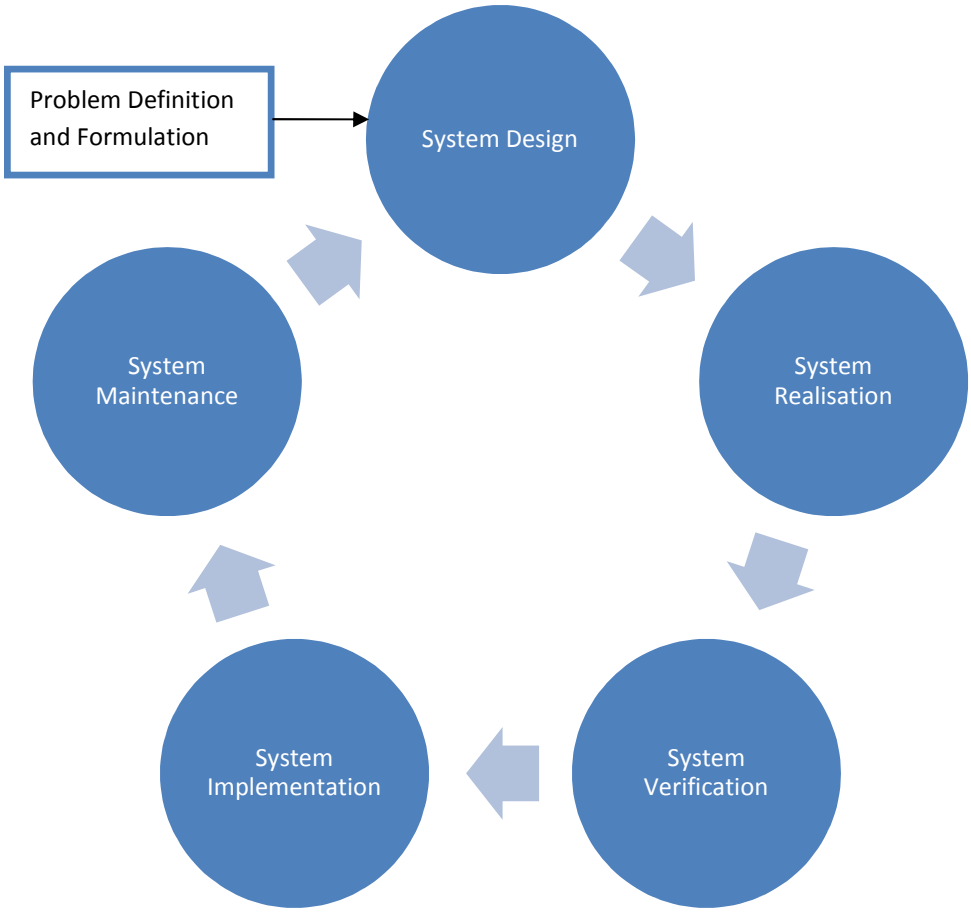


Figure 2 – The phases of ANN development

3.3 Neural network hierarchy

Angus (1991) provides guidelines on selecting the best network for the application at hand. The author suggests that the type of network be guided on its applicability to the problem. Roughly, the problem statement can be split between time-variant and time-invariant problems. A time-variant problem would relate to some spatio-temporal pattern, where the time stamp of the variable(s) in question plays some role in the output. In contrast, a time-invariant problem does not require any dependence on a time stamp.

Generally, feed-forward networks are sufficient for learning time-invariant problems, however, there are particular networks, such as Tapped Delay Neural Networks (TDNNs) that can be used. Angus (1991) argues that the use of hidden states in a neural network (NN) expands the range of applications for the NN. Recurrent Neural Networks (RNNs) can be used to model time-varying problems, recognise patterns or for forecasting purposes. These networks can model non-linear chaotic, dynamic systems and in principle, should be able to predict future values of the output variable. As such, the family of RNNs is considered more applicable to the problem of modelling cyclical market efficiency. In particular, a non-linear autoregressive with exogenous inputs (NARX) RNN will be used.

3.3.1 Non-Linear Autoregressive Models with Exogenous inputs (NARX)

NARX recurrent neural networks are a form of non-linear models which determine current output values from past input and past output values. A NARX network can be described as follows:

$$y(t) = f\big(u(t - D_u), \dots, u(t-1), u(t), y\big(t - D_y\big), \dots, y(t-1)\big) \qquad \{33\}$$

where *u(t)* and *y(t)* represent the input and output respectively of the network, D_u and D_y are the lags of the input and output respectively; and *f* is a non-linear function. The NARX network is illustrated in Figure 3.

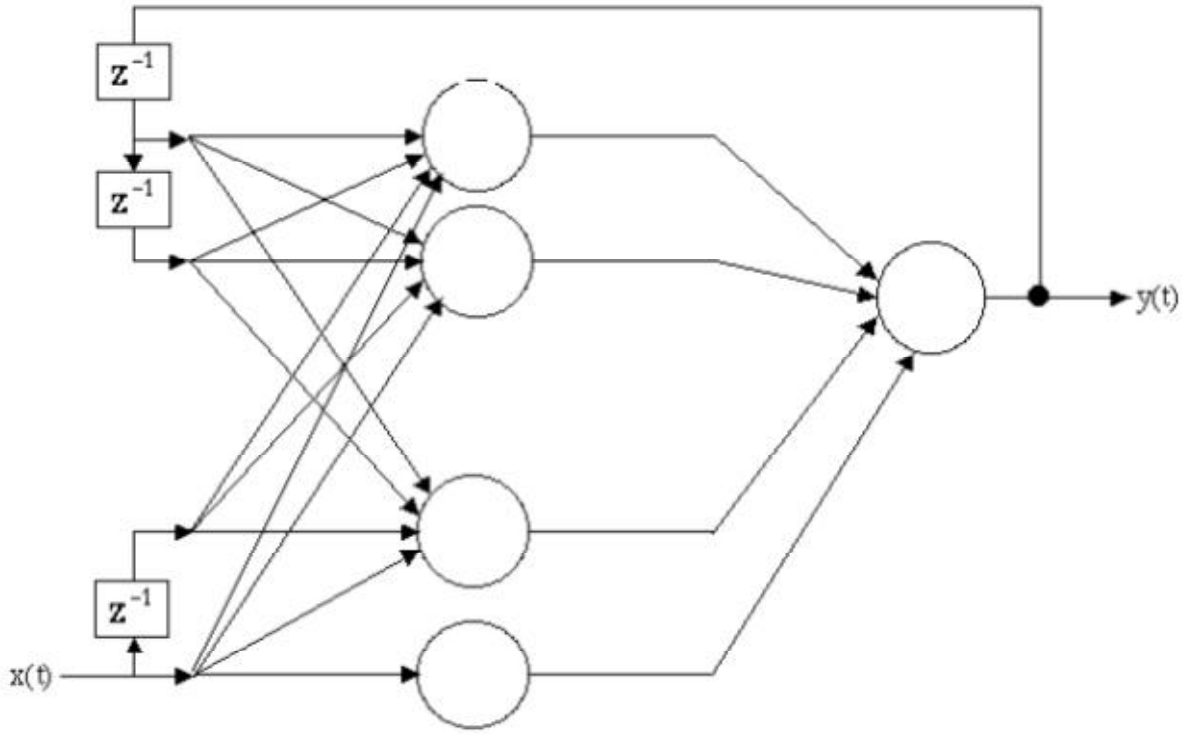


Figure 3 – A Non-linear Autoregressive with Exogenous Inputs Network

NARX NNs train and converge much faster compared to their traditional NN counterparts. They are also quite adept at learning long term dependencies (Lin, Horne, Tino and Giles, 1996) and can store information over extended periods of time. This thesis uses a NARX network for evaluate and forecast the ALSI.

Tino, Horne and Giles (2001) describe a NARX network of zero input order, but suggest that the results can be generalised to higher order inputs. A zero order input would simplify Equation (33) to the following:

$$y(t) = \psi\big(\ u(t), y\left(t - D_y\right), \dots, y(t-1)\ \big) \qquad \{34\}$$

where Ψ represents the mapping performed by an MLP. Graphically, this is depicted in Figure 4 below.

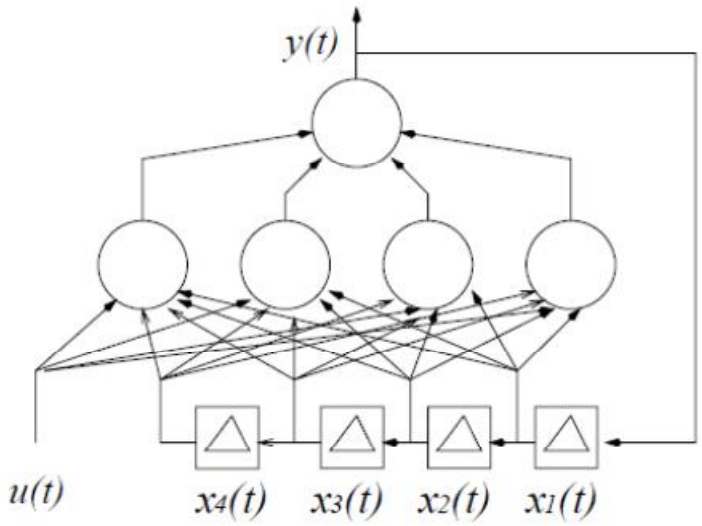


Figure 4 – A Multi-layer perceptron

To assist in describing the network, Kailath (1980) suggests that the equations are transformed into state space format. This assists in examining the Jacobian matrix. As the states of a discrete-time dynamical system can be mapped with the unit-delay elements in the realisation of the system, the NARX NN can be described in state space form as:

$$x_i(t+1) = \begin{cases} \Psi(u(t), \mathbf{x}(t)) & i = 1 \\ x_{i-1}(t) & i = 2, \ldots, D \end{cases} \quad \{70\}$$

and

$$y(t) = x_1(t+1) \quad \{71\}$$

While NARX networks work well with time series data, they are still prone to long-term dependencies that exist within the data. This issue can be mitigated by dividing the dataset into smaller sub-samples. The process for constructing a neural network is now discussed.

3.4 Issues in ANN development

The success or failure of a network at its task is often heavily weighted towards the ability of the researcher to circumvent certain issues in data collection, processing and network training. This section discusses the most pertinent of these issues and their influence on the result of the network in matching its output to the actual data.

3.4.1 Database size and partitioning

Perhaps the most detrimental issue to using an ANN in research, the sample size needs to adequate enough for training and testing, without being too large to affect the accuracy of the ANN. Conceptually, the sample size should be large enough to account for possible known variations in the definition of the problem to be solved. An example would be to use a training sub-sample that covers a market cycle. The sample is partitioned into a training, testing and validation sub-sample. The training sample should be described as above – sufficiently large to cover possible known variations in the data. The testing sub-sample should be sufficiently different to that in the training sub-sample, without being considered completely unrelated. An example would be to use a testing sub-sample that covers a market cycle that is different to the one used in training. Lastly, the validation sub-sample is used after the optimal neural network is modelled. Once again, it must be sufficiently different from the previous data, within reason. An example would again be to use data that covers a market cycle, different from the previous sub-samples.

Looney (1996) suggests that 65% be used for training, 25% for testing and 10% for validation. The latter suggestion is adopted in this thesis by splitting the data into the aforementioned percentages for training, testing and validation as it allows sufficient observations to be used in each stage.

3.4.2 Data pre-processing, balancing and enrichment

To accelerate training of the network, the data often needs to be pre-processed. This can be achieved by removing noise, reducing the number of variables, deletion of outliers and transforming the data (Swingler, 1996).

3.4.3 Data normalisation

Normalisation or scaling of the data within a uniform range prevents larger numbers from overriding smaller ones and premature saturation of the hidden nodes (Basheer and Hajmeer,

2000). Further, for extremely large values, the logarithm of the data may be used prior to normalisation. This would avoid outliers in the data and assist in network training time.

3.4.4 Input and Output representation

Data representation is an important and critical factor in the design of an ANN according to Masters (1994). It may be possible to convert continuous input data to a discrete, binary form to extract rules from a trained network (Fu, 1995). Other specialised algorithms exist for conversion of continuous variables to discrete form based on the distributions (Kerber, 1992). These algorithms allow flexibility in the use of networks as they are capable of handling both discrete and continuous data, transforming the input or output to enhance network accuracy while still providing a tractable means of examining non-linear processes.

3.4.5 Network weight initialisation

Initialising network weights involves assigning an initial, zero-mean random number to each connection (Rumelhart, Hinton and Williams, 1986). The literature does not agree on the importance of selecting the "correct" initial weight. Arguably, while a particular initial weight will assist in speeding up the training time, it can be considered unnecessary if the computer hardware available is sophisticated enough to not be affected by the non-initialised weights.

3.4.6 The backpropagation learning rate, η

Apart from data processing and narrowing the search parameter for neuron weights, one can adjust the parameters of the learning algorithm. While a large value for the learning rate will accelerate training, the search algorithm on the error surface may never converge – leading to over-fitting of the model. However, a small value for the learning rate may result in the network taking too much time to converge on a solution. Authors (Wythoff, 1993; Zupan and Gasteiger, 1991 and Fu, 1995) have suggested learning rates between 0.1 and 1.0; 0.3 and 0.6; and 0.0 to 1.0, respectively. Alternatively, an adaptive learning rate may be used which

will vary along the course of training. This alternative is appealing as, generally, the distance from a minimum cannot be predicted (one will only know the distance from the minimum after it has been reached). Further, when the search algorithm is far away from the minimum, a larger learning rate is required; whereas a smaller learning rate is required when the search algorithm is near the minimum.

3.4.7 The backpropagation momentum coefficient, μ

Haykin (1994) states that the inclusion of a momentum term assists in stabilising the search algorithm for the global minimum. A higher momentum coefficient will accelerate the weight updates and reduce the risk of the search algorithm not converging. However, it also increases the risk of over-fitting. Similar to the learning rate, either a constant or adaptive value can be used. Wythoff (1993) suggests a learning rate between 0.4 and 0.9 whereas Fu (1995) suggests a rate between 0.0 and 1.0. Others, such as Zupan and Gasteiger (1991) suggest a combined learning rate and momentum coefficient approximately equal to unity. An adaptive momentum coefficient will fluctuate as the training progresses. This technique can be used in conjunction with the methods suggested above, in that as the momentum coefficient increases, the learning rate decreases. Practically, the value of the momentum coefficient also impacts the computer storage space of the researcher.

3.4.8 The activation function, σ

A correctly specified activation function is important in the development of an ANN. The choice of activation function is dependent on the objective of the ANN. For example, step functions can be used to indicate whether a neuron is simply activated or not, regardless of the magnitude of activation. ANNs that use backpropagation algorithms usually use a sigmoid function as it has properties of both continuity and differentiability on the real number line. While the advantages of using a particular function over another is not yet understood according to Hassoun (1995), Moody and Yarvin (1992) show that the choice of activation function does affect the success of the ANN. Indeed, if the activation function leads to a saturation of values at its bounds, neurons may be inappropriately activated (inhibited) leading to a larger error term.

3.4.9 Convergence criteria

An ANN may be said to converge if: 1) the training error is acceptable ($\rho \leq \varepsilon$) or 2) the gradient error is acceptable ($\nabla\rho \leq \delta$) or 3) there is cross-validation of the output. Basheer and Hajmeer (2000) state that the last criterion is more reliable, at the cost of computing time, power and abundance of data. Thus, many researchers use the first or second criteria, or a derivation thereof. For example, one can use the coefficient of determination, R^2, or the standard square error, SSE[27], given by:

$$SSE = \frac{1}{N}\sum_{p=1}^{N}\sum_{i=1}^{M}\left(t_{p,i} - O_{p,i}\right)^2 \quad \{35\}$$

Where $O_{p,i}$ and $t_{p,i}$ are the actual and target solution of the i^{th} output node on the p^{th} training example of N examples and M output nodes. Such an approach incorporates a measure of complexity in the network architecture and was introduced by Garth *et al.* (1996).

3.4.10 Number of training cycles

The most intuitive (and perhaps best) approach to determine the optimal number of training cycles is through trial and error. While a large number of training cycles may be beneficial in assisting learning, it may lead to over-training of the network and a complete recall of the data (as opposed to a prediction). While the SSE of the network may not follow a strictly smooth path, one can consider a significant increase in SSE (assuming a decreasing SSE) to indicate that the optimum network configuration has been reached.

[27] Alternate measures, such as the Mean Squared Error (MSE), Root MSE, or Mean Absolute Percentage Error (MAPE), can be used. The first two measures are scale dependent and have the advantage of being the most common and statistically relevant. The latter is scale independent and has the advantage of not being sensitive to outliers (Hyndman and Koehler, 2006).

3.4.11 Modes of training

The optimal number of hidden layers and subsequent hidden nodes is a critical component in network architecture. While the researcher often starts with no *a priori* knowledge on the number of hidden nodes, Basheer (1998) suggests that one hidden layer is sufficient to approximate continuous functions, whereas Masters (1994) suggests two hidden layers for discontinuous functions.

3.4.12 Size of the hidden layer

The optimal number of hidden layers and subsequent hidden nodes is a critical component in network architecture. While the researcher often starts with no *a priori* knowledge on the number of hidden nodes, Basheer (1998) suggests that one hidden layer is sufficient to approximate continuous functions, whereas Masters (1994) suggests two hidden layers for discontinuous functions.

In summary, training a network rests on optimising the parameters of the network. Table 2 below presents a concise view of the effect on each parameter if it incorrectly specified (not all of which is a hindrance to the researcher).

Table 6 – Summary of network parameters

Parameter	Parameter is too large (high)	Parameter is too small (low)
NHN	Over-fitting	Under-fitting
Learning rate, η	Unstable connection weights	Slow training speed
Momentum coefficient, μ	Increased risk of over-shooting minimum error	Entrapment in local error minima.
Number of training cycles	Poor generalisation of untrained data	Incapable of generalising data
Size of training subset	Good generalisation ability	Poor generalisation
Size of testing subset	Can confirm good generalisation ability.	Confirmation of poor generalisation

3.5 Learning algorithms

The ability to learn distinguishes sentient life forms from other biological entities. Similarly, the ability of a network to mimic learning enables the network to increase its accuracy towards the desired output.

> "Learning is defined as the process of updating the internal representation of the system in response to external stimuli so that it can perform a specific task. This includes modifying the network architecture, which involves adjusting the weights of the links, pruning or creating some connection links and changing the firing rules of the individual neurons." (Schalkoff, 1997, p. 128)

ANNs would thus learn through an iterative process by examining the error term generated by the previous network architecture, adapting future network architecture to minimise future error terms. This is similar to the manner in which human beings learn and process information. An ANN is said to have learnt if it can (1) handle imprecise, fuzzy, noisy and probabilistic information without noticeable adverse effects on response quality and (2) generalise from the tasks it has learnt to unknown ones. (Basheer and Hajmeer, 2000).

As per Haykin (1994) and Hassoun (1995), there are four basic learning algorithms. Error-correction learning (ECL) is used in supervised learning in which the arithmetic difference (error) between the ANN solution at any stage during training and the corresponding correct answer is used to modify the connection weights so that the overall network error is gradually reduced. The most popular learning algorithm used in ECL is the backpropagation (BP) algorithm. As a precursor to the BP algorithm, the gradient descent method is used to minimise the error function through updating the weights of the neurons. The method finds the gradient of the weight space and selects the steepest descent at each iteration, finding either a minimum or infinitely decreasing path. When the minimum is found, it is not necessary the global minimum, which can incorrectly lead to premature stopping of the training of the network. The BP algorithm avoids this pitfall by introducing two more parameters (the learning rate and the momentum parameter) that affect the speed at which the system learns. These parameters force the search to consider results from previous iterations,

thereby avoiding the search from finding a local minimum or infinitely decreasing path. This "consideration" is what gives the algorithm its name as it passes information from each output back to the input and hidden layers of the network. The BP algorithm (along with the gradient descent method) is used in training the networks used in this thesis.

3.6 Summary

Recall that the hypothesis of cyclical efficiency will be tested through three phases. Firstly, it is necessary to establish whether share price changes follow a random walk or not. If price changes are random, they cannot be predicted, thus enforcing the notion of weak form market efficiency. However, if price changes are not random, it is then viable to establish whether they can be modelled. The first model requires prior values of the share price as determinants of the current share price. The use of this model is founded on information being contained only in past prices – it can be likened to the semi-strong form of market efficiency where public information is reflected in the share price. Lastly, if the prior step is inadequate, it implies that there exists private information that is not incorporated in past prices and does influence the current share price.

Using five different frequencies of data for fifty equity series, one proceeds to examine the distribution properties of these variables, before testing the random walk hypothesis. Preliminary tests on normality, stationarity and non-linearity are conducted. The latter is of importance as it guides the choice of econometric model to be used. In testing the random walk hypothesis, popular measures from the literature are used - namely the Runs test, Variance Ratio test and the Hurst exponent. Each method is both more complex and more accurate than the preceding, with multiple "versions" of each test also examined. For example, the Chow and Denning (1993) and Wright (2000) modification of the Variance Ratio test is conducted as the former tests for multiple variances whereas the latter is non-parametric in nature. Similarly, the method used to calculate the Hurst exponent differs based on the sample size (thus three methods are used).

If any of the variables are found to not follow a random walk, then an attempt to model the return generating process using autoregressive models is made. Simple ARIMA models, in

addition to threshold models (in particular SETAR models) are generated. The SETAR model is considered a "basic" non-linear econometric model in that it allows for regime changes in the data (recall the non-linearity test was conducted earlier). If these econometric models are found to be lacking, the use of additional risk factors is also considered, along with a neural network model as the former points towards a more complete description of the returns process whereas the latter assists in providing a flexible solution to the modelling problem.

4 Results

This chapter provides and discusses the results in determining if market efficiency is cyclical. Tests and models are run on five differing frequency data to provide robustness to the results. From the sample of 44 shares and 6 indices, select results are displayed in this chapter, with the remaining results being shown in the Appendix. Further, the full sample period for the ALSI only is split evenly into 10 sub-samples that do not overlap and span approximately 21 months of data. This was chosen so as to balance the need for sufficient data points in the NN, circumvent any short term stationarity and to aid in interpreting the performance of the NNs over the short term. According to Moody (1995), the size of the training set assists in dealing with a tradeoff between noise in the data and non-stationarity. A smaller training set makes estimations from the NN more difficult, while a larger training set allows non-stationarity to appear at small time intervals. Beginning with preliminary statistics on the data used, results outline characteristics of the financial variable distributions to inform the choice of model in estimating the data generating process. Pre-specified time series models are then used to model the returns process, followed by unspecified neural network models. In the latter case, exogenous variables are also introduced into the system in the spirit of an APT framework. The neural network models were run in Matlab™ with the remaining analysis conducted in R.

4.1 Preliminary statistics

As with any empirical investigation, it is often useful to be aware of the basic descriptive statistics in the data tested. These simple measures often provide some form of guidance to the researcher in the choice of model to be used and potential caveats of the analysis.

4.1.1 Full sample results

Descriptive statistics for the all frequencies are presented in Table 7 below. The average returns across all frequencies is quite small but increase in magnitude , along with an increasing volatility at lower frequencies. In other words, the standard deviation (Std. Dev) increases as the frequency decreases. Values for skewness and kurtosis vary, in most cases

showing excess kurtosis (defined as kurtosis is excess of a value of 3), but this remains to be statistically shown.

In the case of daily returns, the selection below shows that equities seem to be positively skewed, with indices being negatively skewed. All 10 securities below have excess kurtosis, indicating that returns are more clustered around the mean than expected by pure chance. While the mean returns are quite low (along with their standard deviations), the minimum and maximum returns are quite extreme, in the sense that all minimum returns are negative, with the equities again having greater maximum returns than the two indices. A similar result can be seen with weekly data.

As the frequency of data lowers, the mean, standard deviation, minimum and maximum values typically increase in magnitude, whereas skewness and kurtosis seem to fluctuate, indicative of share specific reasons. For example, the indices have kurtosis values that seem to decrease as the frequency decreases, whereas certain shares have fluctuating kurtosis values (such as NPN). As the frequency decreases, the influence of dividends becomes more evident in creating outlier returns, which, when aggregated to an index level, is not particularly significant.

Table 7 – Descriptive statistics

DAILY RETURNS						
Share Code	Mean	Std. Dev.	Minimum	Maximum	Skewness	Kurtosis
BIL	0.0010	0.0238	-0.1489	0.1973	0.5066	7.6986
MTN	0.0012	0.0261	-0.1875	0.2425	0.2657	8.1143
SOL	0.0009	0.0232	-0.1592	0.1537	0.1558	7.7320
FSR	0.0007	0.0213	-0.1180	0.1356	0.2162	6.5538
SAB	0.0007	0.0177	-0.1278	0.1423	0.2286	7.6986
NPN	0.0011	0.0248	-0.1551	0.1667	0.0788	7.7106
AGL	0.0007	0.0245	-0.1666	0.1544	0.1548	7.1361
TOP40	0.0006	0.0139	-0.1331	0.0882	-0.2386	8.6266
ALSI	0.0006	0.0127	-0.1191	0.0771	-0.3224	8.5930

WEEKLY RETURNS						
Share Code	Mean	Std. Dev	Minimum	Maximum	Skewness	Kurtosis
BIL	0.0049	0.0528	-0.1813	0.4368	0.7894	8.8971
MTN	0.0061	0.0601	-0.3955	0.4732	0.5471	10.9725
SOL	0.0045	0.0531	-0.2117	0.3329	0.5070	7.9750
FSR	0.0040	0.0488	-0.2056	0.3702	0.8103	11.4767
SAB	0.0037	0.0394	-0.1252	0.2303	0.5219	6.0827
NPN	0.0055	0.0584	-0.2830	0.3602	0.0295	7.7439
AGL	0.0035	0.0548	-0.2004	0.2944	0.4019	5.3265
TOP40	0.0032	0.0311	-0.1458	0.1971	0.0700	6.6411
ALSI	0.0032	0.0288	-0.1692	0.1748	-0.1167	7.1818
MONTHLY RETURNS						
Share Code	Mean	Std. Dev	Minimum	Maximum	Skewness	Kurtosis
BIL	0.0202	0.0968	-0.2432	0.3992	0.2818	3.5318
MTN	0.0251	0.1092	-0.5126	0.4971	0.1298	7.3520
SOL	0.0180	0.0940	-0.2896	0.4127	0.3629	4.7137
FSR	0.0161	0.0879	-0.4028	0.3789	0.1372	6.6790
SAB	0.0153	0.0702	-0.2503	0.2305	-0.2436	4.3424
NPN	0.0243	0.1227	-0.4460	0.4382	-0.1322	5.1626
AGL	0.0140	0.1040	-0.3156	0.4711	0.2447	4.8162
TOP40	0.0361	0.0606	-0.3085	0.1739	-0.6145	5.3303
ALSI	0.0140	0.0574	-0.2930	0.1407	-0.7748	6.2000
QUARTERLY RETURNS						
Share Code	Mean	Std. Dev	Minimum	Maximum	Skewness	Kurtosis
BIL	0.0689	0.1316	-0.3399	0.3674	0.5064	5.1107
MTN	0.0818	0.2433	-0.4510	1.3617	0.2656	6.5679
SOL	0.0520	0.1455	-0.3317	0.5113	0.1557	4.7285
FSR	0.0484	0.1401	-0.4106	0.6159	0.2161	3.5509
SAB	0.0689	0.1316	-0.3400	0.3674	0.2285	4.6951
NPN	0.0750	0.2341	-0.5115	0.8581	0.0788	4.7071
AGL	0.0416	0.1812	-0.3980	0.6583	0.1547	4.1329
TOP40	0.0336	0.0960	-0.3200	0.2531	-1.0784	2.0532
ALSI	0.0341	0.0933	-0.3287	0.2384	-1.2030	2.5413
SEMI-ANNUAL RETURNS						
Share Code	Mean	Std. Dev	Minimum	Maximum	Skewness	Kurtosis
BIL	0.1351	0.2545	-0.3232	0.8948	0.6913	0.8557
MTN	0.1854	0.4716	-0.3047	2.3500	2.7543	10.3188
SOL	0.1154	0.2557	-0.3475	0.8495	0.4101	0.3037
FSR	0.1014	0.2180	-0.5062	0.8024	0.2865	2.4805
SAB	0.0833	0.1471	-0.3455	0.2682	-1.2569	1.3387
NPN	0.1656	0.4469	-0.5028	2.0507	1.9539	6.8682
AGL	0.0984	0.3014	-0.6285	0.8088	0.3028	0.2287
TOP40	0.0710	0.1500	-0.3315	0.3615	-0.8220	0.6335
ALSI	0.0711	0.1424	-0.3125	0.3559	-0.8290	0.8147

To expedite training of the network, all of the candidate variables underwent variable selection in Matlab™. This tool selects variables that do not exhibit multicollinearity amongst themselves and the target variable (the ALSI). After variable selection, there were seven remaining variables that were of a daily frequency: oil returns (Oil), gold returns (Gold), change in ALSI dividend yield (DY), change in ALSI earnings yield (EY), S&P 500 returns (S&P), Hang Seng 100 returns (HS), and FTSE 100 returns (FTSE). It is interesting to note that none of the macro-economic variables were selected to be included in the network. It can be inferred, yet remains to be proven, whether the average investor even considers macro-economic data in examining movements in the ALSI. Descriptive statistics on the additional pricing factors are presented in Table 8 below. The mean values reveal little additional information as they are less than 1%. The negative values for the dividend (and to a lesser extent earning) yield occur due to the data being differenced (dividend yield is always a positive number, yet the difference between dividend yields can be negative).

Table 8 – Descriptive statistics of the inputs to the network

	Oil	Gold	DY	EY	S&P	HS	FTSE
Minimum	-0.1444	-0.0891	-0.0881	-0.1884	-0.0947	-0.1473	-0.0926
Maximum	0.1290	0.0964	0.1247	0.2109	0.1096	0.1725	0.0938
Mean	0.0003	0.0003	0.0000	0.0000	0.0002	0.0001	0.0001
Standard deviation	0.0215	0.0115	0.0001	0.0160	0.0126	0.0169	0.0122

4.1.2 ALSI Sub-sample results

Each sub-sample had daily observations in a 21 month period. Table 9 shows that the average daily return is less than 1% in all of the sub-samples, with some evidence of volatility. Further, there is some evidence of kurtosis and skewness, but this remains to be statistically shown.

Table 9 - Descriptive statistics of each sub-sample

Sub-sample	Mean	Std. Dev	Minimum	Maximum	Skewness	Kurtosis
#1	0.0001	0.0170	-0.1191	0.0771	-1.0658	10.6966
#2	0.0010	0.0115	-0.0759	0.0440	-0.5014	8.1980
#3	0.0003	0.0127	-0.0538	0.0608	0.1376	5.2758
#4	0.0002	0.0106	-0.0329	0.0466	0.2715	3.9630
#5	0.0017	0.0091	-0.0383	0.0389	-0.1772	4.7206
#6	0.0009	0.0126	-0.0647	0.0505	-0.4530	6.0672
#7	0.0000	0.0204	-0.0729	0.0709	0.1221	4.1457
#8	0.0007	0.0101	-0.0362	0.0433	-0.1326	4.0770
#9	0.0007	0.0092	-0.0315	0.0373	0.0190	4.9549
#10	0.0006	0.0084	-0.0320	0.0229	-0.4053	4.0974

The descriptive statistics of the data (ALSI returns and pricing factors) reveal that normality is a poor assumption to make for all of the variables under consideration. This is empirically tested in the next sub-section. For ease of reading, the specific results for quarterly and semi-annual data is not presented here, but rather summarised at the end of this chapter (and provided in detail in the Appendix).

4.2 Tests for normality

Three techniques are used to determine if the variables used are normally distributed. The first, the Jarque Bera test is perhaps the most popular in literature and is a parametric test of normality. The second is a graphical inspection of the quantile-quantile (Q-Q) plot, which plots the frequency distribution of the variables against a theoretical quantile distribution. Thirdly, a non-parametric counterpart to the Jarque-Bera test, the Kolmogorov-Smirnov test will be used.

4.2.1 Full sample results

4.2.1.1 Jarque Bera test

The Jarque Bera test for normality is presented in Table 10 and Table 11 below. The table illustrates the χ^2 test statistic, degrees of freedom (D.O.F) and the corresponding p-value. The

null hypothesis of normality is rejected if the p-value is statistically significant. The results of the Jarque Bera test indicate that the null hypothesis of normality is rejected at a significance level of 5% (indeed at a level of 1%) for all equities at a daily and weekly frequency.

Table 10 – Jarque Bera Test for normality using daily frequency data

Share Code	$\chi 2$	D.O.F	P-value
BIL	5 074.0830	2.00	0.00***
MTN	8 115.5240	2.00	0.00***
SOL	4 197.8720	2.00	0.00***
FSR	2 392.3800	2.00	0.00***
SAB	4 159.9920	2.00	0.00***
NPN	4 146.6940	2.00	0.00***
AGL	3 211.2930	2.00	0.00***
TOP40	5 952.0680	2.00	0.00***
ALSI	5 916.8350	2.00	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 11 - Jarque Bera Test for normality using weekly frequency data

Share Code	$\chi 2$	D.O.F	P-value
BIL	1 391.3720	2.00	0.00***
MTN	2 417.6010	2.00	0.00***
SOL	962.3966	2.00	0.00***
FSR	2 780.6240	2.00	0.00***
SAB	395.4563	2.00	0.00***
NPN	840.3172	2.00	0.00***
AGL	226.1857	2.00	0.00***
TOP40	495.6933	2.00	0.00***
ALSI	654.9113	2.00	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

However, under monthly frequency data in Table 12, one share (BIL) statistically follows a normal distribution at the 10% level, whereas all other equities below do not. Two additional mining shares, ANG and GFI also follow a normal distribution at the 10% level, implying that three of six mining shares in the population studied are normally distributed. Further, the Mining and Gold Mining index are also normally distributed. The normality of the returns distribution implies that average returns are expected the majority of the time; and that possibly these returns are randomly generated.

Table 12 - Jarque Bera Test for normality using monthly frequency data

Share Code	$\chi 2$	D.O.F	P-value
BIL	5.1547	2.00	0.08*
MTN	163.1444	2.00	0.00***
SOL	29.7287	2.00	0.00***
FSR	116.8197	2.00	0.00***
SAB	17.5042	2.00	0.00***
NPN	40.7438	2.00	0.00***
AGL	30.3678	2.00	0.00***
TOP40	197.0036	2.00	0.00***
ALSI	108.5021	2.00	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Expanding the results to quarterly and semi-annual data, there are no shares that are normally distributed under quarterly data, but 23 that are normally distributed under semi-annual data. There is no industry specific pattern that emerges, along with no tendency of a share to be normally distributed under multiple frequencies. Similarly, one rejects the null hypothesis of normality at the 5% level of significance for all daily frequency variables considered in the neural network. These values are displayed in Table 13 below.

Table 13 – Jarque Bera test for normality on the input variables

	Oil	Gold	DY	EY	S&P	Hang Seng	FTSE
χ^2	2120.8220	12115.9500	149608.8000	149608.8000	11654.4200	20405.6400	6169.2610
D.O.F	2	2	2	2	2	2	2
P-value	0.00***	0.00***	0.00***	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Therefore, according to the Jarque Bera test for normality, none of the input variables under consideration follow a normal distribution, while in most cases, the equities data do not follow a normal distribution. In classical linear regression modelling, the assumption of normality must be met. However, in the case of neural networks (as well as the non-linear time series models applied here), this assumption is not a concern if it is not met.

4.2.1.2 Quantile-Quantile plot

The Quantile-Quantile (Q-Q) plot displays two probability distributions by plotting their quantiles against each other. If the two distributions have similar plots, then the result would

be a straight line at a 45° angle. Examining the Normal (Q-Q) plots of the daily data in Figure 5 and Figure 6, one finds that all of the equities do not display evidence of normality. Therefore, according to the visual evidence of the Q-Q plot, none of the equities under consideration follow a normal distribution.

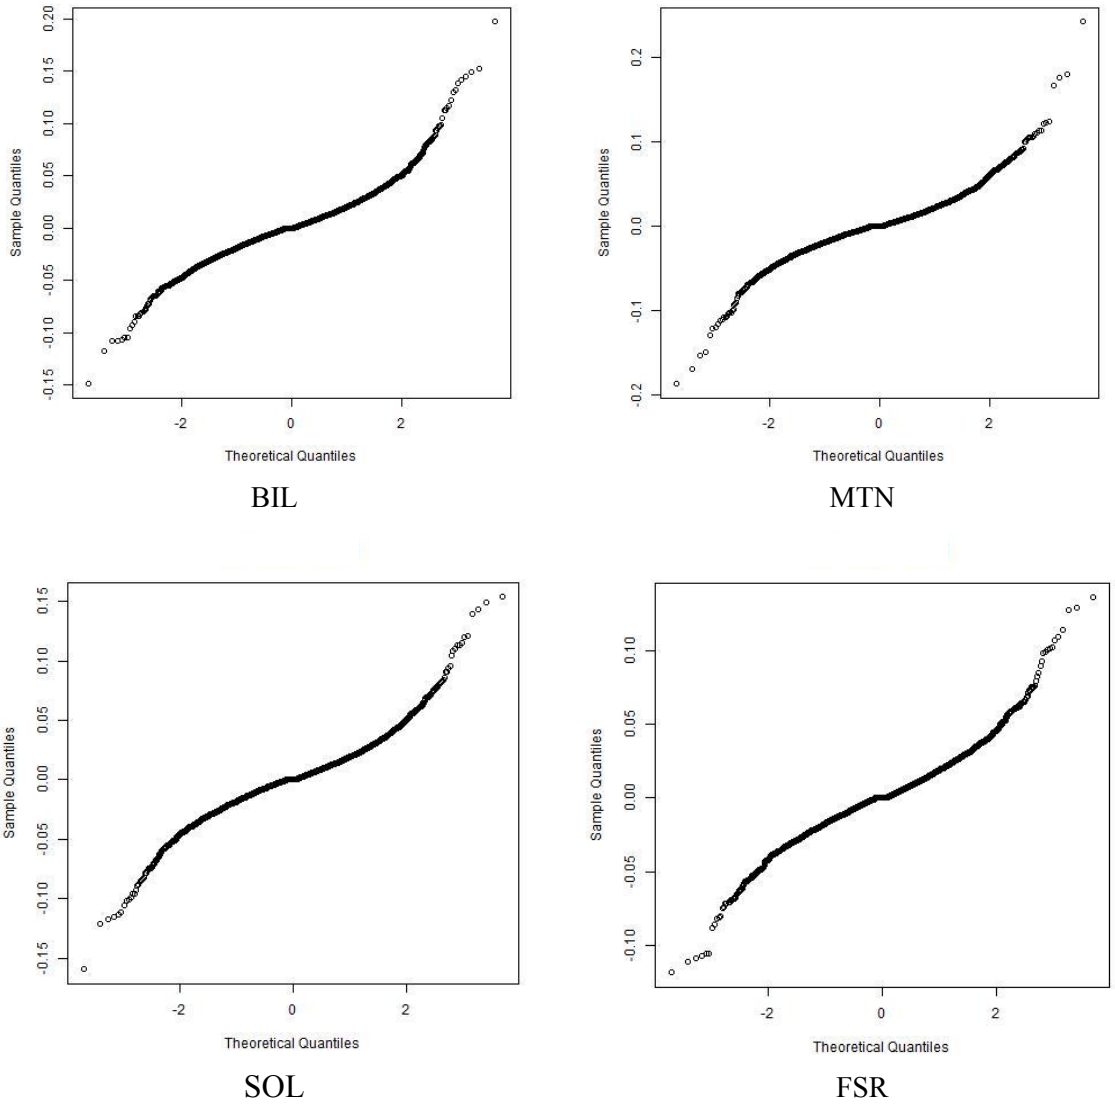
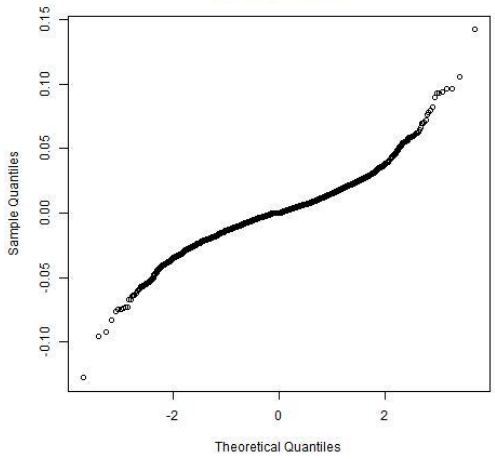
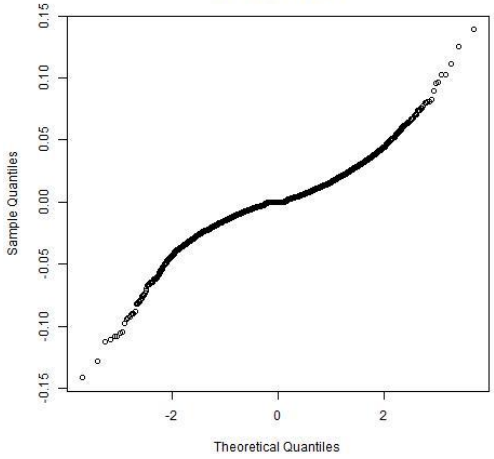


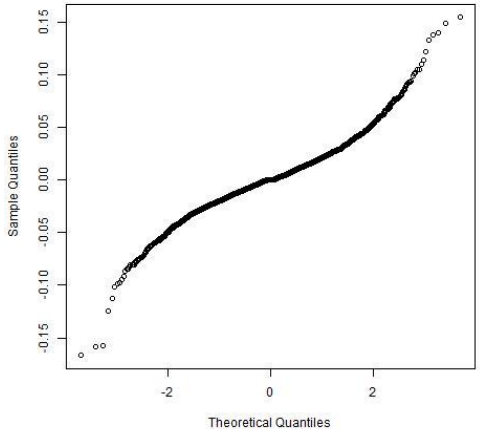
Figure 5 - Q-Q plot of daily frequency data (1)



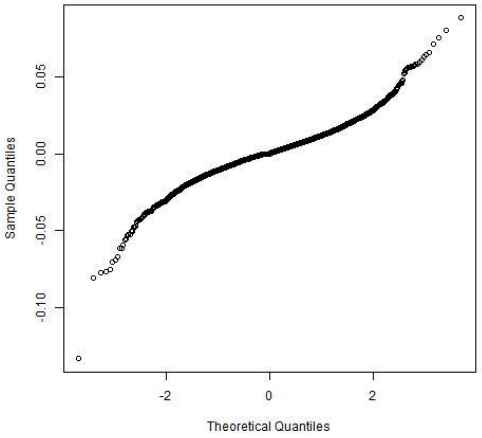
SAB



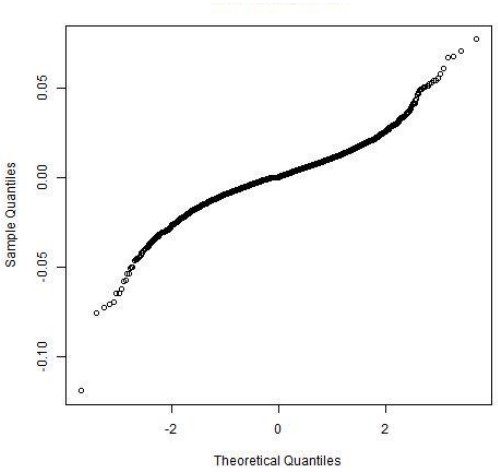
NPN



AGL



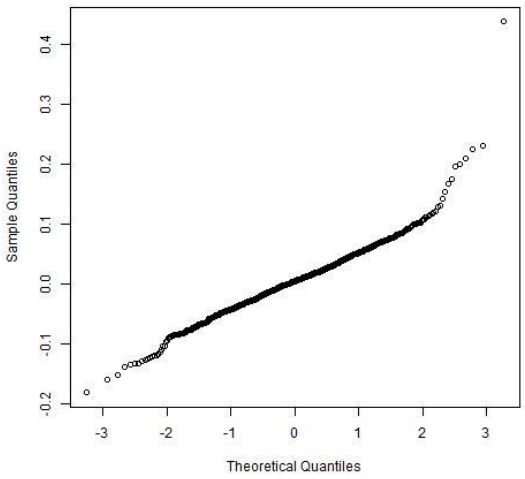
J200



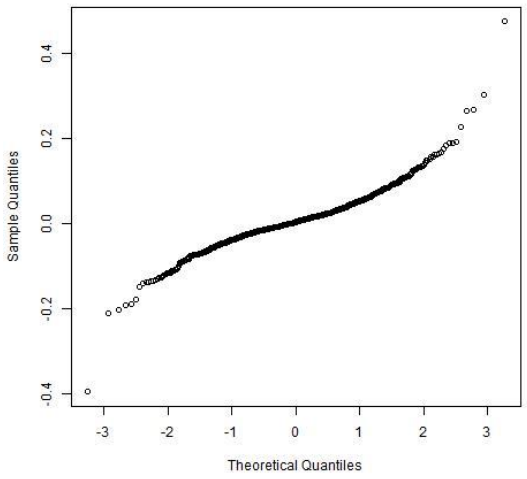
ALSI

Figure 6 – Q-Q plot of daily frequency data (2)

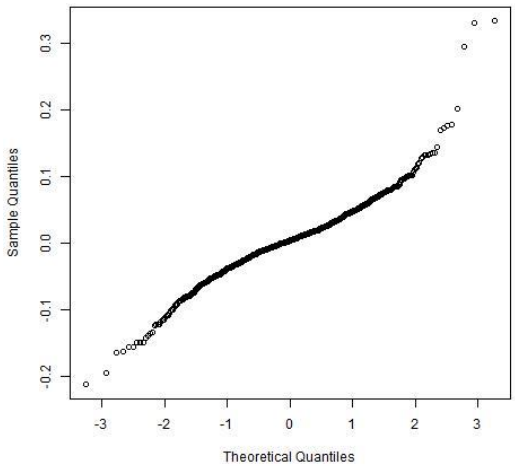
Examining the Normal (Q-Q) plots of the weekly data plots in Figure 7 and Figure 8, one finds that all of the equities do not display evidence of normality. However, there are some equities, such as BIL, SOL, HYP and the ALSI that do weakly follow a normal distribution. As this visual evidence is weak, it can however be concluded that according to the visual evidence of the Q-Q plot as well as the Jarque Bera test, none of the equities under consideration follow a normal distribution.



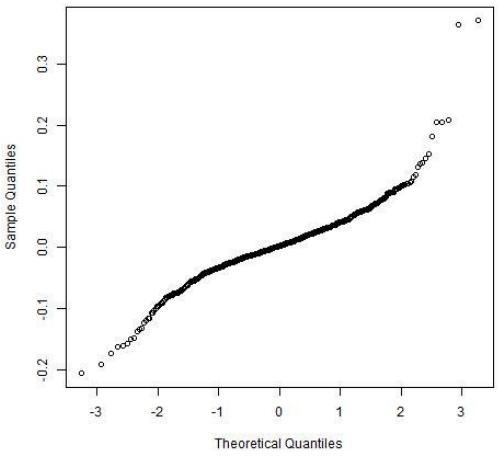
BIL



MTN

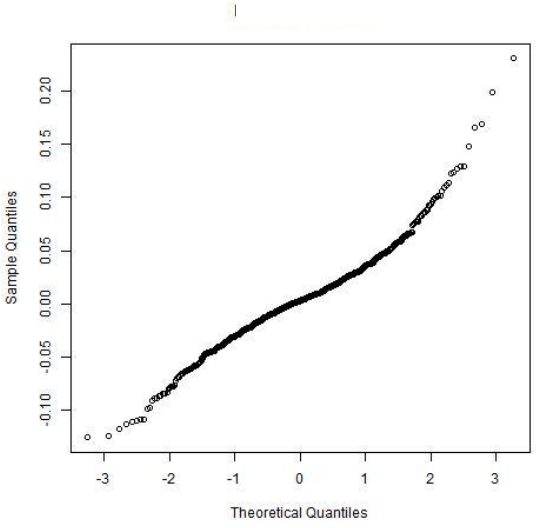


SOL

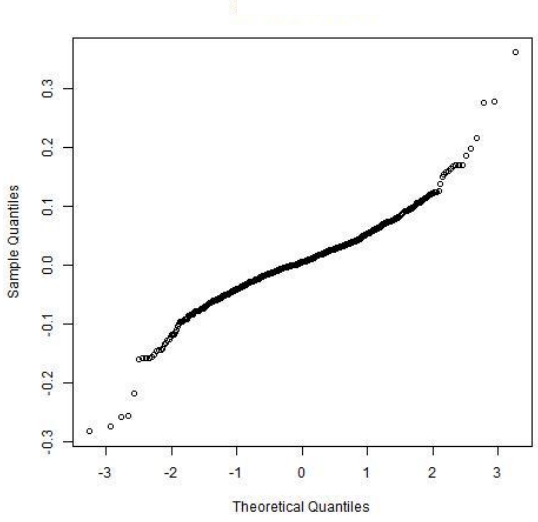


FSR

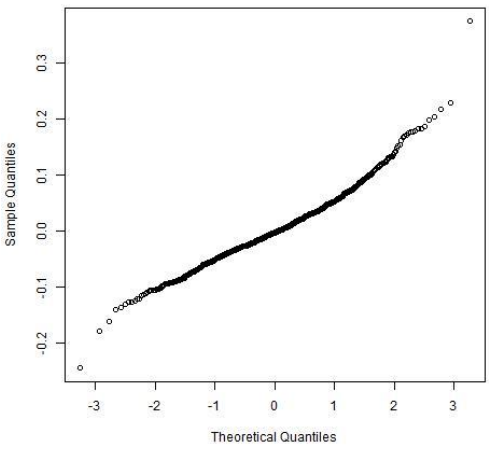
Figure 7 - Q-Q plot of weekly frequency data (1)



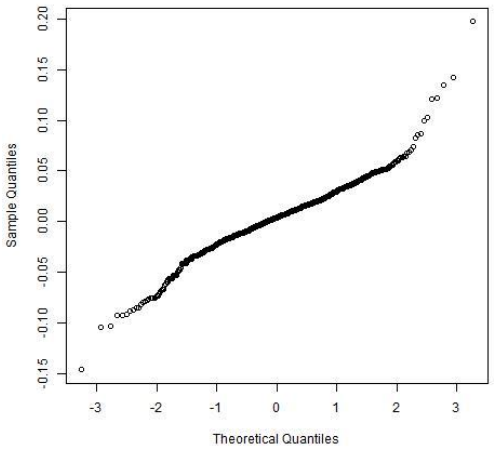
SAB



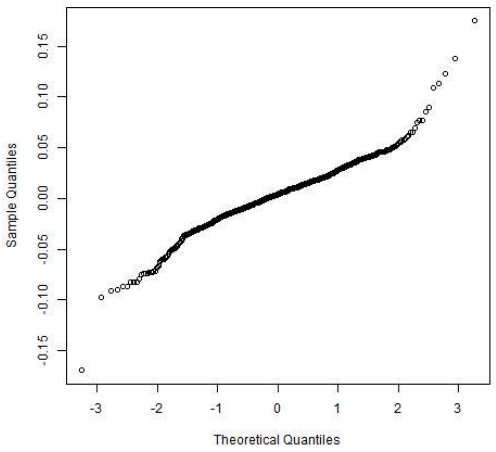
NPN



AGL



J200



ALSI

Figure 8 - Q-Q plot of weekly frequency data (2)

A similar conclusion can be reached when studying the Q-Q plots of the monthly equities data in Figure 9 and Figure 10, albeit the evidence for non-normality is not as significant as with the previous data frequencies.

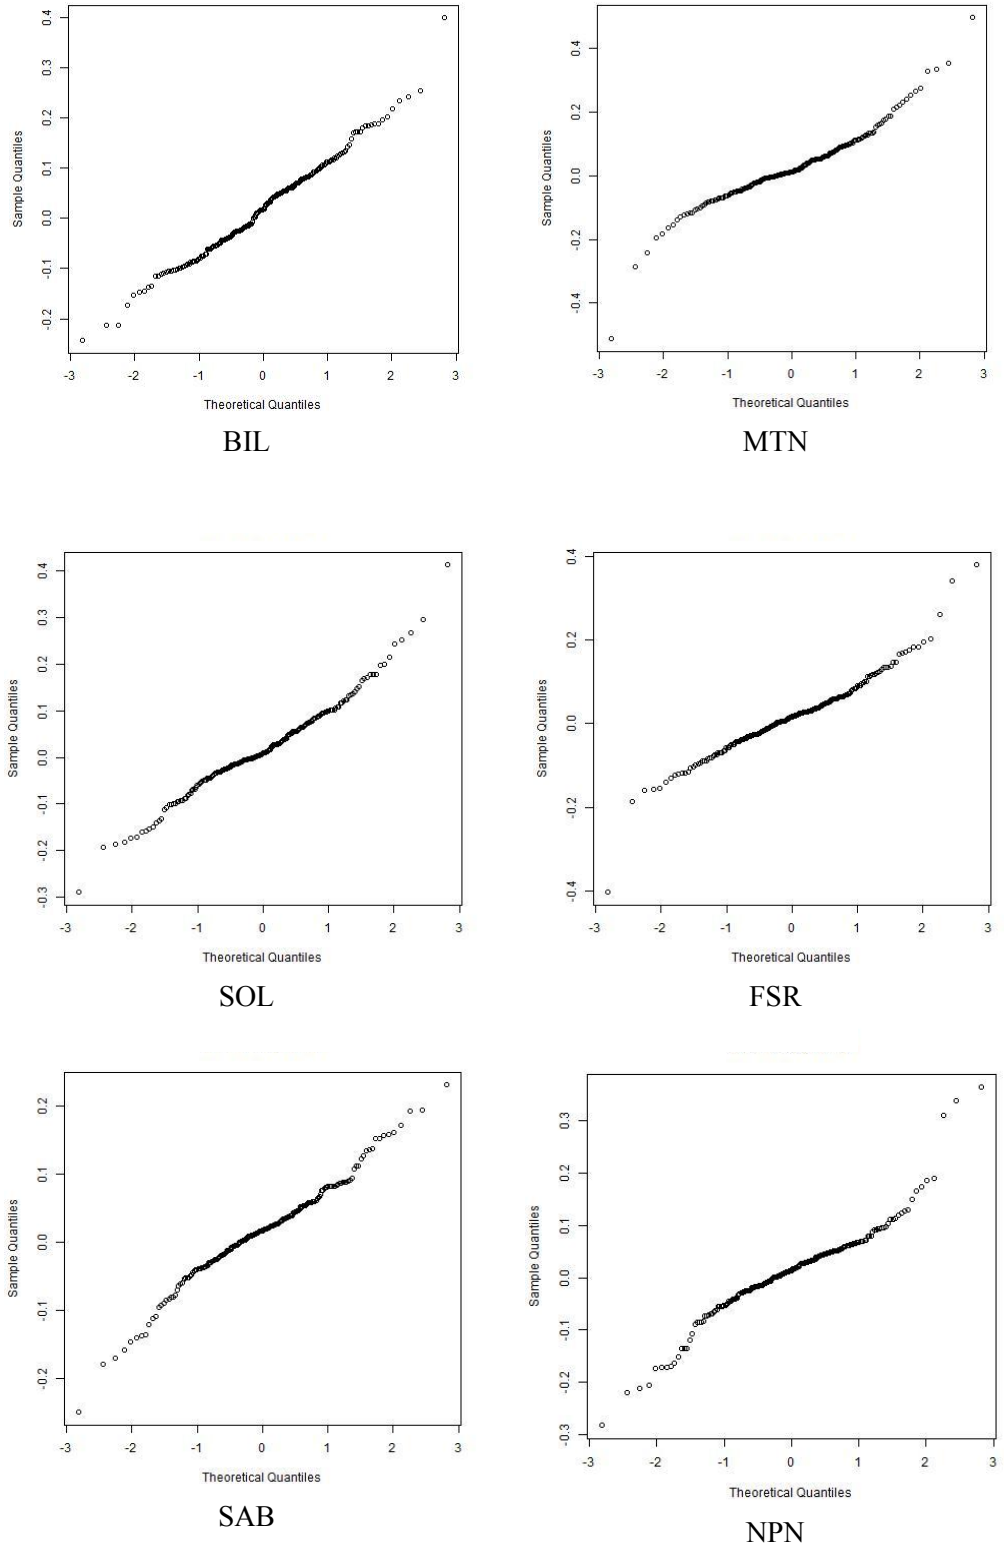
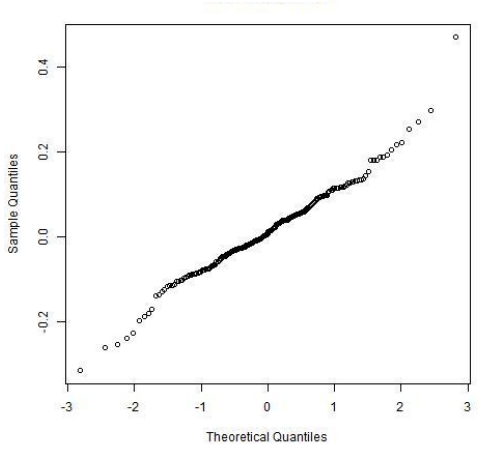
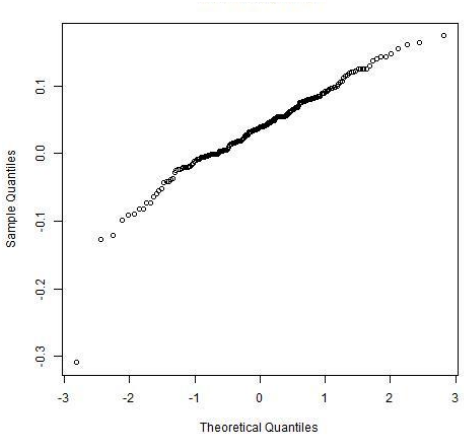


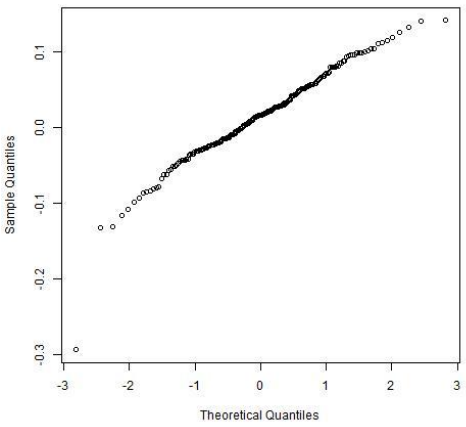
Figure 9 - Q-Q plot of monthly frequency data (1)



AGL



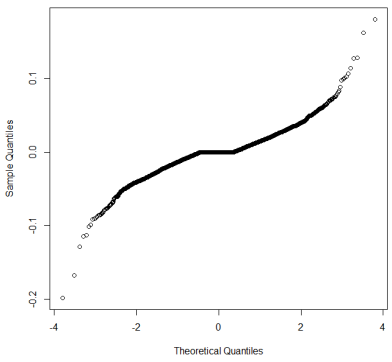
J200



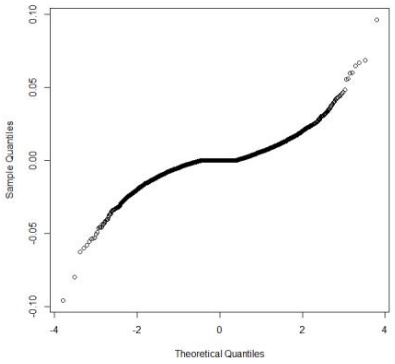
ALSI

Figure 10 - Q-Q plot of monthly frequency data (2)

Examining each of the additional pricing factor Q-Q plots in Figure 11 and Figure 12, one finds that they too do not follow a normal distribution.

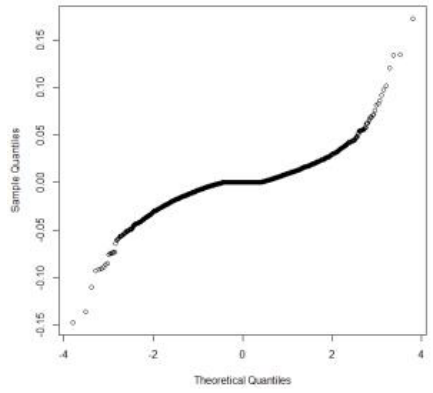


Oil

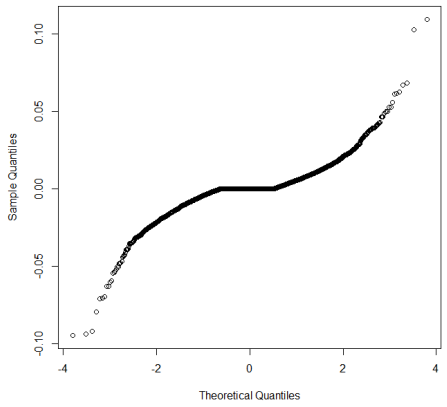


Gold

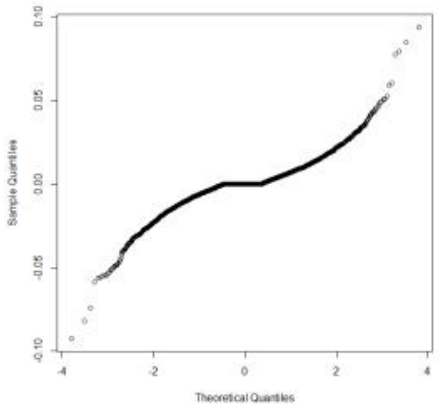
Figure 11 - Q-Q plots for each exogenous variable to be used in the network (1)



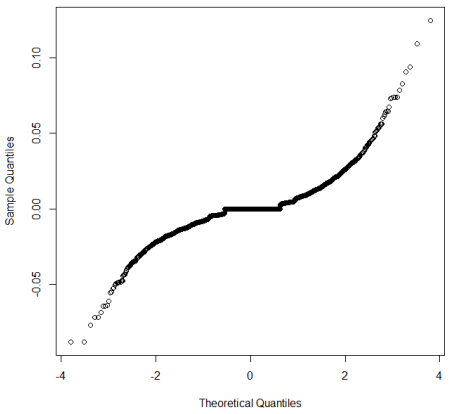
Hang Seng



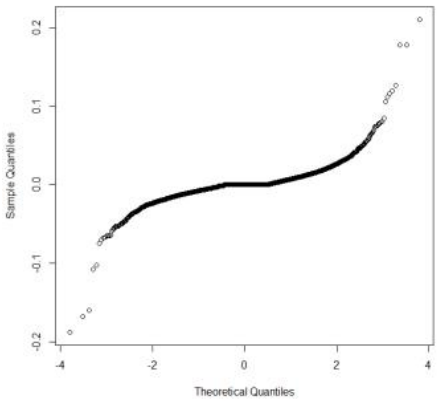
S&P 500



FTSE



ALSI - Dividend Yield



ALSI - Earnings Yield

Figure 12 - Q-Q plots for each exogenous variable to be used in the network (2)

4.2.1.3 Kolmogorov-Smirnov test

The Kolmogorov-Smirnov (K-S) test (also known as the Lilliefors test) is a non-parametric counterpart to the Jarque-Bera test. It is used to describe the difference between the empirical distribution function of the sample and the cumulative distribution function of the reference distribution. The null hypothesis is that both samples are drawn from the same distribution. The results of the K-S test are displayed in Table 14 below. Both the test statistic (D statistic) and the corresponding p-values indicate that all daily equity return series are not drawn from a normal distribution.

Table 14 - K-S results for daily returns

Share Code	D Statistic	P-Value
BIL	0.4785	0.00***
MTN	0.4751	0.00***
SOL	0.4763	0.00***
FSR	0.4798	0.00***
SAB	0.4825	0.00***
NPN	0.4740	0.00***
AGL	0.4765	0.00***
J200	0.4852	0.00***
ALSI	0.4863	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

The results of the K-S test for weekly data are displayed in Table 15 below. Similar to the daily results, there is no evidence of normality.

Table 15 - K-S results for weekly returns

Share Code	D Statistic	P-Value
BIL	0.4521	0.00***
MTN	0.4452	0.00***
SOL	0.4454	0.00***
FSR	0.4506	0.00***
SAB	0.4640	0.00***
NPN	0.4407	0.00***
AGL	0.4508	0.00***
J200	0.4705	0.00***
ALSI	0.4727	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

When examining the monthly return data in Table 16, a similar conclusion to the ones reached above can be drawn – there is no evidence of normality.

Table 16 - K-S results for monthly returns

Share Code	D Statistic	P-Value
BIL	0.4274	0.00***
MTN	0.4195	0.00***
SOL	0.4265	0.00***
FSR	0.4343	0.00***
SAB	0.4316	0.00***
NPN	0.4078	0.00***
AGL	0.4098	0.00***
J200	0.4560	0.00***
ALSI	0.4470	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

A similar conclusion of non-normality under quarterly and semi-annual data can be found in all 50 securities examined. Lastly, the daily frequency variables used as inputs in the neural network also do not display evidence of normality, as shown in Table 17 below.

Table 17 - K-S results for exogenous variables

	Oil	Gold	Hang Seng	S&P	FTSE	DY	EY
D statistic	0.4710	0.4818	0.4762	0.4800	0.4807	0.4828	0.4777
P-Value	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

In summary, considering both visual evidence, parametric and non-parametric tests for normality, it can be concluded that the all of the equities data under daily, weekly and monthly frequencies do not follow a normal distribution. In the case of conflicting results between the Jarque Bera test and the K-S test (such as in some of the mining shares), the difference can be attributed to the non-parametric nature of the latter test. The K-S test determines normality against a randomly generated distribution. If this distribution has a sufficiently large number of observations, then according to the Central Limit Theorem, this distribution is approximately Gaussian. In the case of the monthly returns sample size, the small sample size could have lead to normality being concluded prematurely. Thus, given the superiority of the KS test over the JB test, it can be concluded that none of the equities in the

population are normally distributed. Attention is now drawn to the results of these tests for each non-overlapping sub-sample.

4.2.2 ALSI Sub-sample results

4.2.2.1 Jarque Bera test

The Jarque Bera test for normality in each sub-sample dataset is presented in Table 18 below. The table illustrates the χ^2 test statistic, degrees of freedom and the corresponding p-value. The null hypothesis of normality is rejected if the p-value is statistically significant. The results of the Jarque Bera test indicate that the null hypothesis of normality is rejected at a significance level of 5% (indeed at a level of 1%) for all sub-samples Therefore, according to the Jarque Bera, none of the variables under consideration follow a normal distribution.

Table 18 - Jarque Bera results for each sub-sample

Sub-sample	χ2	D.O.F	P-value
#1	1190.5710	2.00	0.00***
#2	523.1271	2.00	0.00***
#3	98.0930	2.00	0.00***
#4	22.8181	2.00	0.00***
#5	57.6022	2.00	0.00***
#6	190.9274	2.00	0.00***
#7	25.6137	2.00	0.00***
#8	22.9640	2.00	0.00***
#9	71.3668	2.00	0.00***
#10	34.7494	2.00	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

4.2.2.2 Quantile-Quantile plot

The Quantile-Quantile (Q-Q) plot displays two probability distributions by plotting their quantiles against each other. If the two distributions have similar plots, then the result would be a straight line at a 45° angle. Examining the Normal (Q-Q) plots of each sub-sample's data in Figure 13 and Figure 14 against a randomly populated normal distribution, one finds that all return series display departures from normality. Therefore, according to the visual

evidence of the Q-Q plot, none of the sub-samples under consideration follow a normal distribution.

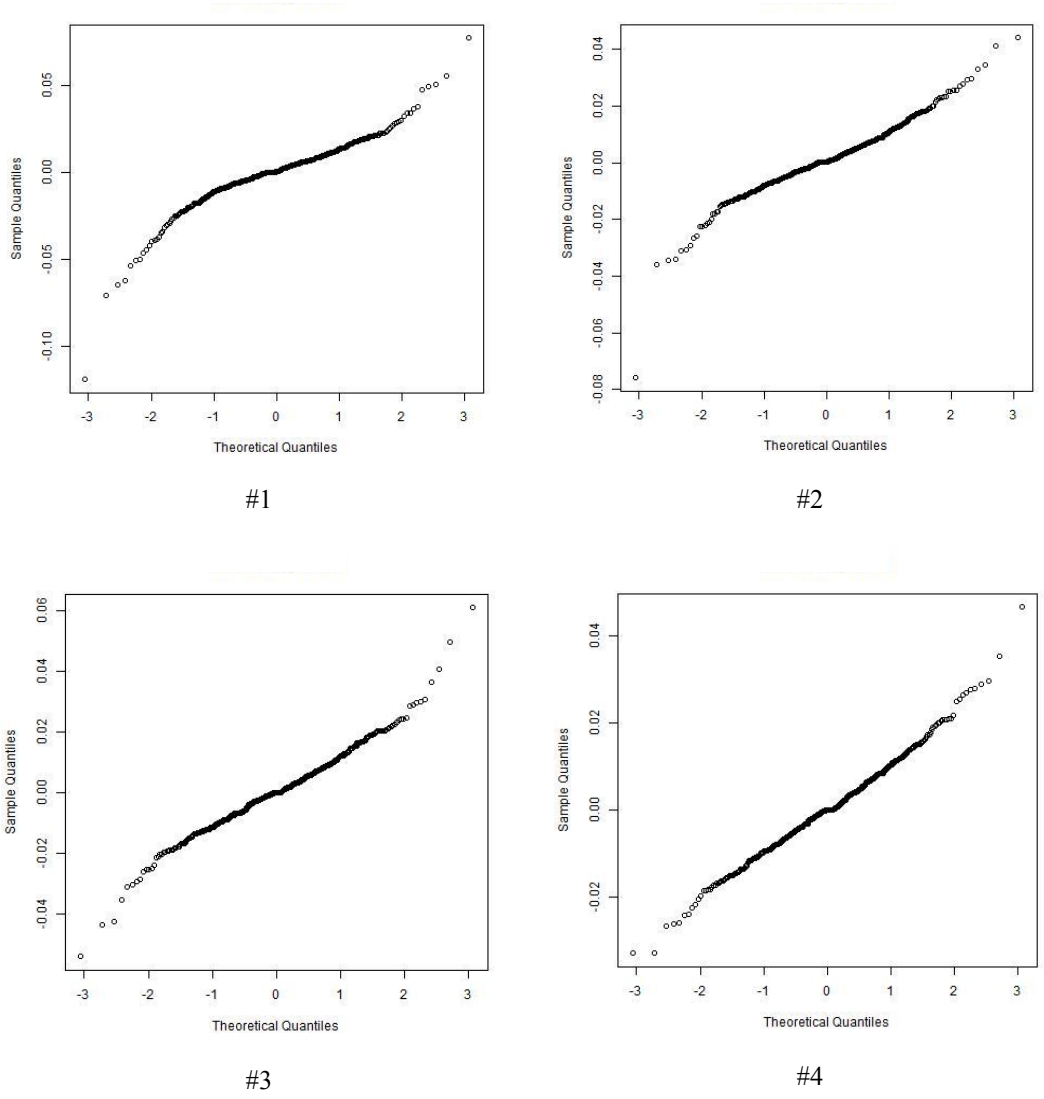
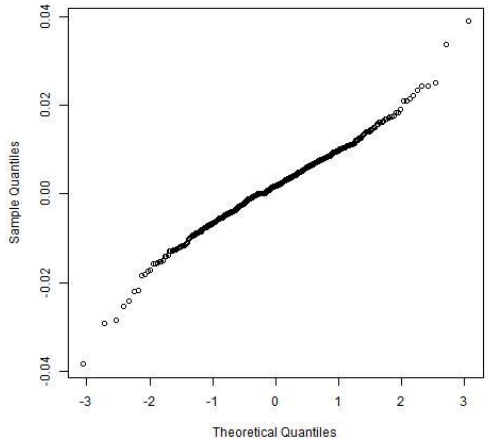
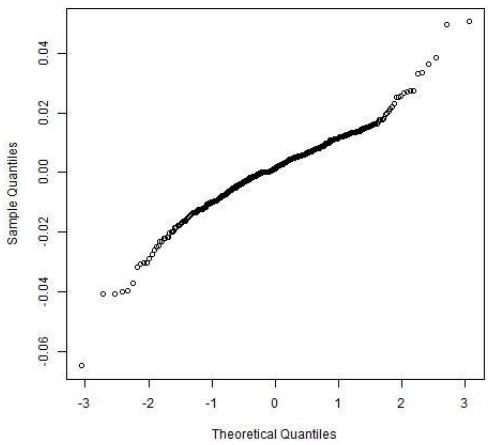


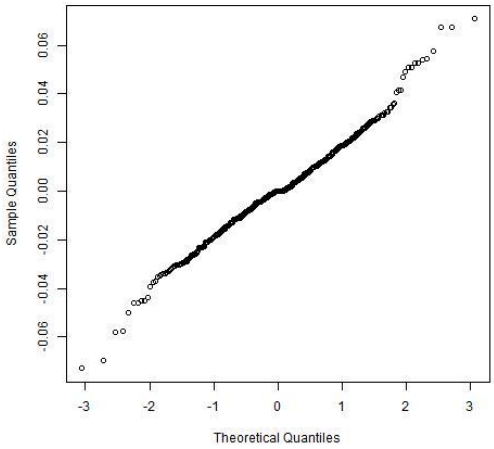
Figure 13 - Q-Q plots for each sub-sample (1)



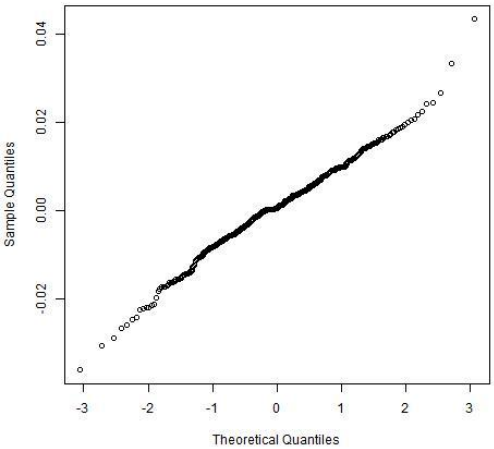
#5



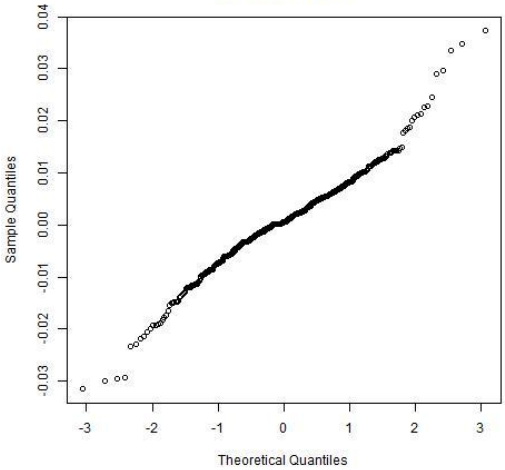
#6



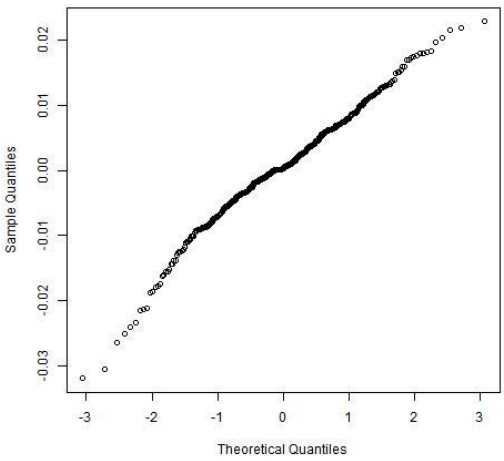
#7



#8



#9



#10

Figure 14 - Q-Q plots for each sub-sample (2)

4.2.2.3 Kolmogorov-Smirnov test

The Kolmogorov-Smirnov (K-S) test is a non-parametric counterpart to the Jarque-Bera test. It is used to describe the difference between the empirical distribution function of the sample and the cumulative distribution function of the reference distribution. The null hypothesis is that both samples are drawn from the same distribution. The results of the K-S test are displayed in Table 19 below. Both the test statistic (D statistic) and corresponding p-value indicate that all return series are not drawn from a normal distribution. Therefore, according to the K-S test, none of the sub-samples under consideration follow a normal distribution.

Table 19 - K-S results for each sub-sample

Sub-sample	D Statistic	P-Value
#1	0.4748	0.00***
#2	0.4854	0.00***
#3	0.4832	0.00***
#4	0.4895	0.00***
#5	0.4891	0.00***
#6	0.4841	0.00***
#7	0.4751	0.00***
#8	0.4883	0.00***
#9	0.4899	0.00***
#10	0.4897	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

In summary, the visual evidence, parametric and non-parametric tests for normality over each non-overlapping sub-sample pointed to the data being drawn from a non-normal distribution. The tests of the random walk hypothesis now begin in earnest, with tests of non-linearity in the return generating process.

4.3 Tests for non-linearity

One test for linearity, the Brock-Dechert-Scheinkman (BDS) test is employed. This test examines linearity around key percentiles of the distribution.

4.3.1 Full sample results

The BDS test examines observations around the 25th percentile, median, 75th percentile and 90th percentile for non-linearity at a lag structure of 2 and 3. At a minimum, the BDS test requires a lag of 2 to be conducted; and a maximum lag of 20. The lag of 3 was chosen as it shows the lowest lag at which the results are statistically significant. The top half of the table provides the percentile values as described above, with the bottom half providing the p-values of those percentiles.

The results of the BDS test show that daily BIL returns (Table 20), daily MTN returns (Table 21) and daily SOL returns (Table 22) exhibit non-linear behaviour as the null hypothesis of linearity is rejected at all common levels of significance.

Table 20 – BDS test for non-linearity of BIL daily returns

BIL	25th percentile	Median	75th percentile	90th percentile
	0.0119	0.0238	0.0357	0.0476
2	8.4002	9.1993	9.8333	10.5223
3	12.126	13.354	14.3111	15.3729

BIL	25th percentile	Median	75th percentile	90th percentile
	0.0119	0.0238	0.0357	0.0476
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 21 – BDS test for non-linearity of MTN daily returns

MTN	25th percentile	Median	75th percentile	90th percentile
	0.013	0.0261	0.0391	0.0522
2	14.2931	15.2998	16.4209	16.8788
3	19.0183	19.3507	19.8789	19.7112

MTN	25th percentile	Median	75th percentile	90th percentile
	0.013	0.0261	0.0391	0.0522
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 22 – BDS test for non-linearity of SOL daily returns

SOL	25th percentile	Median	75th percentile	90th percentile
	0.0116	0.0232	0.0347	0.0463
2	14.163	14.5488	14.4076	14.743
3	17.1412	17.7531	18.0586	18.8018

SOL	25th percentile	Median	75th percentile	90th percentile
	0.0116	0.0232	0.0347	0.0463
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Similar conclusions can be drawn for FSR daily returns (Table 23), SAB daily returns (Table 24), NPN daily returns (Table 25) and AGL daily returns (Table 26) in that there are non-normalities in the data according to the BDS test.

Table 23 – BDS test for non-linearity of FSR daily returns

FSR	25th percentile	Median	75th percentile	90th percentile
	0.0107	0.0213	0.032	0.0426
2	13.6113	14.8994	15.3627	15.1698
3	17.6326	19.0421	19.3718	18.9976

FSR	25th percentile	Median	75th percentile	90th percentile
	0.0107	0.0213	0.032	0.0426
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 24 – BDS test for non-linearity of SAB daily returns

SAB	25th percentile	Median	75th percentile	90th percentile
	0.0089	0.0177	0.0266	0.0355
2	12.358	13.752	14.3582	14.4853
3	14.9321	16.316	16.7457	16.6222

SAB	25th percentile	Median	75th percentile	90th percentile
	0.0089	0.0177	0.0266	0.0355
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 25 – BDS test for non-linearity of NPN daily returns

NPN	25th percentile 0.0124	Median 0.0248	75th percentile 0.0372	90th percentile 0.0497
2	15.5464	16.4342	16.7112	16.0929
3	19.2877	20.01	20.303	20.1143

NPN	25th percentile 0.0124	Median 0.0248	75th percentile 0.0372	90th percentile 0.0497
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 26 – BDS test for non-linearity of AGL daily returns

AGL	25th percentile 0.0122	Median 0.0245	75th percentile 0.0367	90th percentile 0.049
2	9.6824	11.0947	12.7444	13.6597
3	12.4384	14.195	16.3651	17.6314

AGL	25th percentile 0.0122	Median 0.0245	75th percentile 0.0367	90th percentile 0.049
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Lastly, examining the daily returns of the J200 (Table 27), GFI (Table 28) and the ALSI (Table 29), the same conclusion is reached in that there is evidence of non-linearity in all three distributions.

Table 27 – BDS test for non-linearity of J200 daily returns

J200	25th percentile 0.007	Median 0.0139	75th percentile 0.0209	90th percentile 0.0278
2	10.7067	12.3156	13.566	14.0021
3	15.1673	17.4519	19.1856	19.7513

J200	25th percentile 0.007	Median 0.0139	75th percentile 0.0209	90th percentile 0.0278
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 28 – BDS test for non-linearity of GFI daily returns

GFI	25th percentile 0.0147	Median 0.0295	75th percentile 0.0442	90th percentile 0.059
2	10.4435	12.5219	13.9685	14.7544
3	13.5584	15.3725	16.3833	16.5151

GFI	25th percentile 0.0147	Median 0.0295	75th percentile 0.0442	90th percentile 0.059
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 29 – BDS test for non-linearity of ALSI daily returns

ALSI	25th percentile 0.0063	Median 0.0127	75th percentile 0.019	90th percentile 0.0254
2	10.3033	12.1241	13.4984	13.8658
3	15.0979	17.4721	19.3082	19.6816

ALSI	25th percentile 0.0063	Median 0.0127	75th percentile 0.019	90th percentile 0.0254
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

The results of the BDS test show that weekly BIL returns (Table 30), weekly MTN returns (Table 31) and weekly SOL returns (Table 32) exhibit non-linear behaviour as the null hypothesis of linearity is rejected at all common levels of significance.

Table 30 – BDS test for non-linearity of BIL weekly returns

BIL	25th percentile 0.0264	Median 0.0528	75th percentile 0.0792	90th percentile 0.1056
2	4.0587	5.5324	6.4257	7.4672
3	4.2862	5.9743	7.0556	8.0845

BIL	25th percentile 0.0264	Median 0.0528	75th percentile 0.0792	90th percentile 0.1056
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 31 – BDS test for non-linearity of MTN weekly returns

MTN	25th percentile	Median	75th percentile	90th percentile
	0.0300	0.0601	0.0901	0.1202
2	5.6899	5.9294	5.6904	5.2326
3	7.1232	7.4998	7.1586	6.7512

MTN	25th percentile	Median	75th percentile	90th percentile
	0.0300	0.0601	0.0901	0.1202
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 32 – BDS test for non-linearity of SOL weekly returns

SOL	25th percentile	Median	75th percentile	90th percentile
	0.0266	0.0531	0.0797	0.1062
2	2.596	3.4924	4.617	6.2227
3	4.7227	5.3347	5.8603	7.1449

SOL	25th percentile	Median	75th percentile	90th percentile
	0.0254	0.0508	0.0762	0.1016
2	0.01***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Similar conclusions can be drawn for FSR weekly returns (Table 33), SAB weekly returns (Table 33), NPN weekly returns (Table 35) and AGL weekly returns (Table 36) in that there are non-normalities in the data according to the BDS test.

Table 33 – BDS test for non-linearity of FSR weekly returns

FSR	25th percentile	Median	75th percentile	90th percentile
	0.0244	0.0488	0.0733	0.0977
2	4.3307	6.2032	8.0879	9.1905
3	5.3114	7.1288	8.8976	10.0861

FSR	25th percentile	Median	75th percentile	90th percentile
	0.0244	0.0488	0.0733	0.0977
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 34 – BDS test for non-linearity of MDC weekly returns

SAB	25th percentile	Median	75th percentile	90th percentile
	0.0197	0.0394	0.0591	0.0788
2	4.1059	5.5013	5.9267	5.006
3	5.9955	7.3956	7.9474	7.2162

SAB	25th percentile	Median	75th percentile	90th percentile
	0.0197	0.0394	0.0591	0.0788
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 35 – BDS test for non-linearity of PIK weekly returns

NPN	25th percentile	Median	75th percentile	90th percentile
	0.0292	0.0584	0.0875	0.1167
2	5.7037	6.5977	7.1349	7.2322
3	7.032	8.1127	8.409	8.3712

NPN	25th percentile	Median	75th percentile	90th percentile
	0.0292	0.0584	0.0875	0.1167
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 36 – BDS test for non-linearity of HYP weekly returns

AGL	25th percentile	Median	75th percentile	90th percentile
	0.0274	0.0548	0.0823	0.1097
2	3.4574	4.3606	5.612	6.7464
3	4.7389	5.7518	6.8904	7.5565

AGL	25th percentile	Median	75th percentile	90th percentile
	0.0292	0.0584	0.0875	0.1167
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Lastly, examining the weekly returns of the J200 (Table 37), GFI (Table 38) and the ALSI (Table 39), the same conclusion is reached in that there is evidence of non-linearity in all three distributions. While the 25th percentile of ILV returns is only significant at the 5% level

(instead of the 1% level achieved by other percentiles), the conclusion of non-linearity remains.

Table 37 – BDS test for non-linearity of ILV weekly returns

J200	25th percentile	Median	75th percentile	90th percentile
	0.0156	0.0311	0.0467	0.0623
2	4.1644	4.5566	5.6532	6.6236
3	6.443	5.9418	6.7844	7.2664

J200	25th percentile	Median	75th percentile	90th percentile
	0.0292	0.0584	0.0875	0.1167
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 38 – BDS test for non-linearity of WBO weekly returns

GFI	25th percentile	Median	75th percentile	90th percentile
	0.034	0.068	0.102	0.136
2	3.1684	4.3355	5.6391	6.9756
3	3.7893	5.1029	6.191	7.3617

GFI	25th percentile	Median	75th percentile	90th percentile
	0.034	0.068	0.102	0.136
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 39 – BDS test for non-linearity of ALSI weekly returns

ALSI	25th percentile	Median	75th percentile	90th percentile
	0.0144	0.0288	0.0433	0.0577
2	4.093	4.7478	5.6397	6.5969
3	6.1041	5.9944	6.5691	7.0093

ALSI	25th percentile	Median	75th percentile	90th percentile
	0.0144	0.0288	0.0433	0.0577
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

The results of the BDS test show that monthly BIL returns (Table 40), monthly MTN returns (Table 41) and monthly SOL returns (Table 42) exhibit non-linear behaviour as the null hypothesis of linearity is rejected at all common levels of significance.

Table 40 – BDS test for non-linearity of BIL monthly returns

BIL	25th percentile	Median	75th percentile	90th percentile
	0.0484	0.0968	0.1452	0.1936
2	5.4582	3.5637	3.8127	3.2908
3	7.4697	4.7495	4.8581	4.4573

BIL	25th percentile	Median	75th percentile	90th percentile
	0.0484	0.0968	0.1452	0.1936
2	0.00***	0.00***	0.00***	0.001
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 41 – BDS test for non-linearity of MTN monthly returns

MTN	25th percentile	Median	75th percentile	90th percentile
	0.0546	0.1092	0.1638	0.2183
2	4.328	4.8374	5.1284	5.3886
3	4.0142	5.4697	5.6889	6.1249

MTN	25th percentile	Median	75th percentile	90th percentile
	0.0546	0.1092	0.1638	0.2183
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 42 – BDS test for non-linearity of SOL monthly returns

SOL	25th percentile	Median	75th percentile	90th percentile
	0.047	0.094	0.141	0.1879
2	3.0827	2.9635	3.2846	3.3857
3	2.484	2.6562	2.9824	3.2978

SOL	25th percentile	Median	75th percentile	90th percentile
	0.0447	0.0895	0.1342	0.1789
2	0.00***	0.00***	0.00***	0.00***
3	0.01***	0.01***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

The FSR monthly returns (Table 43) produce interesting results. At a lag of 2, the lower two quantiles (25th and 50th) show significant non-linearity, whereas the upper two quantiles do not. In contrast, at a lag of 3, non-linearity is present throughout the distribution. This result is quite interesting as it implies that the return distribution for FSR is non-linear when returns fall below the mean and linear when returns lie above the mean. It suggests that both a linear and non-linear model should be used when examining the monthly returns generating process. Under the constraints of traditional, *a priori* models, one would need to first determine where this "structural break" occurred before proceeding to model these returns. The monthly SAB returns (Table 44) show linearity for the most part. It is only at a lag of 3 that there is non-linearity in the returns series, however this is not statistically strong. The remaining equity returns of PIK (Table 45) and HYP (Table 46) show evidence of non-linearity.

Table 43 – BDS test for non-linearity of FSR monthly returns

FSR	25th percentile	Median	75th percentile	90th percentile
	0.044	0.0879	0.1319	0.1759
2	2.0661	1.874	1.3339	1.105
3	4.083	3.2033	2.6645	2.5203

FSR	25th percentile	Median	75th percentile	90th percentile
	0.044	0.0879	0.1319	0.1759
2	0.04**	0.06*	0.18	0.27
3	0.00***	0.00**	0.01**	0.01**

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 44 – BDS test for non-linearity of SAB monthly returns

SAB	25th percentile	Median	75th percentile	90th percentile
	0.0351	0.0702	0.1052	0.1403
2	1.1637	0.9844	1.2845	1.3668
3	1.8396	1.6193	2.1003	2.4691

SAB	25th percentile	Median	75th percentile	90th percentile
	0.0351	0.0702	0.1052	0.1403
2	0.24	0.32	0.20	0.17
3	0.07*	0.11	0.04**	0.01**

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 45 – BDS test for non-linearity of NPN monthly returns

NPN	25th percentile	Median	75th percentile	90th percentile
	0.0613	0.1227	0.184	0.2454
2	4.4232	4.3691	4.0739	3.3535
3	6.4474	5.8917	6.0759	5.0606

NPN	25th percentile	Median	75th percentile	90th percentile
	0.0351	0.0702	0.1052	0.1403
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 46 – BDS test for non-linearity of AGL monthly returns

AGL	25th percentile	Median	75th percentile	90th percentile
	0.052	0.104	0.156	0.208
2	0.5434	0.6202	0.6227	0.7973
3	2.1024	1.6479	1.5662	1.7088

AGL	25th percentile	Median	75th percentile	90th percentile
	0.0351	0.0702	0.1052	0.1403
2	0.59	0.54	0.53	0.43
3	0.04**	0.10*	0.12	0.09*

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Lastly, examining the monthly returns of the J200 (Table 47), GFI (Table 48) and the ALSI (Table 49). All three display evidence of non-linearity. In particular, it appears that the returns above the mean for the ALSI show weaker evidence of non-linearity than returns below the mean.

Table 47 – BDS test for non-linearity of J200 monthly returns

J200	25th percentile	Median	75th percentile	90th percentile
	0.0303	0.0606	0.0908	0.1211
2	3.247	3.0268	1.9848	1.551
3	7.1257	5.6793	4.6553	4.0452

J200	25th percentile	Median	75th percentile	90th percentile
	0.0351	0.0702	0.1052	0.1403
2	0.00***	0.00***	0.05*	0.12
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 48 – BDS test for non-linearity of GFI monthly returns

GFI	25th percentile	Median	75th percentile	90th percentile
	0.0588	0.1175	0.1763	0.2351
2	2.2641	2.6922	3.1084	2.6192
3	2.3812	2.27	2.4494	2.0114

GFI	25th percentile	Median	75th percentile	90th percentile
	0.0588	0.1175	0.1763	0.2351
2	0.02**	0.01***	0.00***	0.01***
3	0.02**	0.02**	0.01***	0.04**

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 49 – BDS test for non-linearity of ALSI monthly returns

ALSI	25th percentile	Median	75th percentile	90th percentile
	0.0287	0.0574	0.0861	0.1147
2	3.9638	3.2759	2.1868	1.4318
3	8.0483	6.1382	4.6461	3.8241

ALSI	25th percentile	Median	75th percentile	90th percentile
	0.0287	0.0574	0.0861	0.1147
2	0.00***	0.00***	0.03**	0.15
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

The BDS test for non-linearity is now run on each of the daily frequency input variables. The results for the commodities, Oil and Gold, are presented in Table 50 and Table 51 below. They indicate that both series have non-linear components as the p-values for all variables are zero, implying a rejection of the null hypothesis of linearity.

Table 50 – BDS test for non-linearity on Oil

OIL	25^{th} percentile	Median	75^{th} percentile	90^{th} percentile
	0.0107	0.0215	0.0322	0.0429
2	8.7146	8.7922	9.4258	10.3393
3	11.6335	11.245	12.1772	13.7132

OIL	25^{th} percentile	Median	75^{th} percentile	90^{th} percentile
	0.0107	0.0215	0.0322	0.0429
2	0.000***	0.000***	0.000***	0.000***
3	0.000***	0.000***	0.000***	0.000***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 51 - BDS test for non-linearity on Gold

GOLD	25th percentile 0.0057	Median 0.0115	75th percentile 0.0172	90th percentile 0.0229
2	10.6042	12.1133	13.2069	14.0988
3	12.4524	14.9178	16.1743	16.5816

GOLD	25th percentile 0.0057	Median 0.0115	75th percentile 0.0172	90th percentile 0.0229
2	0.000***	0.000***	0.000***	0.000***
3	0.000***	0.000***	0.000***	0.000***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

The results for the ALSI dividend yield earnings yield are presented in Tables 52 and 53 below. Each variable shows that there exists non-linearity in the series.

Table 52 - BDS Test results for the ALSI Dividend Yield

DY	25th percentile 0.0057	Median 0.0113	75th percentile 0.0170	90th percentile 0.0270
2	14.0581	12.8392	10.4675	12.596
3	9.9397	12.2369	15.3878	13.7962

DY	25th percentile 0.0057	Median 0.0113	75th percentile 0.0170	90th percentile 0.0270
2	0.000***	0.000***	0.000***	0.000***
3	0.000***	0.000***	0.000***	0.000***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 53 - BDS Test results for the ALSI Earnings Yield

EY	25th percentile 0.0080	Median 0.0160	75th percentile 0.0240	90th percentile 0.0320
2	10.4938	11.1411	10.4675	9.4944
3	15.1129	16.1595	15.3878	13.8553

EY	25th percentile 0.0080	Median 0.0160	75th percentile 0.0240	90th percentile 0.0320
2	0.000***	0.000***	0.000***	0.000***
3	0.000***	0.000***	0.000***	0.000***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

The results for international indices, the S&P 500, Hang Seng 100 and FTSE 100, are presented in Tables 54, 55 and 56 below. The results from the test indicate that the series has

some non-linear components as the p-values for all variables are zero, implying a rejection of the null hypothesis of linearity.

Table 54 - BDS Test results for S&P 500

S&P	25th percentile 0.0063	Median 0.0126	75th percentile 0.0189	90th percentile 0.0252
2	8.9495	10.279	12.5075	15.2412
3	15.347	16.3849	17.7273	19.9918

S&P	25th percentile 0.0063	Median 0.0126	75th percentile 0.0189	90th percentile 0.0252
2	0.000***	0.000***	0.000***	0.000***
3	0.000***	0.000***	0.000***	0.000***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 55 - BDS Test results for the Hang Seng

Hang Seng	25th percentile 0.0084	Median 0.0169	75th percentile 0.0253	90th percentile 0.0338
2	7.5547	10.2474	12.5962	14.4866
3	10.5999	14.4438	16.8918	18.4769

Hang Seng	25th percentile 0.0084	Median 0.0169	75th percentile 0.0253	90th percentile 0.0338
2	0.000***	0.000***	0.000***	0.000***
3	0.000***	0.000***	0.000***	0.000***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 56 - BDS Test results for the FTSE

FTSE	25th percentile 0.0061	Median 0.0122	75th percentile 0.0183	90th percentile 0.0244
2	13.3668	14.4039	15.3884	15.9411
3	18.6776	19.5417	20.119	20.7226

FTSE	25th percentile 0.0061	Median 0.0122	75th percentile 0.0183	90th percentile 0.0244
2	0.000***	0.000***	0.000***	0.000***
3	0.000***	0.000***	0.000***	0.000***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Therefore, according to the BDS test, all of the variables under consideration follow some form of non-linear data generating process. The tests for non-linearity are now conducted on each sub-sample.

4.3.2 ALSI Sub-sample results

The BDS test examines observations around the 25th percentile, median, 75th percentile and 90th percentile for non-linearity at a lag structure of 2 and 3. The top half of the table provides the percentile values as described above, with the bottom half providing the p-values of those percentiles.

The results of the BDS test (Table 57) show for sub-sample 1 that the series is non-linear, as each BDS statistic is statistically significant at the 5% level of significance.

Table 57 - BDS Test results for the sub-sample 1

#1	25th percentile	Median	75th percentile	90th percentile
	0.0085	0.017	0.0255	0.034
2	5.0347	4.7872	4.6635	4.5558
3	8.1186	7.4139	6.5083	5.626

#1	25th percentile	Median	75th percentile	90th percentile
	0.0085	0.017	0.0255	0.034
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

However, when examining the second sub-sample (Table 58), it is found that the series is linear at a lag of 2 for the 25th percentile, but is non-linear for the remaining quintiles and lags.

Table 58 - BDS Test results for the sub-sample 2

#2	25th percentile	Median	75th percentile	90th percentile
	0.0057	0.0115	0.0172	0.023
2	1.5147	2.5792	3.6968	5.2205
3	1.791	2.6283	3.9164	5.1564

#2	25th percentile	Median	75th percentile	90th percentile
	0.0057	0.0115	0.0172	0.023
2	0.13	0.01**	0.00***	0.00***
3	0.07*	0.01**	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

For sub-sample 3 (Table 59) and sub-sample 4 (Table 60); the former is found to be non-linear, with the latter being linear.

Table 59 - BDS Test results for the sub-sample 3

#3	25th percentile	Median	75th percentile	90th percentile
	0.0063	0.0127	0.019	0.0254
2	2.3244	2.9062	3.1386	3.275
3	3.0965	4.0206	4.0778	4.0392

#3	25th percentile	Median	75th percentile	90th percentile
	0.0063	0.0127	0.019	0.0254
2	0.02**	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 60 - BDS Test results for the sub-sample 4

#4	25th percentile	Median	75th percentile	90th percentile
	0.0053	0.0106	0.016	0.0213
2	0.5109	0.3161	0.0308	-0.0801
3	1.3836	0.6568	0.2101	-0.1055

#4	25th percentile	Median	75th percentile	90th percentile
	0.0053	0.0106	0.016	0.0213
2	0.61	0.75	0.98	0.94
3	0.17	0.51	0.83	0.92

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

The results for sub-sample 5 (Table 61) are interesting in that the majority of the test statistics are insignificant, but there is some hint of significant non-linearity at a lag of 3.

Table 61 - BDS Test results for the sub-sample 5

#5	25th percentile	Median	75th percentile	90th percentile
	0.0045	0.0091	0.0136	0.0182
2	1.316	0.8975	0.5933	0.0168
3	3.795	2.6413	1.6555	0.9235

#5	25th percentile	Median	75th percentile	90th percentile
	0.0045	0.0091	0.0136	0.0182
2	0.19	0.37	0.55	0.99
3	0.00***	0.01***	0.10*	0.36

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

The BDS test for the remaining sub-samples shows that sub-sample 6 (Table 62) to sub-sample 9 (Table 65) display evidence of non-linearity (although somewhat weaker in evidence for sub-sample 7 and sub-sample 9).

Table 62 - BDS Test results for the sub-sample 6

#6	25th percentile	Median	75th percentile	90th percentile
	0.0063	0.0126	0.0188	0.0251
2	3.4406	3.7493	3.4475	3.1801
3	4.8201	5.5963	5.9115	5.9394

#6	25th percentile	Median	75th percentile	90th percentile
	0.0063	0.0126	0.0188	0.0251
2	0.00***	0.00***	0.00***	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 63 - BDS Test results for the sub-sample 7

#7	25th percentile	Median	75th percentile	90th percentile
	0.0102	0.0204	0.0306	0.0408
2	1.9723	2.3699	2.4129	2.4313
3	3.9487	4.9003	5.0655	4.9032

#7	25th percentile	Median	75th percentile	90th percentile
	0.0102	0.0204	0.0306	0.0408
2	0.05*	0.02**	0.02**	0.02**
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 64 - BDS Test results for the sub-sample 8

#8	25th percentile	Median	75th percentile	90th percentile
	0.005	0.0101	0.0151	0.0201
2	3.206	3.079	2.7082	2.8991
3	3.8343	3.8436	3.6062	3.8647

#8	25th percentile	Median	75th percentile	90th percentile
	0.005	0.0101	0.0151	0.0201
2	0.00***	0.00***	0.01**	0.00***
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 65 - BDS Test results for the sub-sample 9

#9	25th percentile	Median	75th percentile	90th percentile
	0.0046	0.0092	0.0138	0.0184
2	2.7114	3.5751	2.7441	1.7464
3	4.5361	5.1548	4.5489	3.8499

#9	25th percentile	Median	75th percentile	90th percentile
	0.0046	0.0092	0.0138	0.0184
2	0.01**	0.00***	0.01**	0.08*
3	0.00***	0.00***	0.00***	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

The results of sub-sample 10 (Table 66) are largely in favour of linearity, with the exception at the 10% level of the 75th percentile test statistic at a lag of 3. While this particular statistic is significant, it is not enough to conclude that the series is non-linear.

Table 66 - BDS Test results for the sub-sample 10

#10	25th percentile	Median	75th percentile	90th percentile
	0.0042	0.0084	0.0126	0.0167
2	-0.0974	0.3926	0.3619	-0.0218
3	1.1141	1.4259	1.7263	1.3085

#10	25th percentile	Median	75th percentile	90th percentile
	0.0042	0.0084	0.0126	0.0167
2	0.92	0.69	0.72	0.98
3	0.27	0.15	0.08*	0.19

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Therefore, according to the BDS test, the majority of the sub-samples under consideration do not follow some non-linear data generating process, implying that linear models would be more adept at capturing their behaviour than non-linear models. In contrast to the overall sample results, it was discovered that while the overall sample may follow a non-linear data generating process, the majority of time periods within the overall sample follow a linear data generating process. This implies that some points in time that act as regime changers - considered turning points in the series. It is quite plausible that the data around these regime changing points exhibit non-linear behaviour, with the data further from these points exhibiting linear behaviour. These results imply that some form of regime switching models should be used to model the returns generating process. Before commencing however, one should first inspect the series for autocorrelation if any time series model is to be used.

4.4 Tests for stationarity

Two complementary tests for stationarity are conducted, namely the Augmented Dickey Fuller (ADF) and the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test. The former tests for the presence of a unit root (which has implications on stationarity), while the latter examines a time series as being composed of multiple components, some of which may not be stationary. Further, a test conducted by Lo (2004), in constructing his argument for cyclical efficiency, is replicated here. This is a simple measure of autocorrelation over time, using a rolling window approach. The graphical inspection of the diagram can reveal points in time where returns were not independent of each other (thus having implications on efficiency) as well as provide robustness to the overall preliminary statistics in adopting an overlapping sample approach.

4.4.1 Full sample results

4.4.1.1 Augmented Dickey Fuller test

The Augmented Dickey Fuller (ADF) test for stationarity was conducted on the ALSI. The test statistic is usually a negative number, with a larger negative number representing a stronger rejection of the null hypothesis of a unit root in the series. Further, the test is

conducted for a mean and trend in the series, to determine if any form of the described stationarity exists. The results from Table 67 to Table 69 show that the ALSI is stationary at the lags tested and at the 1% level of statistical significance. In other words, the daily, weekly and monthly returns data do not show signs of non-stationarity.

Table 67 - ADF Test Results for daily returns

	Test Statistic	Lag	P-value
BIL	-17.3255	16	0.01***
MTN	-16.5742	16	0.01***
SOL	-17.2609	16	0.01***
FSR	-17.0675	16	0.01***
SAB	-16.6482	16	0.01***
NPN	-14.5267	16	0.01***
AGL	-16.5913	16	0.01***
J200	-16.0414	16	0.01***
ALSI	-15.6333	16	0.01***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 68 - ADF Test Results for weekly returns

	Test Statistic	Lag	P-value
BIL	-9.6735	9	0.01***
MTN	-9.4568	9	0.01***
SOL	-10.0748	9	0.01***
FSR	-11.0275	9	0.01***
SAB	-10.0548	9	0.01***
NPN	-10.1591	9	0.01***
AGL	-9.4168	9	0.01***
J200	-9.6489	9	0.01***
ALSI	-9.7272	9	0.01***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 69 - ADF Test Results for monthly returns

	Test Statistic	Lag	P-value
BIL	-6.7246	5	0.01***
MTN	-5.9674	5	0.01***
SOL	-6.021	5	0.01***
FSR	-7.2153	5	0.01***
SAB	-5.4231	5	0.01***
NPN	-5.8839	5	0.01***
AGL	-5.8224	5	0.01***
J200	-5.9984	5	0.01***
ALSI	-6.0981	5	0.01***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Similarly, when one considers quarterly and semi-annual data, the results indicate stationarity. The results for the daily frequency input variables are displayed in Table 70 below. Using the mean and trend version of the test, it is found that all exogenous return variables are stationary.

Table 70 - ADF Test Results for exogenous variables

	Test Statistic	Lag	P-value
Gold	-16.2804	16	0.01***
Oil	-14.7765	16	0.01***
FTSE	-16.412	16	0.01***
HS	-15.5022	16	0.01***
SP	-15.8228	16	0.01***
EY	-15.7285	16	0.01***
DY	-15.7285	16	0.01***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

4.4.1.2 Kwiatkowski-Phillips-Schmidt-Shin test

The Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test is used to test the null hypothesis that a time series is stationary around a deterministic trend. The time series is decomposed into a deterministic trend, random walk and stationary error component and a Lagrange multiplier method is used to test the hypothesis that the random walk component has a zero variance. The KPSS test thus supplements the ADF test in that both test for a unit root and stationarity.

The results of the KPSS test for the daily, weekly and monthly equity returns are displayed in Table 71 to Table 73 below. For all frequencies tested, the KPSS test shows that the null hypothesis of stationarity is not rejected, with the truncated parameter (T.P) being optimally chosen so as to compromise between the sample size and statistical power of the test. Further, each of the daily frequency input variables (Table 74) is also stationary at the 10% level of significance.

Table 71 - KPSS results for daily returns

Share Code	Test statistic value	T.P	P-Value
BIL	0.1227	15	0.1*
MTN	0.099	15	0.1*
SOL	0.0545	15	0.1*
FSR	0.0462	15	0.1*
SAB	0.1715	15	0.1*
NPN	0.2226	15	0.1*
AGL	0.1488	15	0.1*
J200	0.0412	15	0.1*
ALSI	0.0481	15	0.1*

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 72 - KPSS results for weekly returns

Share Code	Test statistic value	T.P	P-Value
BIL	0.126	6	0.1*
MTN	0.1079	6	0.1*
SOL	0.0631	6	0.1*
FSR	0.0351	6	0.1*
SAB	0.1586	6	0.1*
NPN	0.1453	6	0.1*
AGL	0.1812	6	0.1*
J200	0.0392	6	0.1*
ALSI	0.0436	6	0.1*

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Table 73 - KPSS results for monthly returns

Share Code	Test statistic value	T.P	P-Value
BIL	0.1353	3	0.1*
MTN	0.1084	3	0.1*
SOL	0.0735	3	0.1*
FSR	0.0624	3	0.1*
SAB	0.2269	3	0.1*
NPN	0.1491	3	0.1*
AGL	0.1764	3	0.1*
J200	0.0541	3	0.1*
ALSI	0.0479	3	0.1*

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

The results hold for quarterly and semi-annual data, in that all of the securities are found to be stationary.

Table 74 - KPSS results for exogenous variables

	Oil	Gold	Hang Seng	S&P	FTSE	DY	EY
Test statistic value	0.0779	0.1974	0.0587	0.1173	0.0547	0.0177	0.0577
T.P	15	15	15	15	15	15	15
P-Value	0.1*	0.1*	0.1*	0.1*	0.1*	0.1*	0.1*

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Therefore, according to the KPSS test, all of the variables under consideration are stationary at a 10% level of statistical significance.

4.4.1.3 Rolling autocorrelation test

Under the random walk hypothesis, returns should not be correlated; the autocorrelation coefficient should be zero. Thus, Lo (2004) adopts a rolling window approach to test this premise. By a graphical view of the rolling first order autocorrelation (the autocorrelation coefficient calculated over a daily rolling window), the author concludes that over the sample period investigated; there were phases of inefficiency and efficiency. Replicating this approach for the daily, weekly and monthly ALSI returns, the results are displayed in Figure 15 below. The graphs show the plot of each time series' autocorrelation coefficient return

against time. Under conditions of market efficiency, the autocorrelation coefficient across time should be close to zero.

Using daily returns, the graph shows that the rolling autocorrelation coefficient fluctuates around zero, implying that there are small periods of time where the ALSI experiences market inefficiency. Using weekly returns and monthly returns, the results are similar, showing volatility in the autocorrelation coefficient with a cyclical pattern. This therefore implies that over the sample period examined, the ALSI experienced periods of weak form efficiency and periods of weak form inefficiency. As can be expected with lower frequency data, extending the method to quarterly and semi-annual data, the evidence of randomness becomes more clear (there is less of a pattern in the ALSI returns). While not a test in and of itself (one needs to determine corresponding p-values to ascertain whether a deviation from zero is statistically significant), the rolling autocorrelation plots replicate the work of Lo (2004, 2005) in providing the foundation to examine market efficiency in a dynamic context.

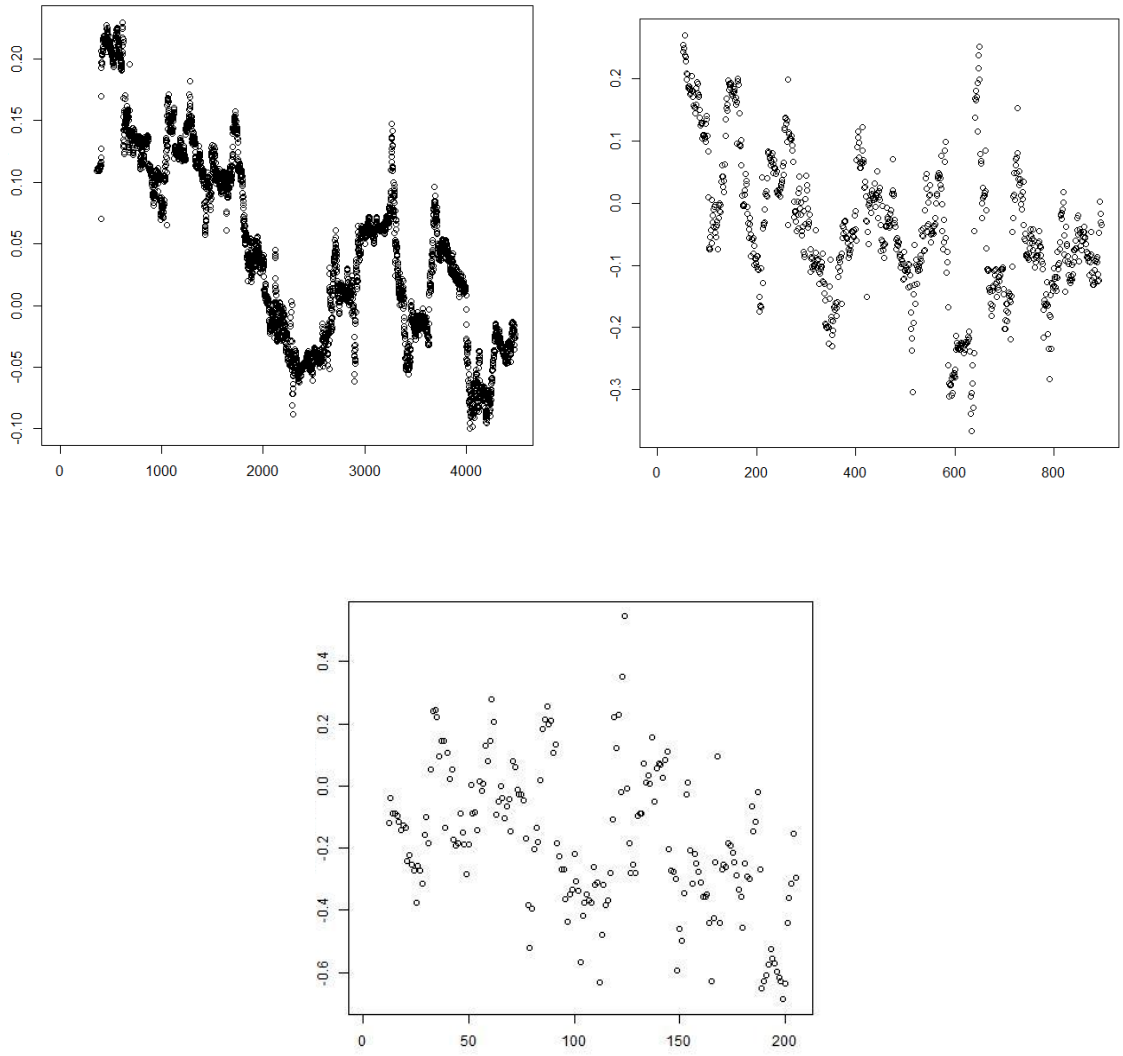


Figure 15 - Rolling autocorrelation at lag 1 for daily (top right), weekly (top left) and monthly (centre) ALSI returns

Therefore, according to a rolling autocorrelation test, the ALSI experienced periods of market efficiency and periods of market inefficiency, across all frequencies examined. Indeed, the monthly rolling autocorrelation figure seems to correspond roughly to that of a business cycle, with the more recent observations (near the global recession) being autocorrelated. While not a strict statistical procedure in and of itself, the rolling autocorrelation graphs provide indirect evidence of cyclical efficiency. In other words, it has been found that there exists some form of memory in the daily, weekly and monthly ALSI series over the sample period under investigation. Further, this memory does not decay rapidly to zero and is persistent. The stationarity tests are now conducted on each sub-sample.

4.4.2 ALSI Sub-sample results

4.4.2.1 Augmented Dickey Fuller test

For each of the sub-samples, the Augmented Dickey Fuller (ADF) test for stationarity was conducted. The test statistic is usually a negative number, with a larger negative number representing a stronger rejection of the null hypothesis of a unit root in the series. Further, the test is conducted for a constant, linear trend in the series, to determine if any form of the described stationarity exists. The results from Table 75 show that the all return sub-samples are stationary at the lags tested and at the 1% level of statistical significance.

Table 75 - ADF results for each sub-sample

Sub-sample	Test Statistic	Lag	P-Value
#1	-7.1274	7	0.01***
#2	-7.1597	7	0.01***
#3	-6.333	7	0.01***
#4	-7.9652	7	0.01***
#5	-7.6637	7	0.01***
#6	-8.2051	7	0.01***
#7	-7.853	7	0.01***
#8	-8.7232	7	0.01***
#9	-8.7456	7	0.01***
#10	-7.3046	7	0.01***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

Therefore, according to the ADF test under a zero mean, single mean and trend, all of the sub-samples under consideration are stationary.

4.4.2.2 Kwiatkowski-Phillips-Schmidt-Shin test

The Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test is used to test the null hypothesis that a time series is stationary around a deterministic trend. The time series is decomposed into a deterministic trend, random walk and stationary error component and a Lagrange multiplier method is used to test the hypothesis that the random walk component has a zero variance. The KPSS test thus supplements the ADF test in that both test for a unit root and stationarity. The results of the KPSS test for all sub-samples of ALSI returns are displayed in Table 76 below. For all frequencies tested, the KPSS test shows that the null hypothesis of stationarity is not rejected.

Table 76 - KPSS results for each sub-sample

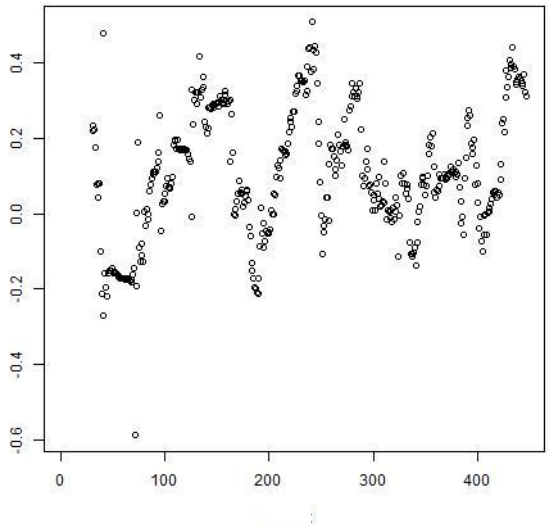
Sub-sample	Test statistic value	T.P	P-Value
#1	0.1387	4	0.1*
#2	0.0720	4	0.1*
#3	0.1128	4	0.1*
#4	0.2062	4	0.1*
#5	0.0441	4	0.1*
#6	0.0918	4	0.1*
#7	0.1462	4	0.1*
#8	0.0415	4	0.1*
#9	0.1249	4	0.1*
#10	0.0628	4	0.1*

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

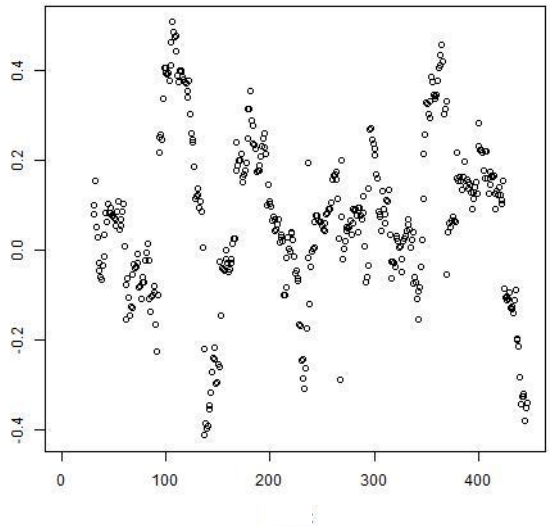
Therefore, according to the KPSS test, all of the sub-samples under consideration are stationary. With stationary variables, one can now proceed to conduct tests on the variance of the data over time; where changes in this variance would have implications for market efficiency.

4.4.2.3 Rolling autocorrelation test

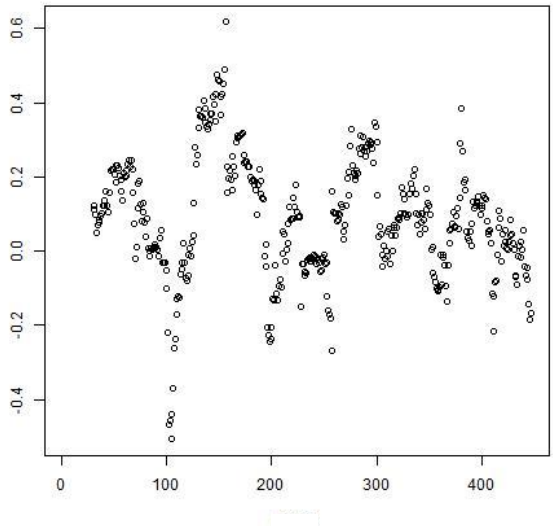
Using the sub-sample returns, Figure 16, Figure 17 and Figure 18 show that the rolling autocorrelation coefficient fluctuates around zero for each sub-sample period, implying that there are small periods of time where the ALSI experiences market inefficiency. This therefore implies that over the sample period examined, the ALSI experienced periods of weak form efficiency and periods of weak form inefficiency.



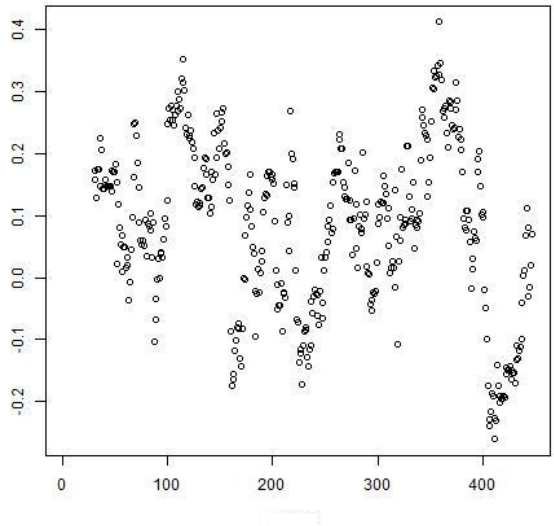
#1



#2

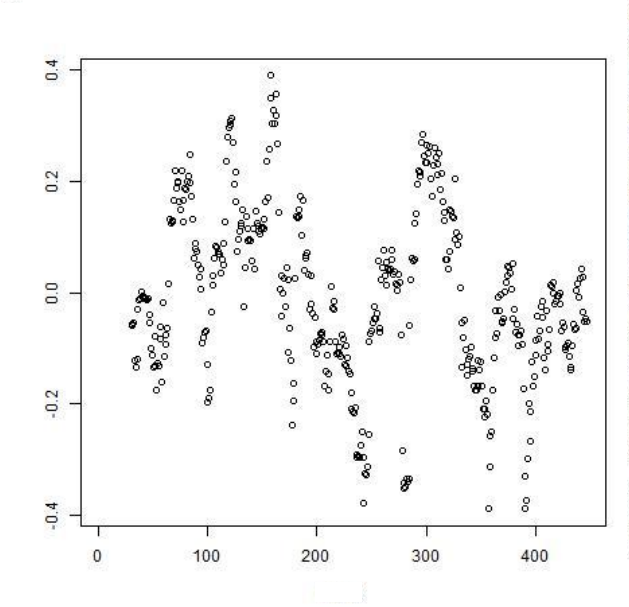


#3

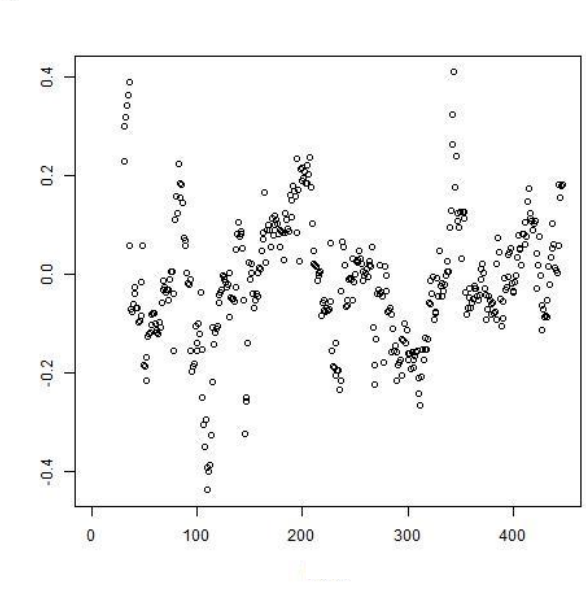


#4

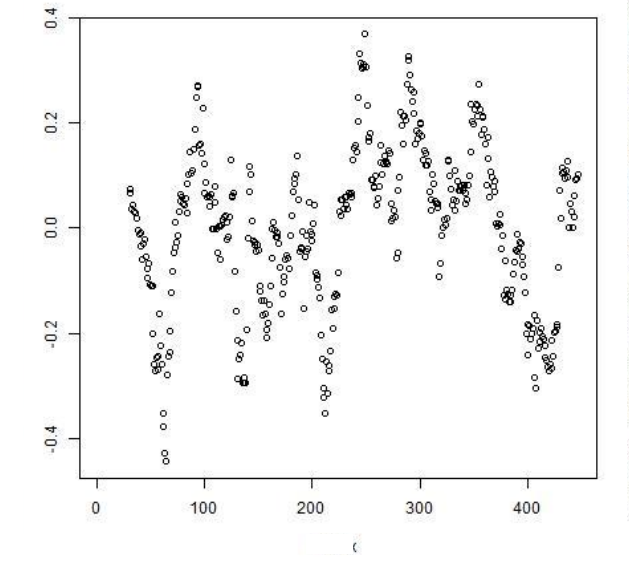
Figure 16 - Rolling autocorrelation for each sub-sample (1)



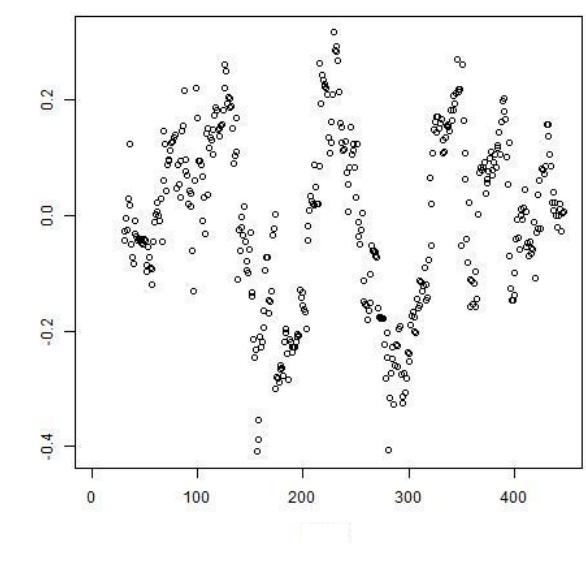
#5



#6

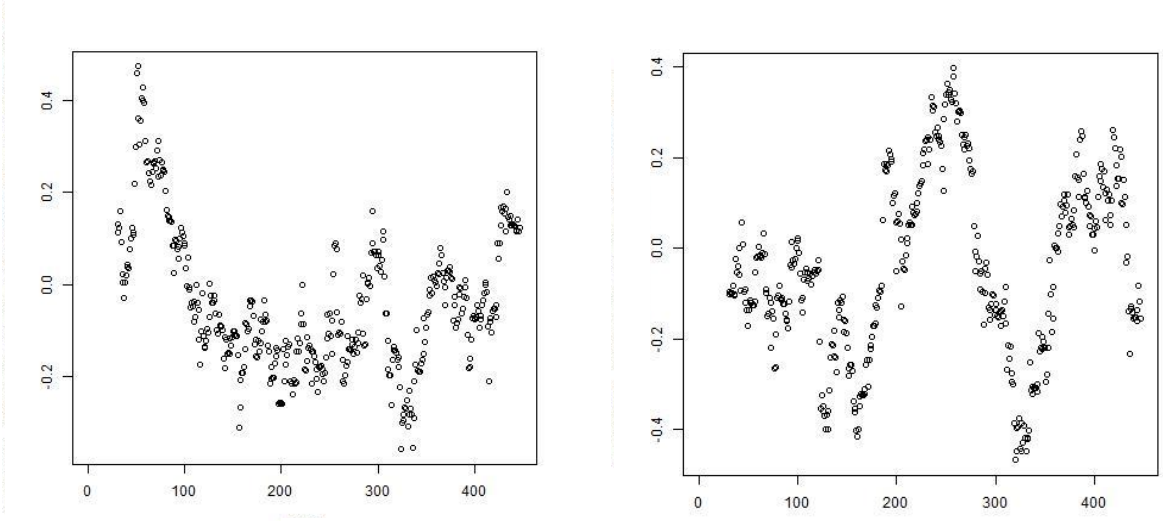


#7



#8

Figure 17 - Rolling autocorrelation for each sub-sample (2)



#9 #10

Figure 18 - Rolling autocorrelation for each sub-sample (3)

Therefore, according to a rolling autocorrelation test, the ALSI experienced periods of market efficiency and periods of market inefficiency, across all frequencies examined. While not a strict statistical procedure in and of itself, the rolling autocorrelation graphs provide indirect evidence of cyclical efficiency.

4.5 Testing the random walk hypothesis

To examine possible random walk behaviour of the data, the runs test, variance ratio tests (specifically the Chow Denning and Wright modifications) and the Hurst exponent are used for the different frequencies of equities data only. These tests allow for multiple variances in the data and are an improvement over the popular Lo-MacKinlay variance ratio tests (as a result, the Lo-MacKinlay tests are not conducted). Further, graphical methods, namely plots of variance decomposition and the Hurst exponent are employed to add robustness to the results. The former allows one to simply view the evolution of variance over time (similar to a Q-Q plot in its purpose) while the latter is a sophisticated method used to detect if the sample distribution is random or deterministic (both in the short term or long term).

4.5.1 Full sample results

4.5.1.1 Runs test

The Runs test is a simple, non-parametric means of assessing whether a series is randomly generated or not. It can be considered a precursor to the Hurst exponent, tested later in this study.

The results of the Runs test for the daily returns data are shown in Table 77 below. The null hypothesis of randomness in the data is rejected if the p-value is statistically significant. The results show that few of the daily equities are randomly generated. In particular, BIL, MTN, SAB (and FSR, at a 10% level of significance) are randomly generated. For the majority (43) of the 50 securities that are not randomly generated, all have runs that are less than expected by chance, implying less fluctuation in the return generating process. This however does not imply that the volatility is lower, but rather that there are fewer instances where the change from a positive run to a negative run, enabling a "smoother" process, with the same amount of deviation from the mean. This smoother process should imply a greater level of confidence in being able to fit a model – which is to be investigated later.

Table 77 - Runs test for daily returns

Share Code	Test Statistic	P-value
BIL	-0.269	0.79
MTN	-0.9862	0.32
SOL	-2.57	0.01**
FSR	-1.7034	0.09*
SAB	-0.6574	0.51
NPN	-3.855	0.00***
AGL	-3.2275	0.00***
J200	-2.3608	0.02**
ALSI	-2.7194	0.00***

Note: * denotes a 10% level of significance, ** denotes a 5% level of significance and *** denotes a 1% level of significance

Using weekly returns data, it is found in Table 78 that only one return series is not randomly generated. While one of these (SAB) has the same conclusion under daily data, it is interesting to note that the ALSI under weekly data is randomly generated (whereas under daily data it was found to not be randomly generated). This could imply any pricing

anomalies that manifest themselves in the short term are eliminated over a week. Further, it conceptually shows that a non-randomly generated series can be a subset of a randomly generated series. From the population of securities studied, nine are non-randomly generated, with their runs being greater than expected by chance. This is in contrast to the daily results, implying that there is more fluctuation in the returns process under lower frequency data.

Table 78 - Runs test for weekly returns

Share Code	Test Statistic	P-value
BIL	0.9359	0.35
MTN	0.9359	0.35
SOL	3.0084	0.00***
FSR	0.9359	0.35
SAB	0.7354	0.46
NPN	-0.0669	0.95
AGL	-0.5348	0.59
J200	-0.6685	0.50
ALSI	-0.4011	0.69

Note: * denotes a 10% level of significance, ** denotes a 5% level of significance and *** denotes a 1% level of significance

Lastly, by examining monthly returns data in Table 79, two of the series are not randomly generated, with SAB maintaining the same conclusion under all three frequencies. It is again found that the ALSI, under monthly data, is randomly generated using the Runs test. Here, only 8 securities are non-randomly generated, with the majority of these 8 having runs that are greater than chance. Only three shares, AFE, ASR and TFG have runs that are less than expected by chance. This is perhaps due to the lower liquidity.

Table 79 - Runs test for monthly returns

Share Code	Test Statistic	P-value
BIL	1.8159	0.07*
MTN	-1.3969	0.16
SOL	2.235	0.03**
FSR	-0.5588	0.58
SAB	0.6984	0.48
NPN	-1.1175	0.26
AGL	0.6984	0.48
J200	-0.1397	0.89
ALSI	-0.4191	0.68

Note: * denotes a 10% level of significance, ** denotes a 5% level of significance and *** denotes a 1% level of significance

Under quarterly data, 12 securities were non-random, with four having runs that were greater than chance. These shares were across three industries, so there is no discernible sector-specific pattern. Under semi-annual data, five securities were non-random, with two of them having runs thata were greater than chance. From all securities, there was no instance where it was found to be non-random under all frequencies. One share, SOL, was non-random under four of the frequencies (from daily to quarterly), and had the middle two frequencies of runs being greater than chance. Other exceptions were ASR and PIK, where ASR had daily, weekly and monthly runs being less than chance, and PIK being greater than chance under weekly, monthly and quarterly data.

4.5.1.2 Wright test

To test individual variance ratios, the non-parametric version of the variance ratio test by Wright (2000) is used. Results from the Wright test (Wright, 2000) for the daily return series in Table 80 provide mixed evidence of whether the series are identically and independently distributed. All series (indeed from the population of securities examined) have some lags that are statistically significant and others that are not. The ALSI appears to not follow a random walk for lower lags, yet becomes a randomly generated series at higher lags.

Table 80 – Wright test on daily returns

BIL	R 1	R 2	S 1
k=2	0.9703	1.1769	1.1653
k=5	-1.1089	-1.0571	0.8238
k=10	-2.5023**	-3.0621***	0.4885

NPN	R 1	R 2	S 1
k=2	4.4304***	5.0889***	5.4981***
k=5	2.1352**	2.4871**	5.2645***
k=10	0.9546	1.246	5.0798***

MTN	R 1	R 2	S 1
k=2	2.4868**	3.2687***	1.6136*
k=5	0.1109	0.6225	1.9149*
k=10	-1.3267	-1.171	2.347**

AGL	R 1	R 2	S 1
k=2	3.8024***	3.6054***	3.496***
k=5	2.2982**	2.088**	2.9296***
k=10	-0.0492	-0.3076	1.0195

SOL	R 1	R 2	S 1
k=2	4.9494***	4.9959***	3.8546***
k=5	0.8977	0.8344	1.9585**
k=10	-1.5297	-1.5512	0.8691

J200	R 1	R 2	S 1
k=2	3.7386***	3.7625***	4.1833***
k=5	1.697*	1.4468	3.4315***
k=10	0.1868	-0.3184	3.0461***

FSR	R 1	R 2	S 1
k=2	2.3786**	3.2079***	1.9721**
k=5	-0.4105	-0.1823	1.042
k=10	-2.0241**	-2.0241**	0.3735

WBO	R 1	R 2	S 1
k=2	4.1177***	4.3166***	11.8029***
k=5	4.6183***	4.3087***	18.1939***
k=10	4.0851***	3.4396***	23.6257***

SAB	R 1	R 2	S 1
k=2	0.5856	1.2056	0.747
k=5	-1.3569	-1.5002	0.1473
k=10	-1.7494*	-2.0011**	0.6142

ALSI	R 1	R 2	S 1
k=2	4.3878***	4.4929***	4.0638***
k=5	2.5**	2.3815**	3.5079***
k=10	1.0381	0.6781	3.4957***

Note: * denotes a 10% level of significance, ** denotes a 5% level of significance and *** denotes a 1% level of significance

The results from the weekly return data (Table 81) show a similar trend of mixed conclusions. One share, SOL, shows significant test statistics at all lag levels, whereas other shares show no significance at any lag level. Two shareas, AGL and INP, are randomly generated as they have no significant lags. In particular, the ALSI only shows significant S1 statistics, indicating that the distribution of positive and negative values deviates from a Gaussian distribution. There are 28 securities that have a random distribution according to the Wright test, and these seem to be concentrated in the financials, healthcare and industrials sectors. These align to the intuition behind the Wright test – that a longer sample interval should portray a greater variance, implying more chance of randomness.

Table 81 - Wright test on weekly returns

BIL	R 1	R 2	S 1
k=2	-2.0676*	-3.1129***	-0.735
k=5	-1.9639*	-2.5353**	-0.6465
k=10	-1.7012	-1.9373*	-0.661

NPN	R 1	R 2	S 1
k=2	-0.3292	-0.2493	0.8018
k=5	0.5336	0.7231	2.0128**
k=10	0.5826	0.6368	1.6029

MTN	R 1	R 2	S 1
k=2	-2.3533**	-2.5426**	-1.6704
k=5	-1.7911*	-1.663	-1.2077
k=10	-1.3556	-1.1296	-1.1201

AGL	R 1	R 2	S 1
k=2	-0.6004	-1.0219	0.6013
k=5	-1.1083	-1.3166	-0.1342
k=10	-0.724	-0.9005	0.186

SOL	R 1	R 2	S 1
k=2	-3.4887***	-3.3531***	-3.608***
k=5	-2.998***	-2.7942***	-2.6227***
k=10	-2.8997***	-2.7633***	-1.5277

J200	R 1	R 2	S 1
k=2	-0.088	-0.382	1.4031
k=5	-0.4902	-0.5097	1.72*
k=10	-0.3209	-0.2582	1.8879*

FSR	R 1	R 2	S 1
k=2	-2.1948**	-2.3394**	0.3341
k=5	-2.4159**	-2.3513**	-0.0854
k=10	-2.0651**	-2.173**	0.0198

INP	R 1	R 2	S 1
k=2	-0.015	-0.5246	1.0022
k=5	-1.08	-1.5261	0.1098
k=10	-0.6852	-1.0845	0.4354

SAB	R 1	R 2	S 1

ALSI	R 1	R 2	S 1

k=2	-1.5748	-1.7602*	0.1336
k=5	-2.2639**	-2.233**	-0.1342
k=10	-1.617	-1.6787	0.4551

k=2	0.615	0.3154	2.3385**
k=5	0.3397	0.2709	2.33**
k=10	0.4748	0.4359	3.1148***

Note: * denotes a 10% level of significance, ** denotes a 5% level of significance and *** denotes a 1% level of significance

Lastly, the monthly returns data (Table 82) show that few of the shares show any significant statistics in not being randomly generated series. In particular, the ALSI has one significant S1 statistic at a lag of 10. This by itself is not enough evidence to conclude that the ALSI is not randomly generated. From the population of 50 securities, most of the randomly generated series are from the financials and consumer goods sectors.

Table 82 - Wright test on monthly returns

BIL	R 1	R 2	S 1
k=2	-1.7834	-1.702	-0.2787
k=5	-1.0599	-0.8763	0.7887
k=10	-1.2232	-1.0787	0.6108

NPN	R 1	R 2	S 1
k=2	1.0758	1.3344	0.9754
k=5	-0.6353	-0.4634	0.4834
k=10	-0.6539	-0.4876	1.1226

MTN	R 1	R 2	S 1
k=2	1.517*	1.5697*	0.418
k=5	1.6108*	1.589*	1.4501
k=10	0.8962	1.0552	1.2959

AGL	R 1	R 2	S 1
k=2	-0.8	-0.3054	-0.8361
k=5	-0.6667	-0.0552	-0.1272
k=10	-0.8175	-0.4713	0.0743

SOL	R 1	R 2	S 1
k=2	-1.9726*	-1.8627*	-2.0902**
k=5	-0.6955	-0.4218	-0.229
k=10	-0.3767	-0.1709	-0.0083

J200	R 1	R 2	S 1
k=2	-1.0258	-0.888	3.7624***
k=5	-1.0513	-0.7697	6.5892***
k=10	-1.0637	-0.9702	8.6091***

FSR	R 1	R 2	S 1
k=2	-1.2322	-1.4995	0.6967
k=5	-0.9655	-1.3903	0.9413
k=10	-0.7091	-1.2945	1.3702

OCE	R 1	R 2	S 1
k=2	0.1627	0.4787	0.1393
k=5	-1.0146	-0.963	0.4325
k=10	-0.7551	-0.6361	0.9822

SAB	R 1	R 2	S 1
k=2	-1.0325	-1.0739	1.1148
k=5	-1.4111	-1.5371	2.5696**
k=10	-0.7535	-0.9524	4.2096***

ALSI	R 1	R 2	S 1
k=2	-0.5511	-0.4369	0.6967
k=5	-0.6392	-0.478	0.7378
k=10	-0.6058	-0.6118	1.7499*

Note: * denotes a 10% level of significance, ** denotes a 5% level of significance and *** denotes a 1% level of significance

Therefore, according to the Wright test, there is no discernible conclusion that can be reached about the shares in general, which in the strictest sense leads one to fail to reject the null hypothesis of a randomly generated series. Under quarterly data, only two securities (IMP and GND) were non-random. These two were also non-random for the higher frequencies (daily, weekly and monthly). However, under semi-annual data, no security was non-randomly generated. In particular, the daily and weekly ALSI show weak evidence of departures from the independence and identically distributed assumption (with the daily results being a somewhat stronger rejection of random walk behaviour than the weekly results), while the monthly ALSI has no evidence to show it deviates from a randomly generated series. As such, a test for multiple variances is now employed with the aim of providing clearer results.

4.5.1.3 Chow Denning test

The Chow Denning test (Chow and Denning, 1993) for multiple variances is employed, in which the null hypothesis is that the series follows a random walk. The CD2 test statistic is also provided as a heteroscedasticity-robust version of the CD1 statistic. Therefore, one can reject the null hypothesis of a random walk if both versions of the test statistic are significant.

Using daily returns (Table 83), one of the returns series, SHF has insignificant test statistics, implying that the series follow a random walk. The remaining return series, have significant values of either the CD1 or CD2 statistic, indicative non-random behaviour. In the example of MTN and SAB, the CD2 statistic is not significant at the 5% level, implying that while non-random behaviour might be present, it is quite likely "masked" by multiple variances as the CD1 statistic implied non-random behaviour, but when controlling for multiple variances, no significant evidence of non-random behaviour was found. A contrasting result to the Runs test is found with HYP - it was found to not be randomly generated under the Runs test yet was found to be randomly generated under the Chow Denning test. Applying the same logic used in the case of MTN, the conflicting results are most likely due to the presence of multiple variances. Examining all 50 securities, only six are randomly generated with no particular sector pattern emerging.

Table 83 – Chow Denning test on daily returns

Share Code	CD1	CD2
BIL	3.799***	2.5526**
MTN	3.083***	1.7195
SOL	4.1877***	2.535**
FSR	4.0026***	2.9092**
SAB	2.5954**	1.8282
NPN	5.3459***	3.3411***
AGL	2.8803**	1.8526
J200	3.1154***	1.7223
SHF	0.6658	0.4059
ALSI	4.0791***	2.2615*

Note: * denotes a 10% level of significance, ** denotes a 5% level of significance and *** denotes a 1% level of significance

The Chow Denning test is now run on weekly returns and is shown in Table 84 below. Only two equity series, BIL and SOL, have significant test statistics at the 5% level, implying a rejection of random walk behaviour. The remaining series, including the ALSI, appear to follow a random walk. Similar to the case when daily returns were used, MTN and in this example, FSR, appear to have non-random behaviour but this is masked by multiple variances. Of the population of securities, 39 are randomly generated, mostly from the financials, healthcare and industrials sectors. The results of the ALSI thus far correlate to those of the Runs test in that the daily series was found to not be randomly generated in contrast to the weekly series.

Table 84 - Chow Denning test on weekly returns

Share Code	CD1	CD2
BIL	4.482***	2.676**
MTN	2.8216**	1.9156
SOL	3.5039***	2.465**
FSR	2.8641**	1.5216
SAB	2.359*	1.7845
NPN	1.0138	0.7086
AGL	1.8433	1.3347
J200	1.393	0.8941
SHF	3.1031***	1.7764
ALSI	0.5369	0.3501

Note: * denotes a 10% level of significance, ** denotes a 5% level of significance and *** denotes a 1% level of significance

Using monthly returns (Table 85), all of the series do exhibit random walk behaviour under the CD2 statistic. Again, in the example of the ALSI, this is in line with the findings of the Runs test.This pattern is also found in the population of 50 securites – only three are non-random (PIK and TFG from consumer services and GND from industrials). Their non-random distributions could be due to low liquidity of the shares.

Table 85 - Chow Denning test on monthly returns

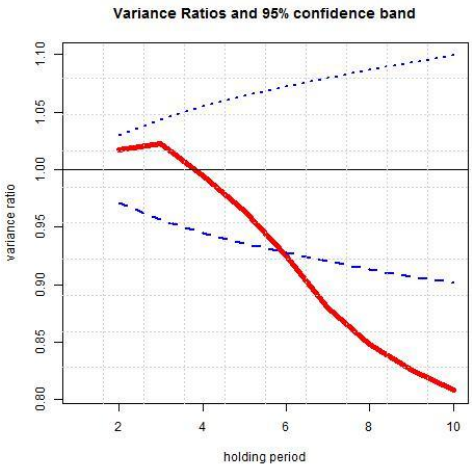
Share Code	CD1	CD2
BIL	1.6928	1.4282
MTN	1.8515	1.1523
SOL	1.5421	1.1219
FSR	2.1371*	1.6069
SAB	1.3548	1.132
NPN	1.5917	1.2155
AGL	0.273	0.253
J200	0.9148	0.8647
SHF	1.0325	0.79
ALSI	0.6482	0.5468

Note: * denotes a 10% level of significance, ** denotes a 5% level of significance and *** denotes a 1% level of significance

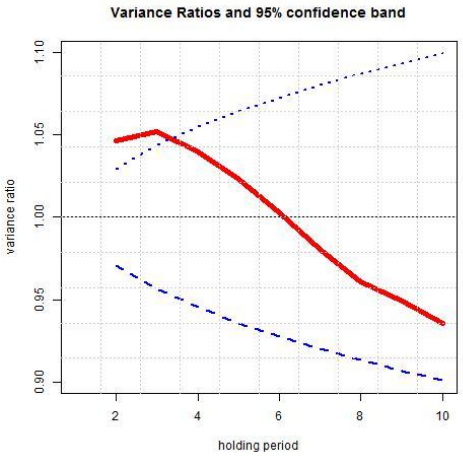
Under quarterly data, eight securities are non-randomly generated, with three under semi-annual data. No particular share is non-randomly generated under all frequencies, along with no discernible industry that stands out. In summary, according to the Chow Denning test for multiple variances, the results did correlate to the Runs test in most instances. In particular, the ALSI did not follow a random walk under daily data, but did follow a random walk under weekly and monthly data.

4.5.1.4 Variance decomposition plots

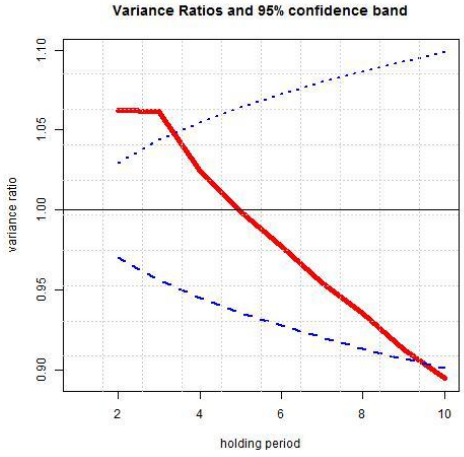
Figure 19 and Figure 20 below shows the variance ratios over time of the daily returns series. The variance ratio for BIL becomes significant at higher lags, indicating possible long term memory. In contrast, the variance ratio for MTN, SOL, FSR, NPN, AGL, the J200 and the ALSI have significant ratios at lower lags, indicating possible short term memory. Extending the analysis to the 50 securities, the majority (41) are non-random. The evidence therefore implies that the ALSI is not weak form efficient when considering daily returns data in the short term. This is in line with the results from the Runs test and variance ratio test.



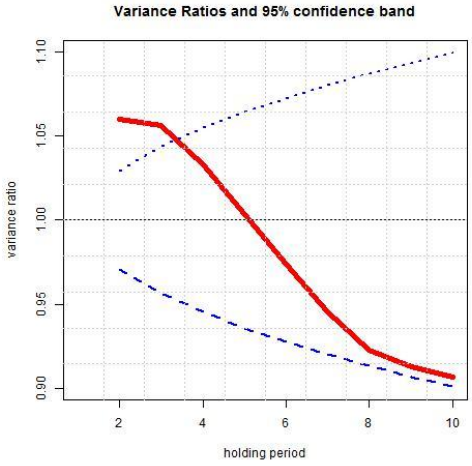
BIL



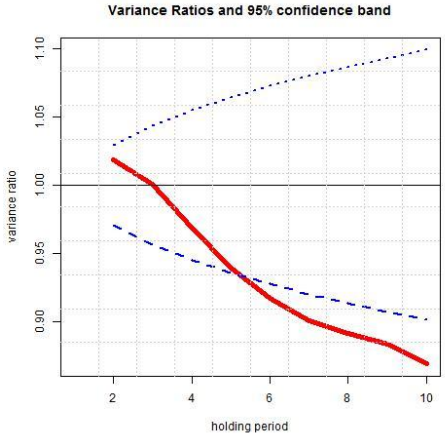
MTN



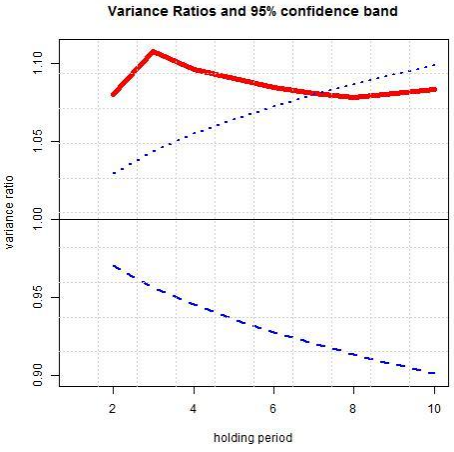
SOL



FSR

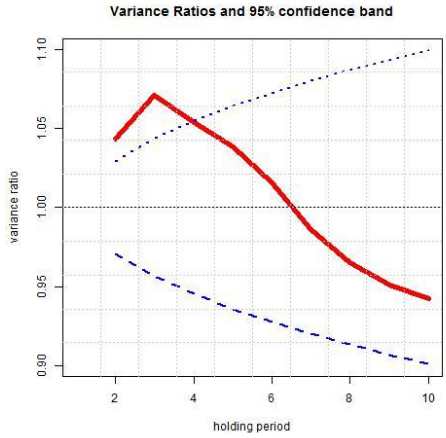


SAB

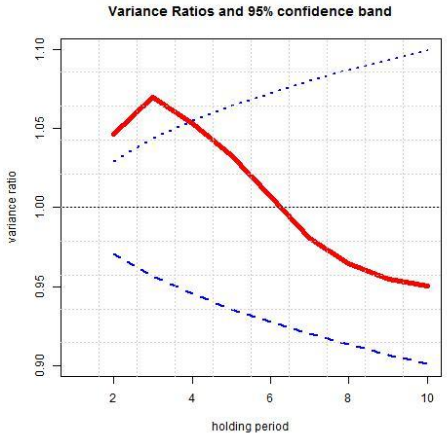


NPN

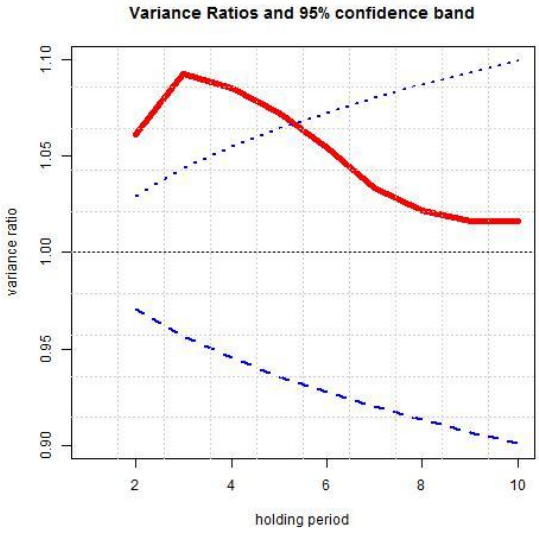
Figure 19 - Graphical representation of variance ratio for daily returns (1)



AGL



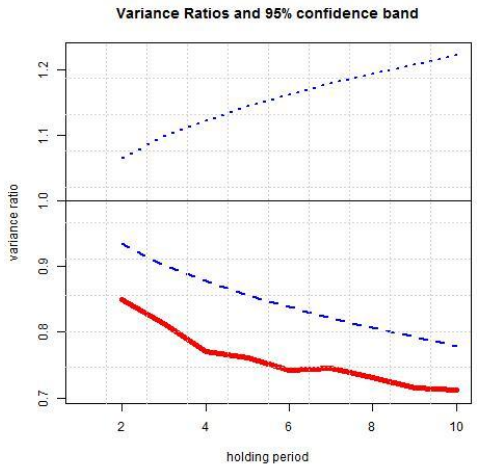
J200



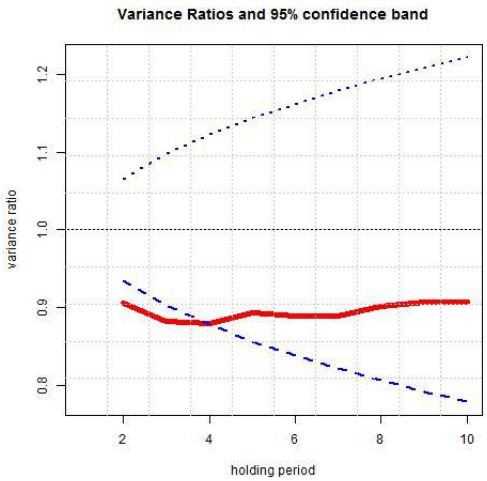
ALSI

Figure 20 - Graphical representation of variance ratio for daily returns (2)

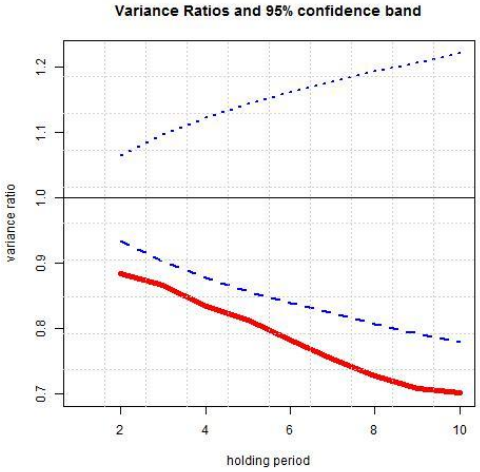
Figure 21 and Figure 22 below shows the variance decomposition of the weekly returns over time. Four return series, BIL, SOL, FSR and SAB, exhibit significant variance ratios, implying some form of memory exists in these weekly series. The remaining equities however, including the ALSI, do not display any significant memory characteristics. Indeed, only 23 securities display evidence of non-random behaviour across all industries. The evidence therefore implies that the ALSI is weak form efficient when considering weekly returns data, in line with results of previous tests. A possible reason for the ALSI being weak form inefficient under daily data and efficient under weekly data could lie in the speed of adjustment of stock prices to new information. On a daily basis, perhaps investors do not incorporate all information into stock prices, yet by the end of a week that information is priced into the stock.



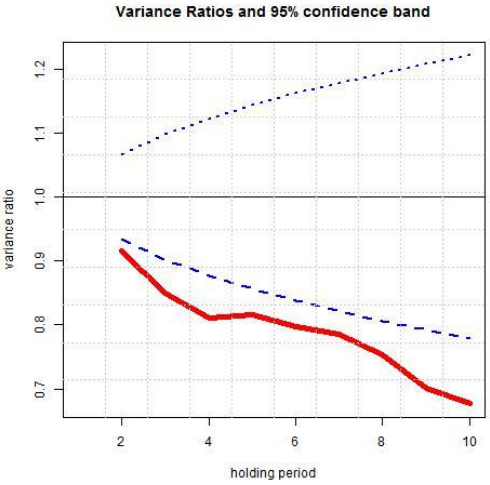
BIL



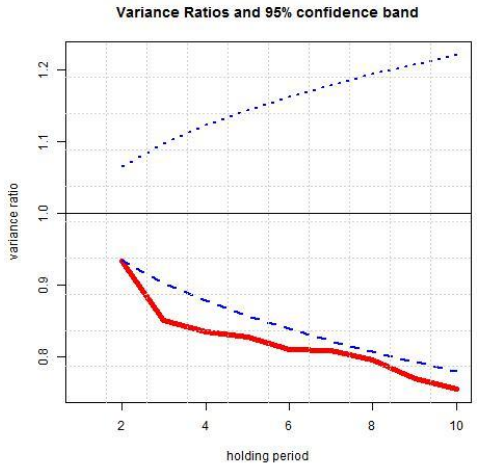
MTN



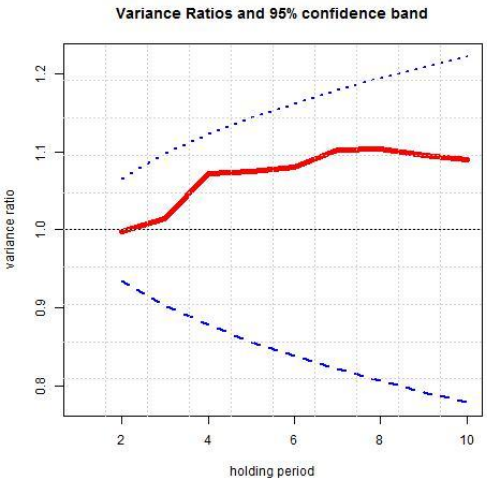
SOL



FSR

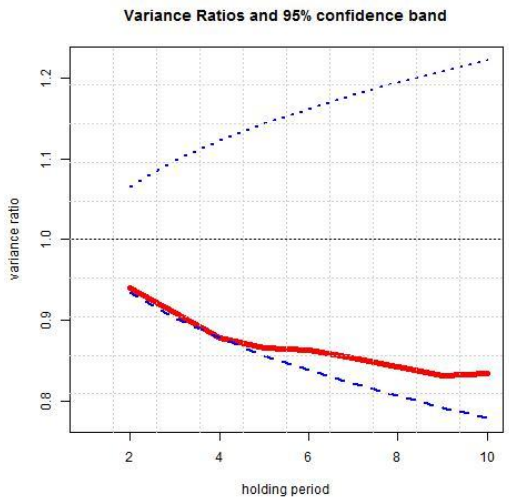


SAB

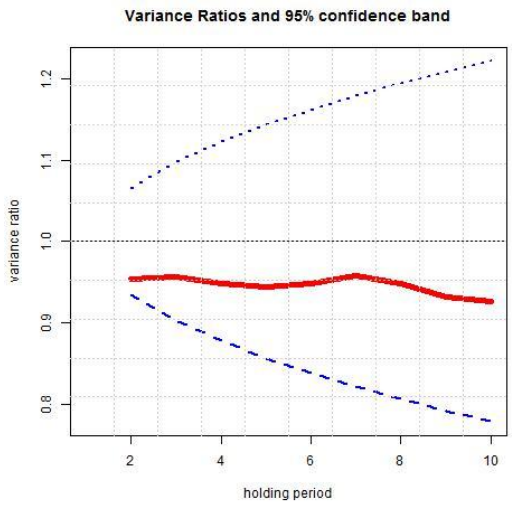


NPN

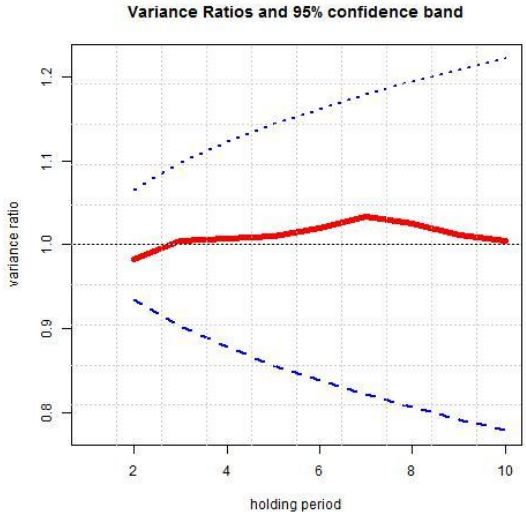
Figure 21 - Graphical representation of variance ratios for weekly returns (1)



AGL



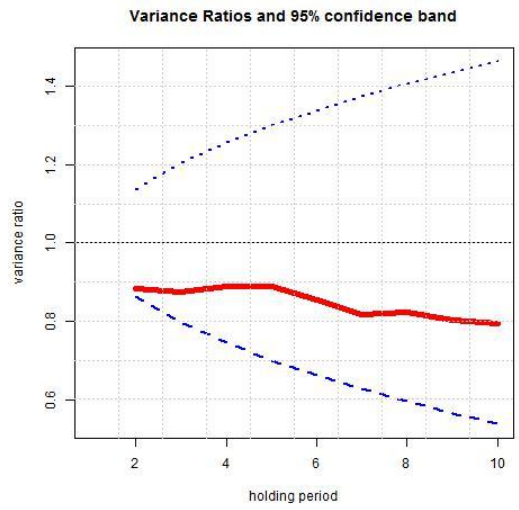
J200



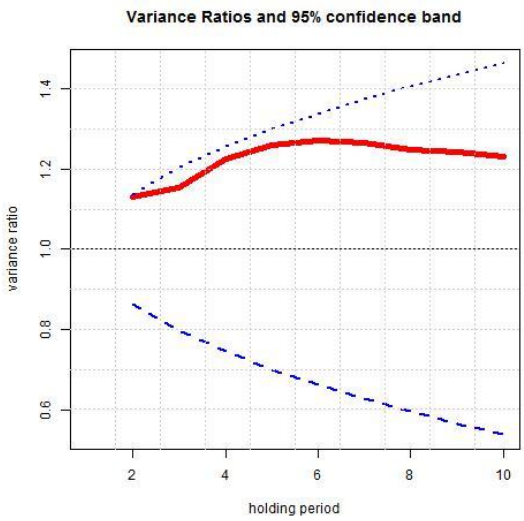
ALSI

Figure 22 - Graphical representation of variance ratios for weekly returns (2)

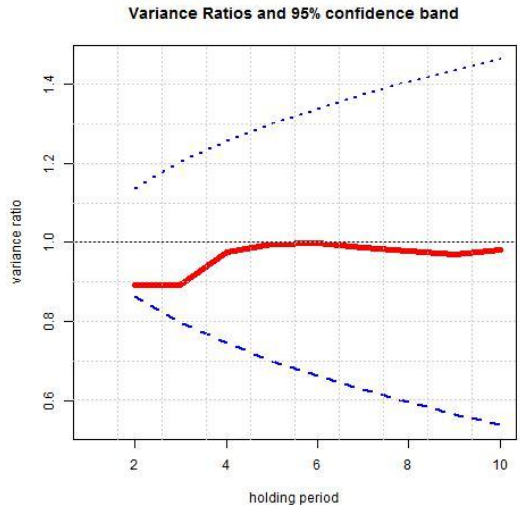
Figure 23 and Figure 24 below shows the variance decomposition of the monthly returns over time. Apart from one share (FSR), none of the other 49 securities display evidence of memory; the results are not statistically significantly different from random walk behaviour. The evidence therefore implies that the ALSI is weak form efficient when considering monthly returns data. This corroborates the weekly data test results and is in contrast to the daily data test results. The same line of reasoning in explaining why the weekly ALSI returns follow a random walk can be applied here.



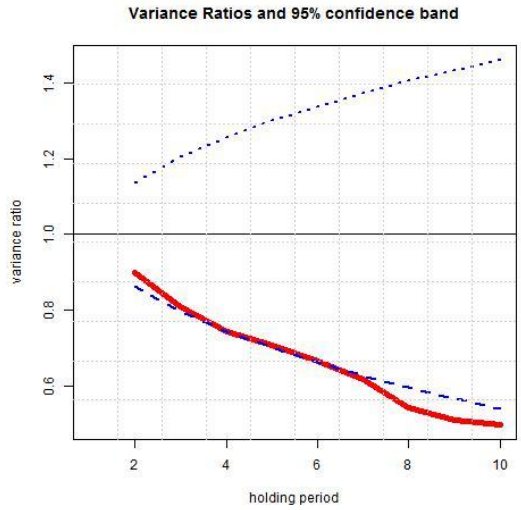
BIL



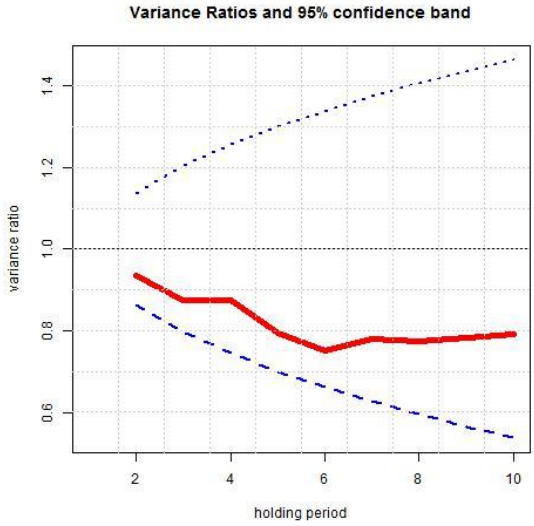
MTN



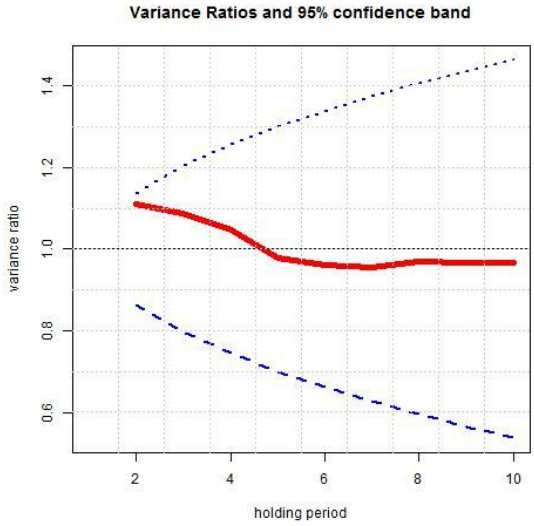
SOL



FSR

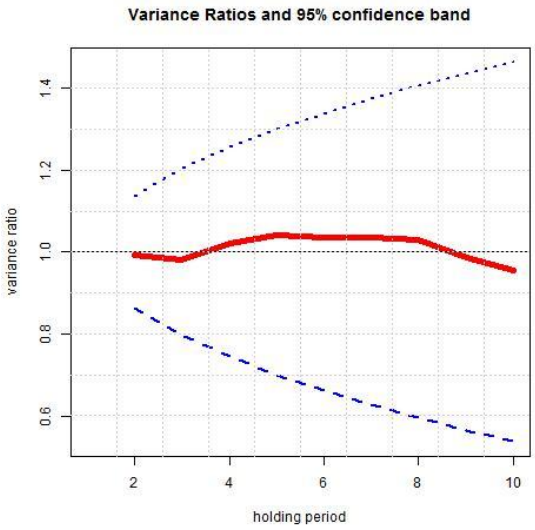


SAB

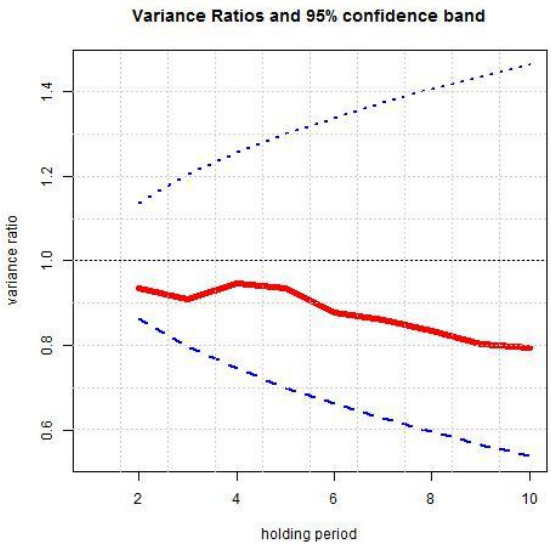


NPN

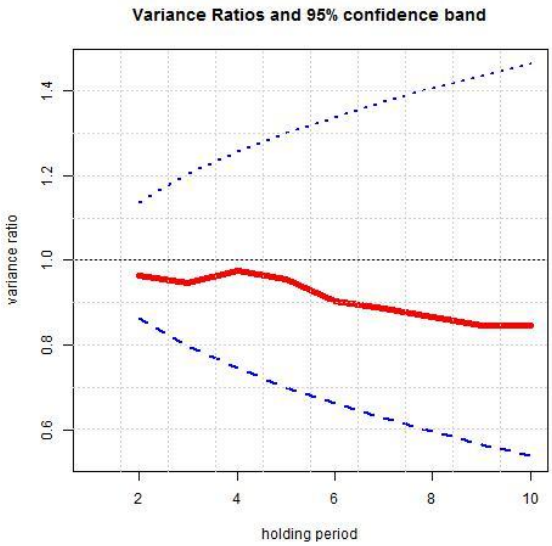
Figure 23 - Graphical representation of variance ratios for monthly returns (1)



AGL



J200



ALSI

Figure 24 - Graphical representation of variance ratios for monthly returns (2)

There are two securities (IMP and GND) that are non-random under quarterly data and none that are non-random under semi-annual data. Therefore, according to the visual evidence of the variance decomposition over time, the daily ALSI return series shows some (weak) element of predictability, but the weekly and monthly return series appears to behave like a random walk. This implies that investors do not fully incorporate all information into daily prices, but only in weekly and monthly prices. Further, it is quite possible (yet remains to be tested), whether an investor can earn abnormal profits on a daily portfolio strategy.

As a preliminary and basic test of market efficiency, a simple technical analysis rule (that of moving averages) was applied to the daily ALSI return series. Briefly, if an *x*-day moving average rule were to earn an investor profits after costs have been taken into account, then it points towards market inefficiency, particularly against the notion of share price behaviour following a random walk. While the results are not presented in detail here, the results showed that an investor would earn 4.2306% before costs on a 50 day moving average rule, 5.9636% before costs on a 150 day moving average rule and 7.5150% before costs on a 200 day moving average rule. Once costs have been taken into account, these returns decrease significantly, to the point of being negative. Thus, in the simplest scenario of practically testing whether markets are efficient according to the EMH, one finds in favour of the weak form of market efficiency. However, as argued in Chapter 2, given the multitude of trading rules available, one cannot reject or fail to reject the EMH unless one comprehensively tests all possible trading rules in existence. By conducting a simple moving average test using popular trading windows, this thesis found that the ALSI is weak form efficient. However, this result could be by pure chance and thus more robust statistical methods are required.

4.5.1.5 Hurst exponent test

The Hurst Exponent is further used to test for randomness in the data series. The results of the Peng method[28] in calculating the Hurst exponent are shown in Figure 25 and Figure 26 below. The corresponding 95% confidence interval for the Hurst exponent is in the range (0.4429, 0.5515). In other words, for a series to be classified as a random walk, the Hurst exponent should be in the range specified. The results from Figure 36 below show that BIL and SOL exhibit mean reversion, whereas the remaining series, including the daily ALSI return series, follows a random walk. Indeed, the results of the ALSI Hurst exponent are robust across all confidence intervals. This result is contradictory to the Runs test, yet can be explained by the method of sampling used in each - the Hurst exponent uses a rolling window approach to calculate the test statistic whereas the Runs test does not. Hence, the accuracy of the Hurst exponent can be considered superior to that of prior tests. In total, there are 8 securities that are non-random, specifically they are mean-reverting according to the Hurst exponent and span across sectors. The low number of non-randomly generated securities

[28] Recall that according to Taqqu *et al.* (1995), the Peng estimator should be used for series with 4000 to 7000 data points.

point is intuitive, as one would not expect patterns to emerge from high frequency data that can be realistically capitalised upon by an investor.

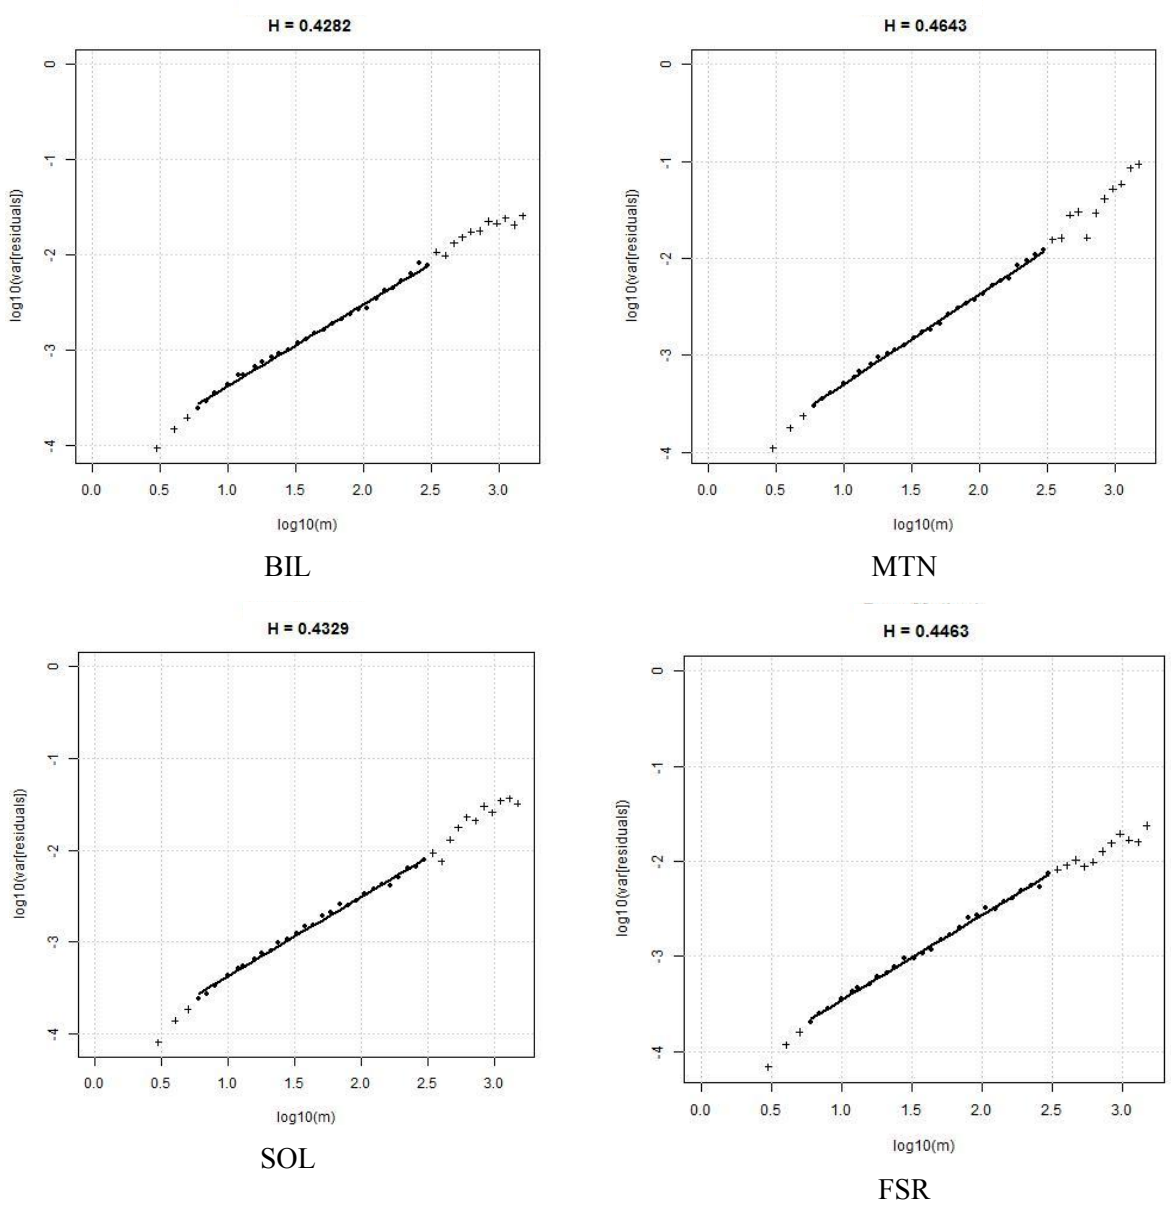
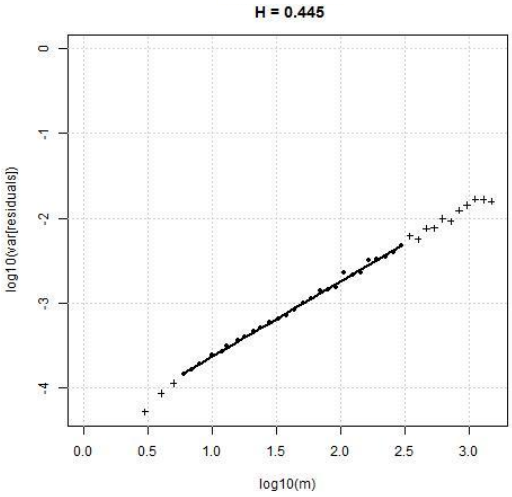
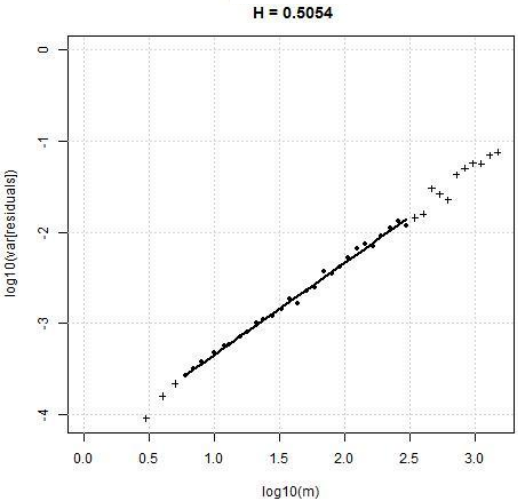


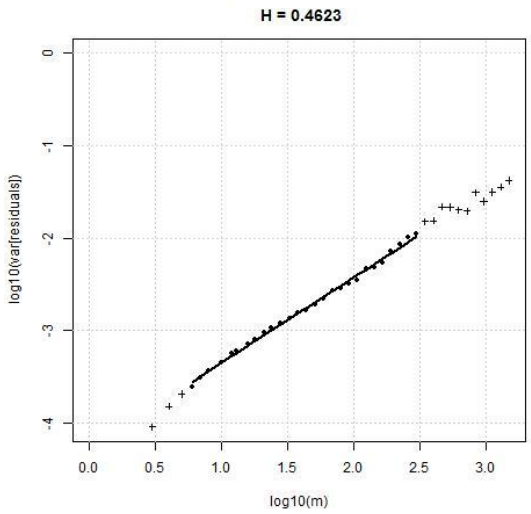
Figure 25 - Hurst Exponent (Peng method) for daily returns (1)



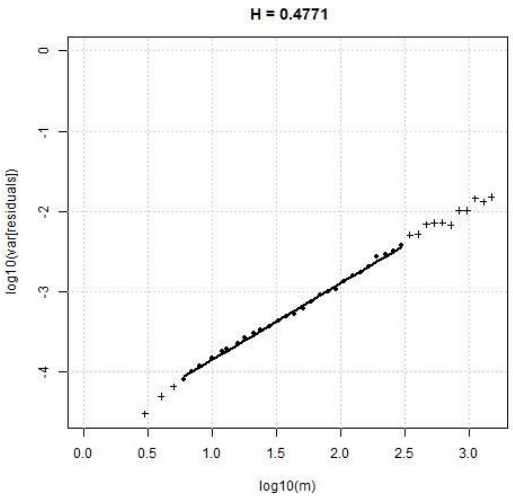
SAB



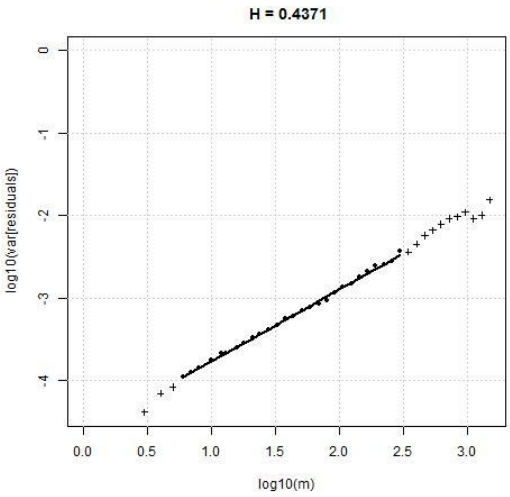
NPN



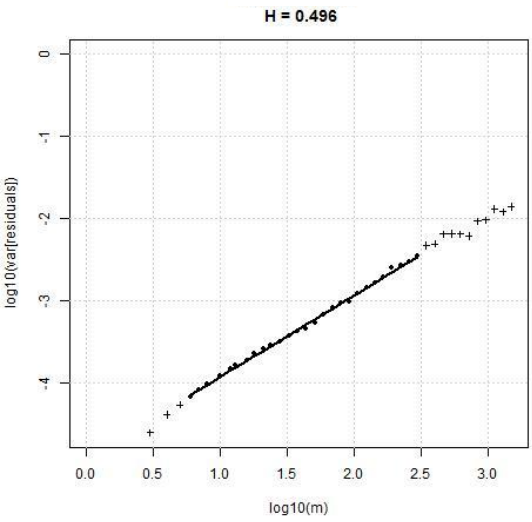
AGL



J200



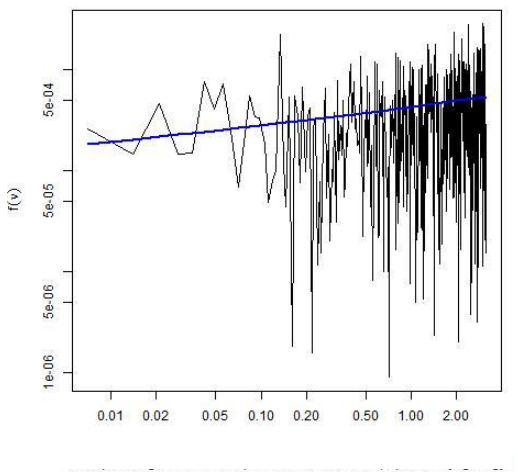
FPT



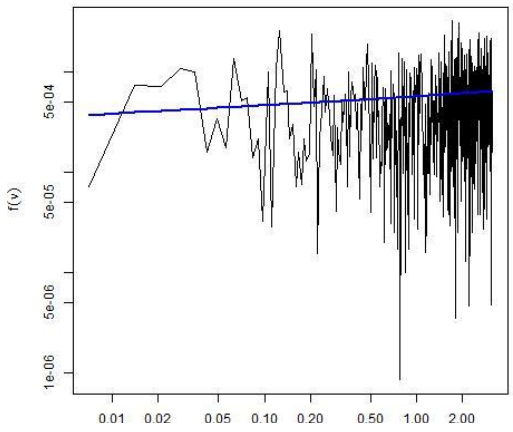
ALSI

Figure 26 - Hurst Exponent (Peng method) for daily returns (2)

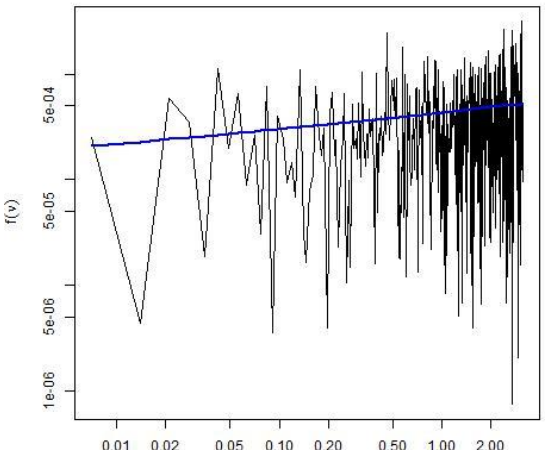
When examining the results of weekly returns via the Hurst exponent[29], shown in Figure 27 and Figure 28 below, the corresponding 95% confidence interval is (0.2027, 0.7630). It is found that all of the series below follow a random walk. Some of these results are in line with prior tests (including the R/S method), however, others are contradictory. Given the wide confidence intervals, none of the 50 securities are non-random. However, this is possibly due to the confidence intervals of the Whittle method. In particular, the ALSI has the same conclusion of following a random walk under both daily and weekly data.



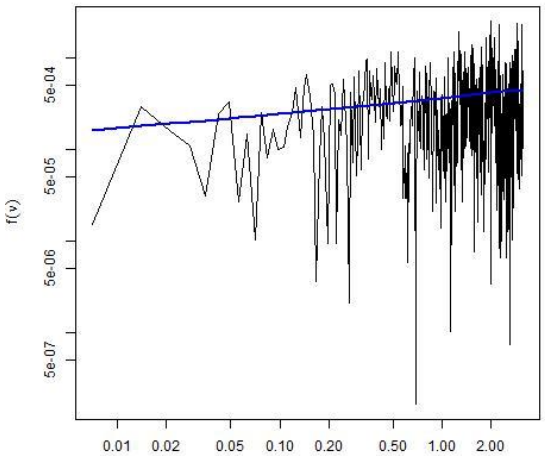
BIL (H = 0.4190)



MTN (H = 0.4594)



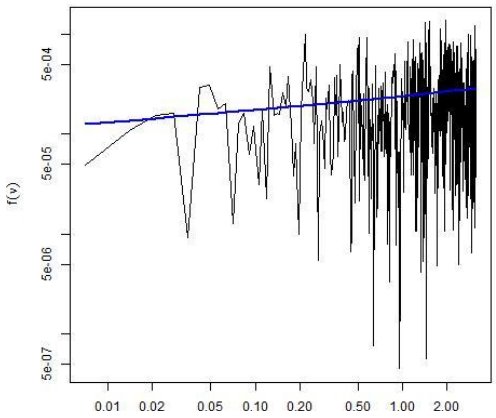
SOL (H = 0.4301)



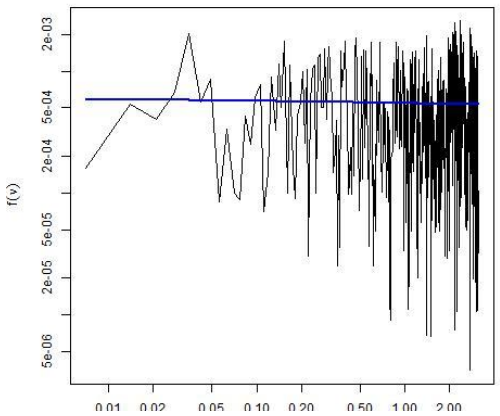
FSR (H = 0.4234)

Figure 27 - Hurst Exponent (Whittle Estimator) for weekly returns (1)

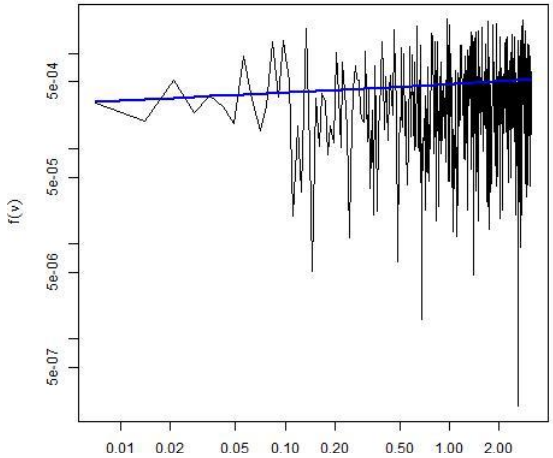
[29] Recall that according to Taqqu *et al.* (1995), the Whittle Estimate should be used for series with 700 to 1000 data points.



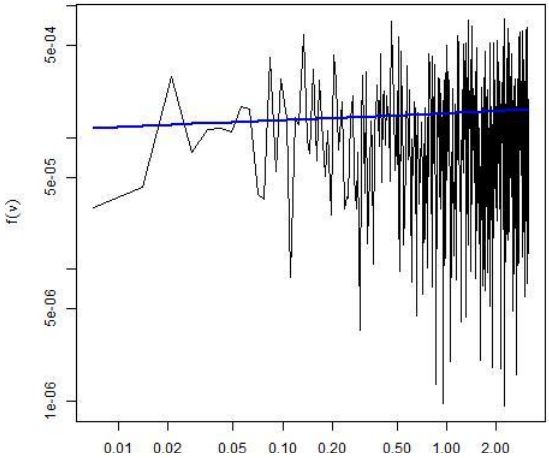
SAB (H = 0.4383)



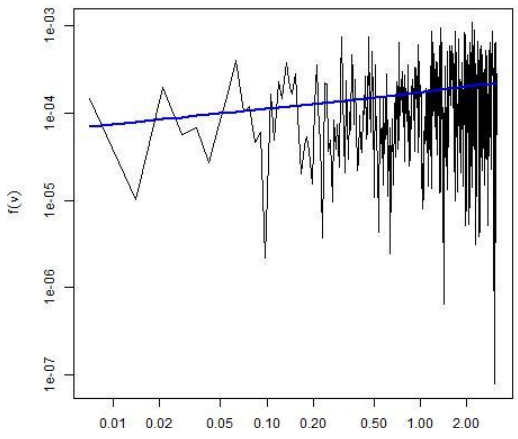
NPN (H = 0.5084)



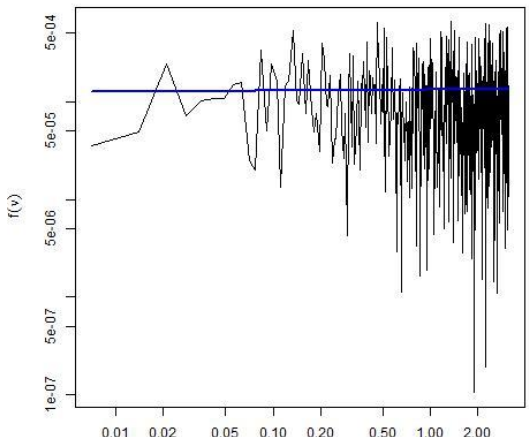
AGL (H = 0.4600)



J200 (H = 0.4763)



FPT (H=0.4154)

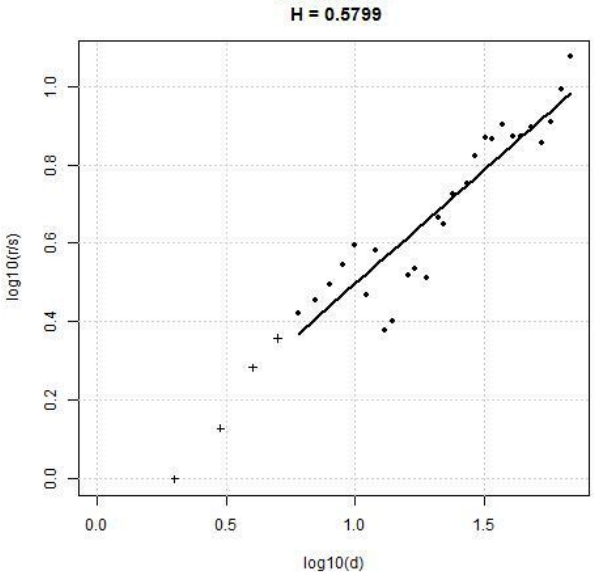


ALSI (H = 0.4950)

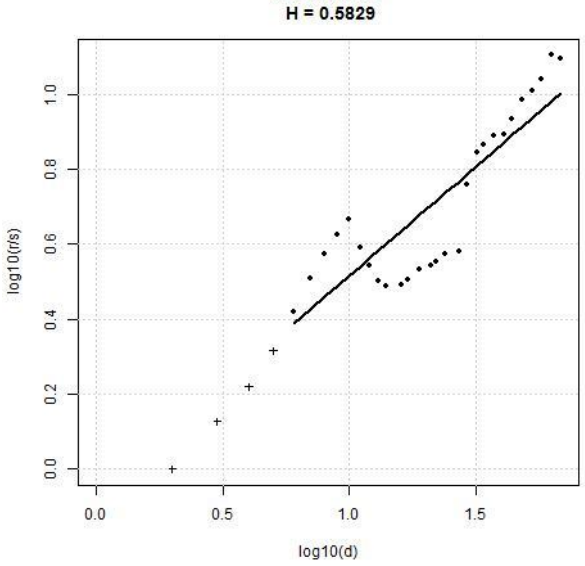
Figure 28 - Hurst Exponent (Whittle Estimator) for weekly returns (2)

When examining the Hurst exponent according to monthly data, the corresponding 95% confidence interval is (0.4503, 0.6405). According to Figure 29 and Figure 30 below[30], only four series, MTN, SOL, FSR, the J200 and the ALSI show signs of memory. Indeed, 22 securities of the 50 studied are non-random, with the indices and financial shares being mean-reverting and the remaining not. The mean-reversion property of the indices and shares implies a level of sophistication in the traders of those securities. By perhaps following technical analysis, there are cyclical profits to be obtained, as there is a tendency of the share to oscillate between the minimum and maximum resistance levels. These financial shares also had non-linear distributions under the BDS test, which is supportive evidence of the results from the Hurst exponent. The results of the ALSI are in contrast to those found in previous tests but are robust across all confidence intervals. In other words, the monthly returns of the ALSI do not follow a random walk, but rather show signs of short term memory. Again, this is contradictory to some of the prior test results and is reconciled in the summary section of this chapter.

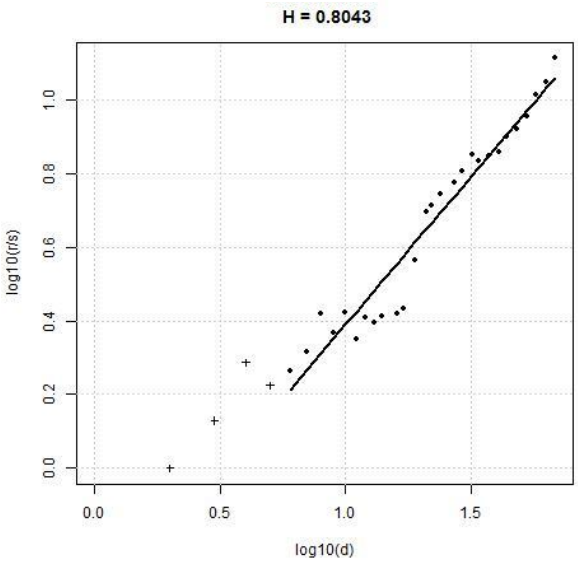
[30] Recall that according to Rea *et al.* (1995), the R/S method should be used for series with less than 700 data points.



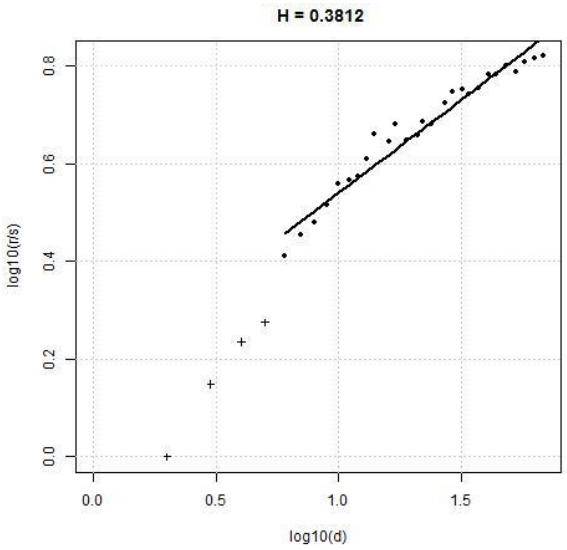
BIL



MTN

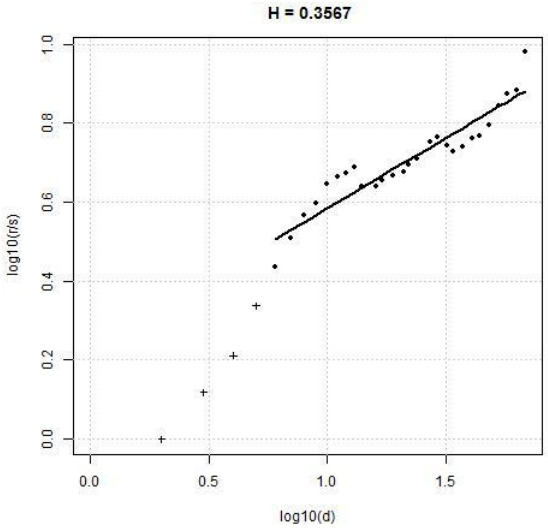


SOL

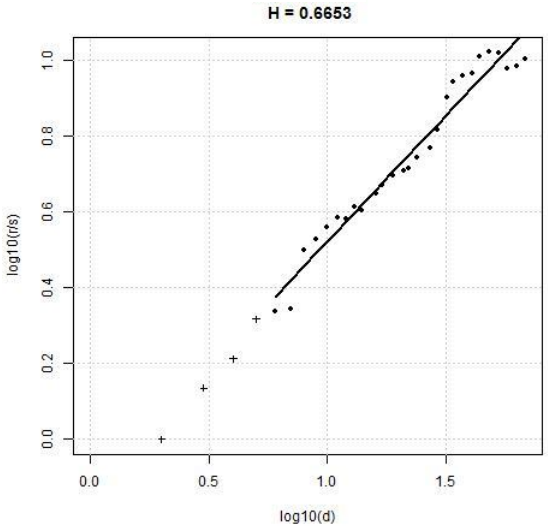


FSR

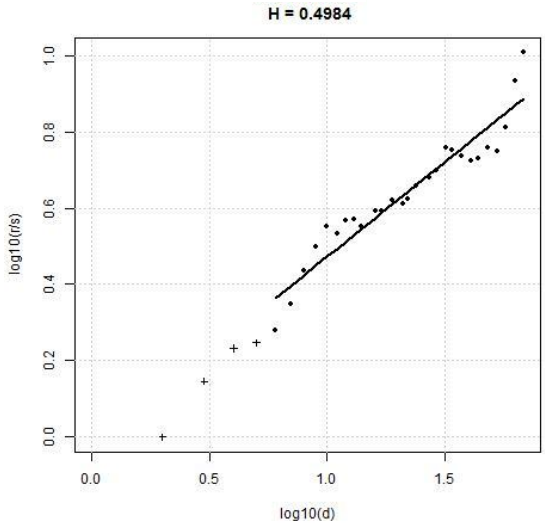
Figure 29 - Hurst Exponent (R/S method) for monthly returns (1)



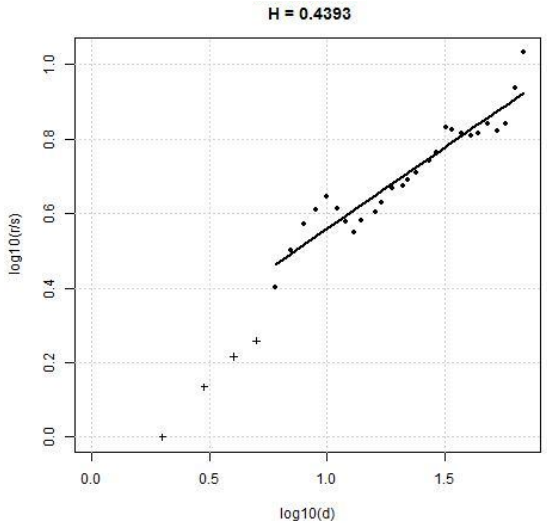
SAB



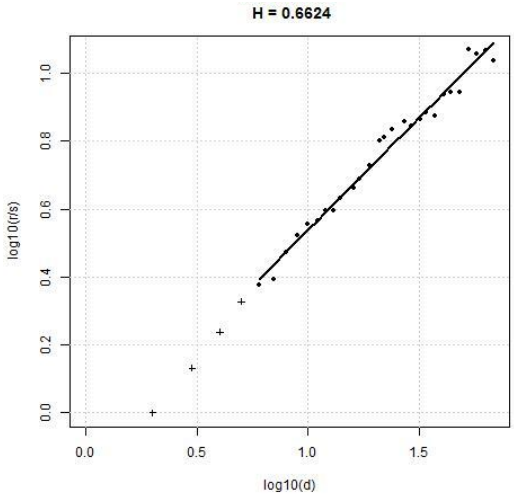
NPN



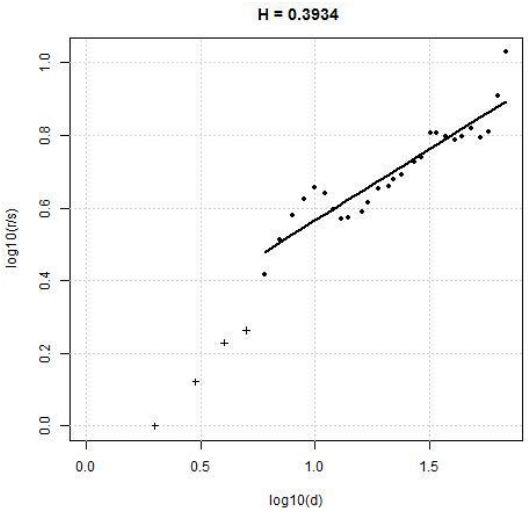
AGL



J200



FPT



ALSI

Figure 30 - Hurst Exponent (R/S method) for monthly returns (2)

Under quarterly and semi-annual data, there are 39 and 38 securities respectively that are non-random. Shares that are part of the healthcare, mining along with the indices tend to be persistent (not mean-reverting). There was only one share, MMI, that was mean reverting under both quarterly and semi-annual frequencies. The remaining shares that were non-random under one frequency did not follow the same conclusion under the lower frequency. In other words, the financial shares that were mean reverting under monthly data, were also mean reverting under quarterly data, but then found to be randomly generated under semi-annual data. This implies that the degree of mean-reversion decreased as the frequency lowered. This is intuitive as these shares have all seen significant increases in price over time, implying an increasing mean price as oppose to a constant mean price. Therefore, according to the Hurst exponent, the monthly return series exhibits mean reversion like behaviour, but the daily and weekly return series do not. This implies that the ALSI is weak form efficient if one examines daily or weekly data, but weak form inefficient if one examines monthly data.

4.5.2 ALSI Sub-sample results

4.5.2.1 Runs test

The Runs test is a simple, non-parametric means of assessing whether a series is randomly generated or not. It can be considered a precursor to the Hurst exponent, tested later in this sub-section. The results of the Runs test for the sub-sample returns data is shown in Table 86 below. The null hypothesis of randomness in the data is rejected if the p-value is statistically significant. The results are quite interesting as they show that all but three of the sub-samples (sub-samples 1, 2 and 4) exhibit randomness as all p-values are above the 5% level of statistical significance. As the daily returns were found to not be randomly generated according to the Runs test, the sub-sample results point towards the daily series being comprised of both random and non-random components.

Table 86 - Runs test result for each sub-sample

Sub-sample	Test Statistic	P-value
#1	-3.0271	0.00***
#2	-2.1757	0.03**
#3	-0.8533	0.39
#4	-1.9865	0.05**
#5	0.7568	0.45
#6	-0.2838	0.78
#7	-0.7568	0.45
#8	-0.3784	0.71
#9	1.0406	0.30
#10	0.0000	1.00

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

4.5.2.2 Wright test

Results from the Wright test (Wright, 2000) for each sub-sample are displayed in Table 87. Most sub-samples have insignificant test statistics, implying that the particular sub-sample does follow a random walk. The exceptions are the first four sub-samples, as one of the R1, R2 or S1 statistic is significant at a lower lag but becomes insignificant at higher lags. This implication reveals that all of the sub-samples can be concluded to follow a random walk, in contrast with the daily ALSI return series result.

Table 87 - Wright results for each sub-sample

#1	R 1	R 2	S 1
k=2	4.0674***	3.6529***	3.3072***
k=5	3.2903***	3.1032***	2.2945**
k=10	2.4745**	2.3165**	1.6008

#2	R 1	R 2	S 1
k=2	3.3194***	3.0144***	3.4017***
k=5	3.4264***	3.0489***	3.2261***
k=10	2.6882***	2.2941**	3.2799***

#3	R 1	R 2	S 1
k=2	2.723***	2.7636***	0.8504
k=5	2.016**	2.3592**	0.3968
k=10	1.0638	1.6888*	-0.2239

#6	R 1	R 2	S 1
k=2	0.2495	0.146	1.3229
k=5	-0.8264	-0.8667	1.6734*
k=10	-1.6553	-1.6943	1.8751*

#7	R 1	R 2	S 1
k=2	0.7302	0.9501	1.0394
k=5	-0.4752	-0.4974	0.6728
k=10	-0.7021	-0.9474	-0.0112

#8	R 1	R 2	S 1
k=2	-0.1658	-0.1348	0.4725
k=5	-0.9682	-0.9393	0.8108
k=10	-1.0175	-1.1426	1.0523

#4	R 1	R 2	S 1
k=2	2.4645**	2.3931**	1.5119
k=5	1.3569	1.3448	0.4658
k=10	0.4329	0.2149	-0.1287

#9	R 1	R 2	S 1
k=2	-0.3477	0.1141	-0.5669
k=5	-1.011	-0.9422	-0.6728
k=10	-1.4951	-1.6032	-0.347

#5	R 1	R 2	S 1
k=2	-0.6925	-0.787	0.6614
k=5	-0.3433	-0.3526	1.4319
k=10	-0.1623	-0.2579	2.9497***

#10	R 1	R 2	S 1
k=2	-0.6417	-1.0796	0.6614
k=5	-1.2155	-1.2327	0.0863
k=10	-1.3202	-1.2418	0.1007

Note: * denotes a 10% level of significance, ** denotes a 5% level of significance and *** denotes a 1% level of significance

4.5.2.3 Chow Denning test

The Chow Denning test (Chow and Denning, 1993) for multiple variances is employed, in which the null hypothesis is that the series follows a random walk. The results are shown in Table 88 below. The null hypothesis is not rejected for all sub-samples at the 5% level, implying that these sub-samples have returns that are randomly generated. This result is contradictory to the one found using the daily ALSI returns series.

Table 88 - Chow Denning results for each sub-sample

Share Code	CD1	CD2
#1	2.5808**	1.3105
#2	2.3256*	1.4594
#3	2.6034**	2.1088
#4	2.1761*	2.3419*
#5	0.6377	0.6546
#6	1.5309	1.0297
#7	1.1344	0.9791
#8	1.2552	1.0401
#9	1.7577	1.3959
#10	1.2122	1.0887

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

4.5.2.4 Variance decomposition plots

Figure 31 and Figure 32 below shows the variance decomposition over time. In sub-samples 1, 2, 3 and 4 there is some evidence of significant variance ratios. However, this does not

remain so as the lags increase. The remaining sub-samples have insignificant variance ratios, implying that returns cannot be forecasted in these sub-samples.

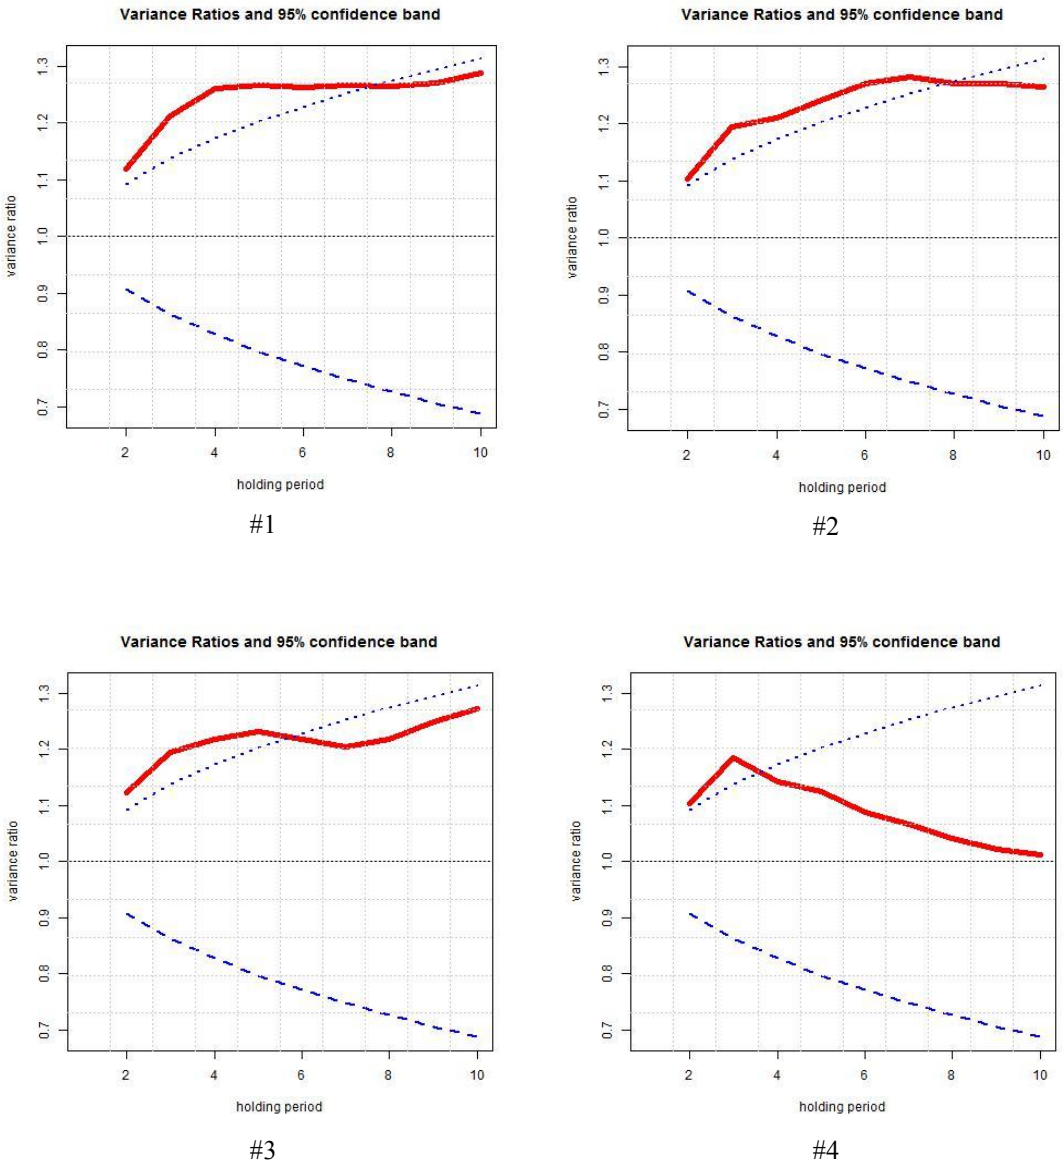
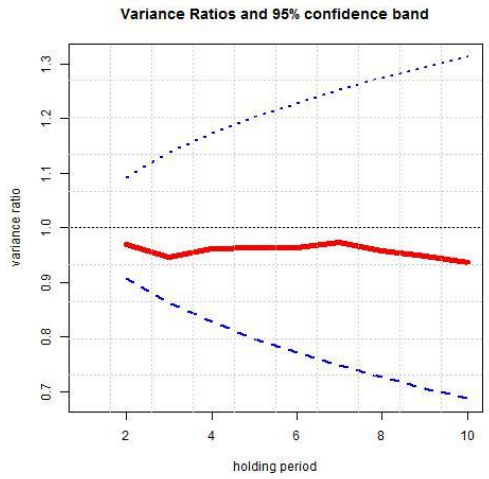
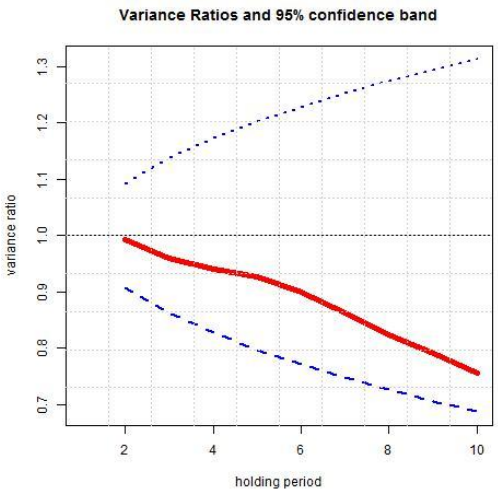


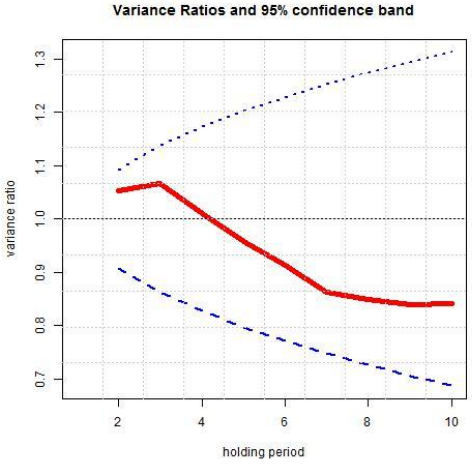
Figure 31 - Graphical representation of variance ratios of each sub-sample (1)



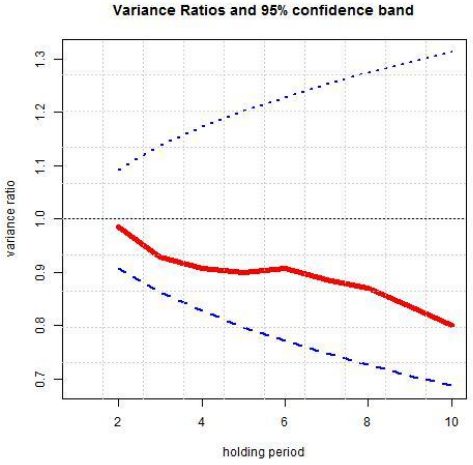
#5



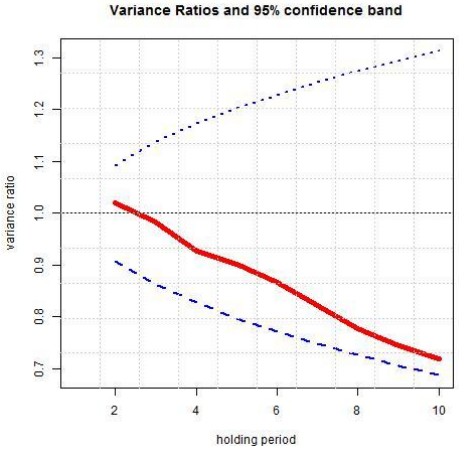
#6



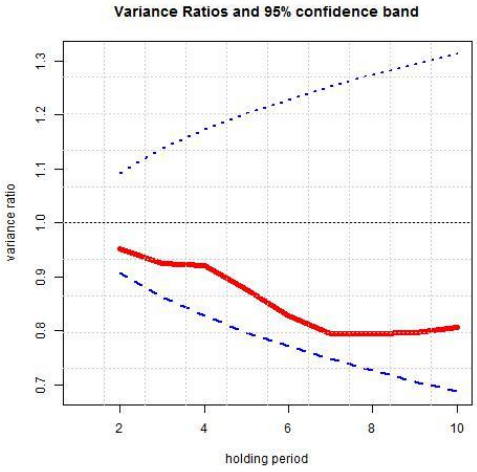
#7



#8



#9



#10

Figure 32 - Graphical representation of variance ratios of each sub-sample (2)

Therefore, according to visual evidence of the variance decomposition over time for each sub-sample, only four out of the ten shows evidence of predictability, albeit the evidence is weak. In comparison to the overall daily sample result, in that there is initial short term memory that fades at longer lags, the tests on each sub-sample provide a supportive result - the first four sub-samples had instances of non-random behaviour at lower lags.

4.5.2.5 Hurst exponent test

The Hurst exponent test was run on each of the ten sub-samples using the R/S method. Figure 33 below shows the result for the first four sub-samples. The null hypothesis of the test is that the exponent is equal to 0.5 (the series follows a random walk), under a 95% confidence interval of (0.4503, 0.6405). From the figure below, it can be concluded that the third sub-sample exhibits long term memory, with the remaining exhibiting random behaviour. The results presented are significant at the 5% level.

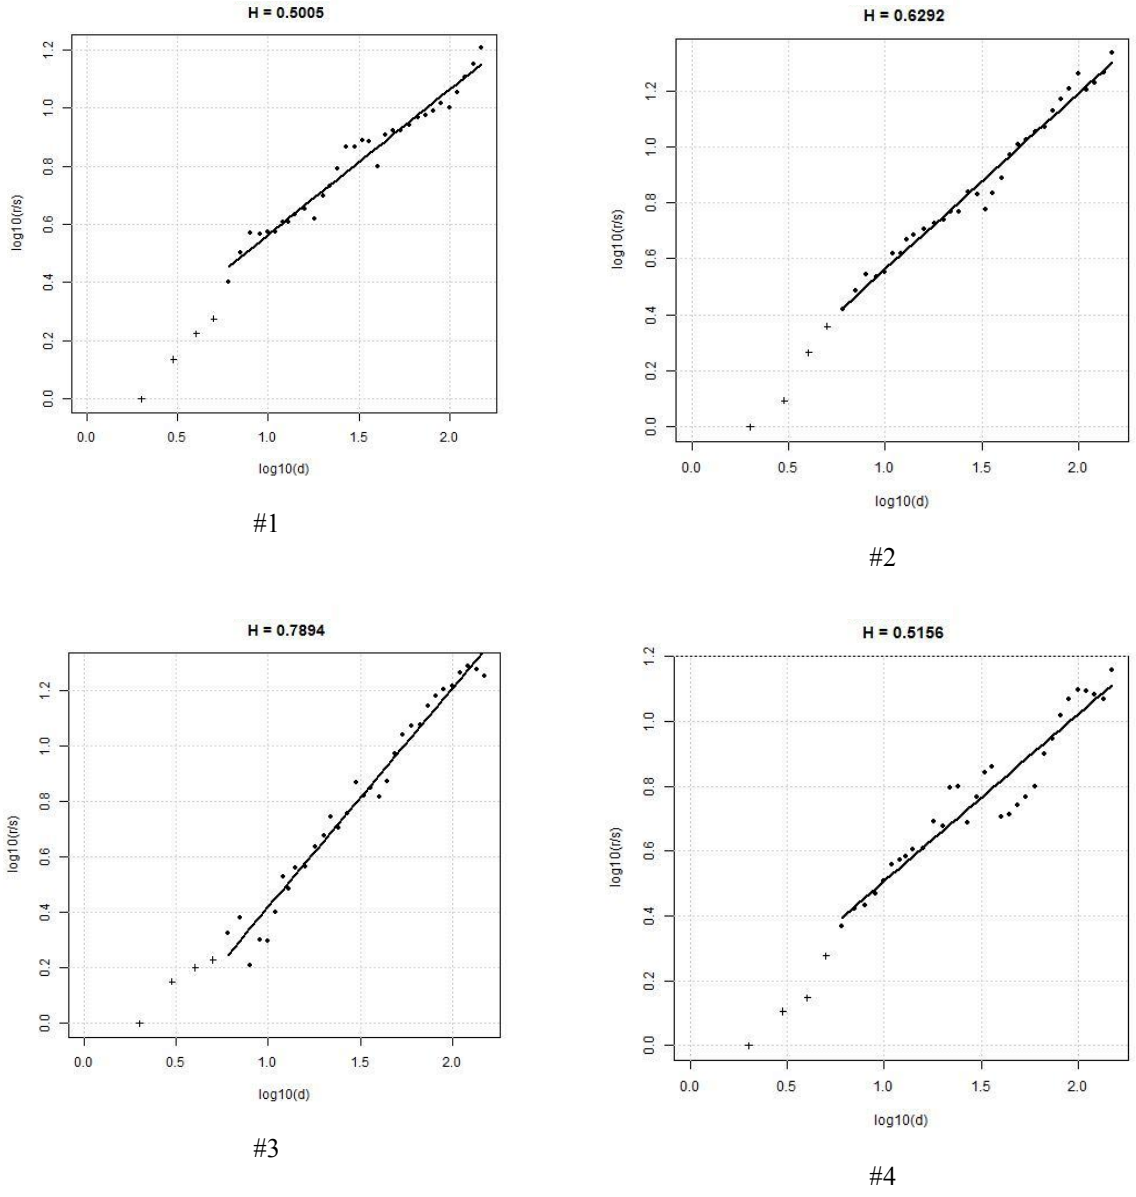
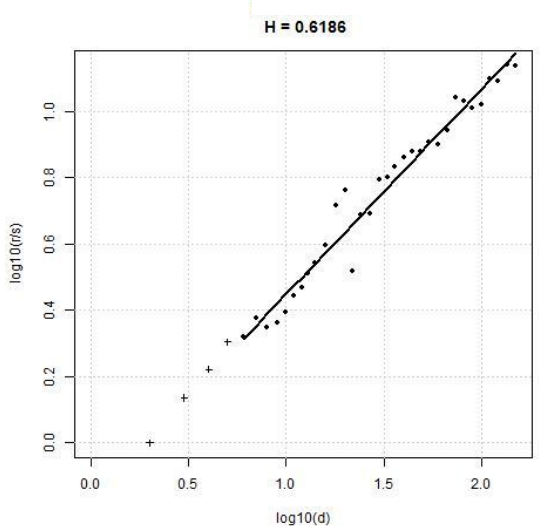
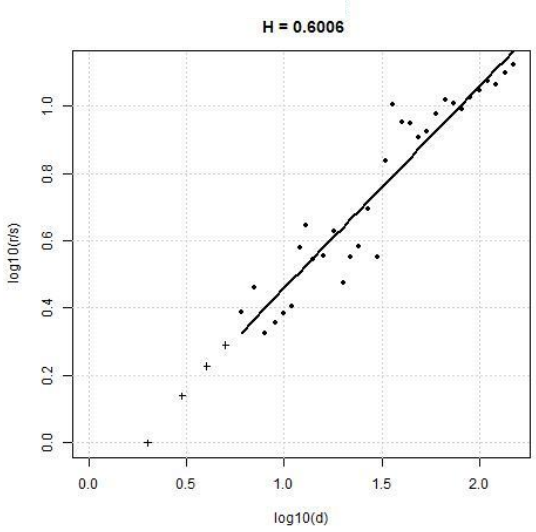


Figure 33 - Hurst exponents for the first set of four sub-samples

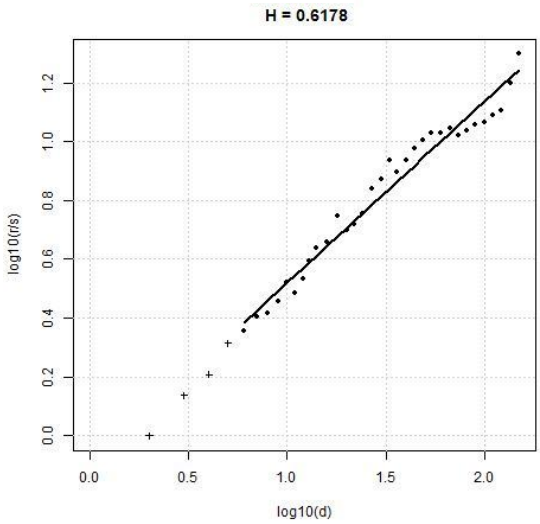
Figure 34 below shows the result for the second set of four sub-samples. The null hypothesis of the test is that the exponent is equal to 0.5 (the series follows a random walk). From the figure below, it can be concluded that the second set of four sub-samples all exhibit random behaviour. The results presented are significant at the 5% level.



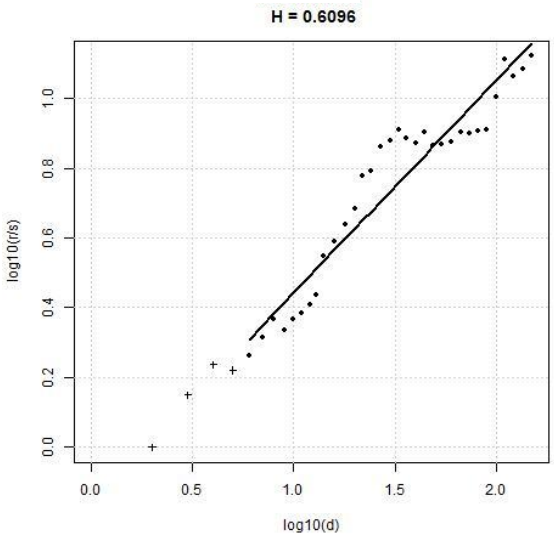
#5



#6



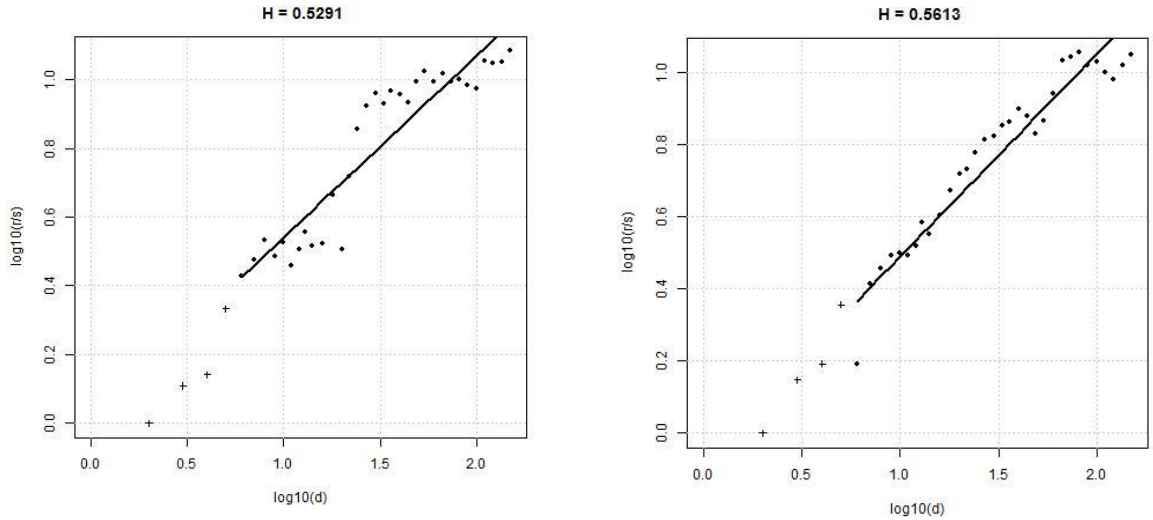
#7



#8

Figure 34 - Hurst exponents for the second set of four sub-samples

Figure 35 below shows the result for the last set of sub-samples. The null hypothesis of the test is that the exponent is equal to 0.5 (the series follows a random walk). From the figure below, it can be concluded that both exhibit random walk behaviour. The results presented are significant at the 5% level.



#9 #10

Figure 35 - Hurst exponents for the last set of sub-samples

Therefore, according to the Hurst exponent, the majority (nine) of the samples follow a random walk, with one sub-sample exhibiting non-random behaviour. The results are somewhat contradictory to that of the overall daily returns sample in that the daily ALSI was found to follow a random walk under the Hurst exponent. The contradictory results show that one can have non-random sequences of data in a series that is randomly generated.

4.5.3 Summary of results for the test of random walk behaviour

By investigation of the distributional properties of the daily, weekly and monthly ALSI share return series, it was found that: all three return series are non-normal, stationary and non-linear. These results provided a backdrop to investigate the random walk hypothesis on the equity market in South Africa. According to various parametric, non-parametric and graphical approaches used, one could build up a foundation on which to evaluate the accuracy of each test. By testing the weak form of the EMH using the Chow and Denning (1993) multiple variance ratio test, it was found that the daily return series is not randomly generated whereas the weekly and monthly return series are randomly generated.

However, by testing the weak form of the EMH using the Hurst exponent, it was concluded that the daily ALSI return series does follow a random walk, whereas the weekly and

monthly return series do not follow a random walk. While the Chow and Denning (1993) test does consider multiple variances, it still does not consider aspects from the non-parametric variance ratio test of Wright (2000), namely signs and ranks. Thus, one can rely more on the Hurst exponent due to its rolling sample approach and non-parametric nature.

The results from the Hurst exponent are intuitive as it points towards lower frequency data being more "predictable" than higher frequency data and can be explained by the interaction of investors and information. As information arrives, it is plausible that this information is not assimilated and reflected into share prices immediately. The daily fluctuations of share prices could be due to noise, whereas once information is interpreted correctly, investors tend to act in accordance, causing some trend to form on the return series.

A limitation of the approach used thus far relates to the range of tests available for normality, linearity, stationarity and testing the random walk hypothesis. One can easily argue that an alternative test should be used instead of those employed here. While reasonable attempts were made to provide robust results in the use of both parametric and non-parametric tests, it is easily conceivable that a new, more powerful test can be employed. Further, the division of the overall sample into equally sized sub-samples were made based on the size of the dataset alone. One can argue that this division could be done differently. However, the choice was to simply provide an alternative view of the results and can be improved by future research.

The increasing amount of attention given to emerging markets behoves researchers to test the most basic assumptions taken for granted by global finance academia. There is no intuitive reasoning to suspect that share prices in emerging markets, for example, do follow a random walk. In particular, many countries in Africa have stock exchanges that are quite young, with more exchanges to emerge in the near future. Given the plethora of hypotheses and theories in existence, it would be interesting to test these against the backdrop of an emerging exchange such as the JSE. A minor link back to literature discussed previously, the results thus far do correlate with Alagidede and Panagiotidis (2009) in that the authors find non-linearity in the daily returns generating process on the JSE.

4.6 Modelling the data generating process without additional factors

Given the results of the previous section, modelling the returns generating process is now attempted. Prior to using SETAR models to explain the returns generating process, an attempt is made to fit a simple ARIMA model to each of the returns series. This provides a foundation on which to base the success of the SETAR model. Further, none of the series showed signs of non-stationarity according to the tests performed previously, enabling one to proceed to model the returns. Further, due to the small number of observations for quarterly and semi-annual data, those results are not discussed here, but presented in the Appendix.

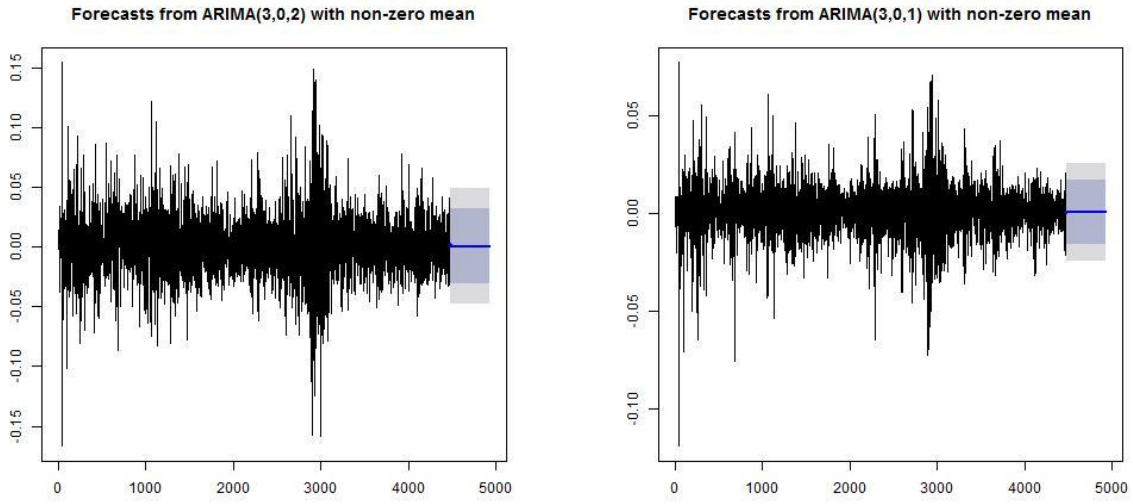
4.6.1 ARIMA models

The summary results of the ARIMA model for each daily return series is shown in Table 89 below. All models have an intercept and the MAPE value is considerably large. In theory, the MAPE value has no upper bound. However, for practical purposes, an upper bound of 100% is imposed. This implies that the closer the MAPE value to 100%, the poorer the fit of the model. All of the models are a poor fit according to the MAPE or even the R^2 criterion (as the R^2 is close to zero). An interesting observation emerges in that the majority of models have an intercept term, suggesting additional factors that could influence the returns process.

Table 89 - ARIMA for daily returns

Share code	ARIMA (p,q,r)	Intercept	R^2 (%)	MAPE (%)
BIL	0,0,0	Yes	2.38	100
MTN	0,0,1	Yes	2.61	100
SOL	0,0,3	Yes	2.30	100
FSR	2,0,1	Yes	2.12	100
SAB	1,0,2	Yes	1.77	100
NPN	3,0,0	Yes	2.47	100
AGL	3,0,2	Yes	2.44	100
J200	3,0,1	Yes	1.39	100
ALSI	3,0,1	Yes	1.26	100

Further, if one were to examine the forecasting ability of these models, the result is quite poor. The forecast of the AGL and ALSI models are provided in Figure 36 below.



AGL ALSI

Figure 36 - Forecasts of daily returns

The summary results of the ARIMA model for each weekly return series is shown in Table 90 below. Similar to the daily results, all models have an intercept and high MAPE. While the MAPE and R^2 values indicate a poor fit of the data, albeit the RMSE values are higher than their daily counterparts. The same observation of significant intercepts also applies in the case of the weekly returns series. Given the high MAPE values, forecasting results are not presented.

Table 90 - ARIMA models for weekly returns

Share code	ARIMA (p,q,r)	Intercept	R^2	MAPE
BIL	0,0,1	Yes	5.22	100
MTN	0,0,1	Yes	5.98	100
SOL	2,0,2	Yes	5.24	100
FSR	0,0,2	Yes	4.82	100
SAB	0,0,2	Yes	3.91	100
NPN	3,0,2	Yes	5.77	100
AGL	0,0,1	Yes	5.47	100
J200	0,0,0	Yes	3.11	100
ALSI	0,0,0	Yes	2.88	100

The summary results of the ARIMA model for each monthly return series is shown in Table 91 below. All models had an intercept and the MAPE value for each model is considerably large. Most of the models indicate that an ARIMA form in itself is not suitable as the ARIMA terms are all zero. The same observation of significant intercepts also applies in the case of the weekly returns series. Given the high MAPE values, forecasting results are not presented.

Table 91 - ARIMA models for monthly returns

Share code	ARIMA (p,q,r)	Intercept	R^2	MAPE
BIL	2,0,3	Yes	9.24	100
MTN	0,0,1	Yes	10.79	100
SOL	3,0,2	Yes	9.10	100
FSR	2,0,2	Yes	8.65	100
SAB	0,0,0	Yes	7.00	100
NPN	0,0,1	Yes	12.14	100
AGL	2,0,2	Yes	10.11	100
J200	2,0,2	Yes	5.97	100
ALSI	2,0,2	Yes	5.67	100

4.6.2 SETAR Test of linearity

As expected, a simple ARIMA model does not adequately capture the returns generating process of any of the data under investigation. Before attempting to fit a SETAR model to the data, the SETAR test is performed to determine if there are multiple regimes in the dataset. The SETAR test results for the daily returns data is displayed in Table 92 below. The first test examines a linear AR model against a SETAR model with one threshold (regime) and is labelled as "1vs2", whereas the second test examines a linear AR model against a SETAR model with two thresholds and is labelled as "1vs3". The results show that a SETAR model (either with one or two thresholds) is favoured over a linear AR model. The SETAR parameters and their significance will therefore lead to picking either a one or two threshold model. If both are found to be significant, the one with more regimes is chosen as this would fit the data better.

Table 92 - SETAR test for daily returns

BIL	Test Statistic	P-value
1vs2	69.9117	0.00***
1vs3	112.0303	0.00***

NPN	Test Statistic	P-value
1vs2	42.1421	0.00***
1vs3	77.1801	0.00***

MTN	Test Statistic	P-value
1vs2	58.9606	0.00***
1vs3	101.9386	0.00***

AGL	Test Statistic	P-value
1vs2	71.8768	0.00***
1vs3	110.8973	0.00***

SOL	Test Statistic	P-value
1vs2	57.6538	0.00***
1vs3	99.2896	0.00***

J200	Test Statistic	P-value
1vs2	45.2868	0.00***
1vs3	80.6688	0.00***

FSR	Test Statistic	P-value
1vs2	33.5100	0.00***
1vs3	62.5522	0.00***

ILV	Test Statistic	P-value
1vs2	47.7605	0.00***
1vs3	77.8584	0.00***

SAB	Test Statistic	P-value
1vs2	44.8145	0.00***
1vs3	72.0092	0.00***

ALSI	Test Statistic	P-value
1vs2	44.6253	0.00***
1vs3	86.6647	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

The SETAR test results for the weekly returns data is displayed in Table 93 below. The results show that a SETAR model is favoured over a linear AR model for all of the equities returns. However, for MDC and SAB, a single threshold model is favoured as opposed to a two threshold model for the remaining series. The SETAR parameters and their significance will therefore lead to picking either a one or two threshold model.

Table 93 - SETAR test for weekly returns

BIL	Test Statistic	P-value
1vs2	44.0224	0.00***
1vs3	74.9299	0.00***

NPN	Test Statistic	P-value
1vs2	62.1547	0.00***
1vs3	112.6071	0.00***

MTN	Test Statistic	P-value
1vs2	38.3828	0.00***
1vs3	76.7992	0.00***

AGL	Test Statistic	P-value
1vs2	40.0725	0.00***
1vs3	85.2041	0.00***

SOL	Test Statistic	P-value
1vs2	47.7922	0.00***
1vs3	72.8696	0.00***

J200	Test Statistic	P-value
1vs2	64.9358	0.00***
1vs3	97.1576	0.00***

FSR	Test Statistic	P-value
1vs2	62.8911	0.00***
1vs3	89.5761	0.00***

MDC	Test Statistic	P-value
1vs2	35.2596	0.00***
1vs3	55.4638	0.2

SAB	Test Statistic	P-value
1vs2	37.5445	0.00***
1vs3	55.8008	0.4

ALSI	Test Statistic	P-value
1vs2	64.5103	0.00***
1vs3	96.8455	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

The SETAR test results for the monthly returns data is displayed in Table 94 below. The results show that a SETAR model (with one threshold) is favoured over a linear AR model for BIL and FSR. The results of ILV point toward a linear AR model being preferred over a SETAR model. This result is interesting as previously it was shown that an ARIMA model on ILV returns was a poor fit. The next possible model to use for ILV returns would be a conditional variance (ARCH) model. The remaining four series favour some SETAR model over a linear AR counterpart. The SETAR parameters and their significance will therefore lead to picking either a one or two threshold model.

Table 94 - SETAR test on monthly returns

BIL	Test Statistic	P-value
1vs2	40.4084	0.00***
1vs3	82.2812	0.2

NPN	Test Statistic	P-value
1vs2	59.4276	0.00***
1vs3	122.6431	0.00***

MTN	Test Statistic	P-value
1vs2	63.0379	0.00***
1vs3	97.8842	0.00***

AGL	Test Statistic	P-value
1vs2	57.3277	0.00***
1vs3	97.1541	0.00***

SOL	Test Statistic	P-value
1vs2	60.0938	0.00***
1vs3	106.4460	0.00***

J200	Test Statistic	P-value
1vs2	73.5673	0.00***
1vs3	129.1300	0.00***

FSR	Test Statistic	P-value
1vs2	53.9565	0.00***
1vs3	97.5631	0.2

BGA	Test Statistic	P-value
1vs2	41.4190	0.4
1vs3	67.5253	0.2

SAB	Test Statistic	P-value
1vs2	53.6157	0.00***
1vs3	113.7620	0.00***

ALSI	Test Statistic	P-value
1vs2	81.6789	0.00***
1vs3	126.5855	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

4.6.3 SETAR model

Having established that some form of SETAR model is appropriate for all except the monthly ILV return series, attempts to model these returns are now made. The SETAR model requires a set of starting values and a number of regimes in which to begin the modelling process. This selection is done via the R command "selectSETAR", in which the command conducts a grid search over user specified values of *gamma* and *d*, along with the number of regimes. This input is then used in the SETAR modelling procedure to derive coefficients of the SETAR model. An example of the model output using the daily ALSI returns is presented in Table 95 below. The output does not contain any significant terms apart from the high regime constant, indicating that the SETAR model is a poor fit.

Table 95 - SETAR model for daily ALSI returns

ALSI	Estimate	Standard Error	T-Statistic	P-Value
Constant (low regime)	-0.0005	0.0004	-1.2136	0.22
$\varphi_{low,1}$	0.0004	0.0301	0.0126	0.99
$\varphi_{low,2}$	0.0180	0.0224	0.8010	0.42
Constant (high regime)	0.0013	0.0004	3.5537	0.00***
$\varphi_{high,1}$	0.0235	0.0295	0.7961	0.43
$\varphi_{high,2}$	0.0117	0.0202	0.5782	0.56
Residuals Variance	0.0001596		MAPE	100%

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

The number of significant coefficients, significant constants and the MAPE for each model is presented in Table 96 below, with the remaining shares presented in the Appendix. In the case of the SETAR models under daily returns, it is found that all models have considerably high MAPEs, along with few significant coefficients. Indeed, there are cases where the intercept terms are significant, indicative of further unknown factors that may play a role in explaining that particular return generating process. In contrast, a model with no significant coefficients or intercept terms points towards an alternate model form that is required.

Table 96 - SETAR model diagnostics for selected shares using daily returns

Share code	Number of significant coefficients	Number of significant intercepts	MAPE
BIL	1	0	100
MTN	1	1	100
SOL	2	2	100
FSR	2	0	100
SAB	1	1	100
NPN	1	2	100
AGL	2	0	100
J200	0	2	100
ALSI	0	1	100

The success of the SETAR model using daily returns was found to be ineffective - a plausible conclusion given that the daily ALSI return series followed a random walk under the Hurst exponent. Attention is now drawn to deriving SETAR models using weekly returns. An

example of the model output using weekly ALSI returns is shown in Table 97 below. It is again found that the SETAR model is a poor fit as it has a high MAPE along with no significant coefficients or intercepts. Similarly, given the random walk conclusion of the weekly ALSI return series, the inability of a SETAR model to fit the data is a plausible finding.

Table 97 - SETAR model for weekly ALSI returns

ALSI	Estimate	Standard Error	T-Statistic	P-Value
Constant (low regime)	-0.0004	0.0021	-0.18	0.86
$\varphi_{low,\ 1}$	-0.1036	0.0713	-1.45	0.15
$\varphi_{low,\ 2}$	0.1102	0.0480	2.30	0.02
Constant (high regime)	0.0061	0.0019	3.18	0.00
$\varphi_{high,\ 1}$	-0.0924	0.0673	-1.37	0.17
$\varphi_{high,\ 2}$	-0.0282	0.0472	-0.60	0.55
Residuals Variance	0.0008202		MAPE	100%

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

The number of significant coefficients, significant constants and the MAPE for each model is presented in Table 98 below, with the remaining shares presented in the Appendix. In the case of the SETAR models under weekly returns, it is found that all models have considerably high MAPEs, along with few significant coefficients. Indeed, there are cases where the intercept terms are significant, indicative of further unknown factors that may play a role in explaining that particular return generating process. In contrast, a model with no significant coefficients or intercept terms points towards an alternate model form that is required.

Table 98 - SETAR model diagnostics for selected shares using weekly returns

Share code	Number of significant coefficients	Number of significant intercepts	MAPE
BIL	2	2	100
MTN	1	1	100
SOL	1	1	100
FSR	2	1	100
SAB	1	1	100
NPN	1	2	100
AGL	0	1	100
J200	2	1	100
ALSI	1	1	100

The results of the SETAR model using monthly ALSI returns are presented in Table 99 below. The model does have a significant constant in the high regime , albeit at the 10% level of significance. However, given the high MAPE, this model is also not suitable to explain the monthly ALSI return generating process.

Table 99 - SETAR model for monthly ALSI returns

ALSI	Estimate	Standard Error	T-Statistic	P-Value
Constant (low regime)	0.0142	0.0096	1.48	0.14
$\varphi_{low,1}$	-0.0605	0.1567	-0.39	0.70
$\varphi_{low,2}$	-0.0574	0.1158	-0.50	0.62
Constant (high regime)	0.0168	0.0092	1.82	0.07*
$\varphi_{high,1}$	-0.0631	0.1518	-0.42	0.68
$\varphi_{high,2}$	0.0429	0.0899	0.48	0.63
Residuals Variance	0.003205		MAPE	100%

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

The number of significant coefficients, significant constants and the MAPE for each model is presented in Table 100 below. Also recall that the output for ILV is not available, as the SETAR test showed that a regime changing model was not suited to this return series. In the case of the SETAR models under monthly returns, it is found that all models have considerably high MAPEs, along with few significant coefficients. Indeed, there are cases where the intercept terms are significant, indicative of further unknown factors that may play

a role in explaining that particular return generating process. In contrast, a model with no significant coefficients or intercept terms points towards an alternate model form that is required.

Table 100 - SETAR model diagnostics for selected shares using monthly returns

Share code	Number of significant coefficients	Number of significant intercepts	MAPE
BIL	0	1	100
MTN	1	0	100
SOL	0	1	100
FSR	1	1	100
SAB	1	1	100
NPN	1	0	100
AGL	1	1	100
J200	1	1	100
ALSI	0	0	100

The results of the SETAR model using quarterly ALSI returns are presented in Table 101 below. The model does not have any significant coefficients, perhaps due to either the small sample size, or the inability of the model to fit the return generating process. This is corroborated by the high MAPE.

Table 101 - SETAR model for quarterly ALSI returns

ALSI	Estimate	Standard Error	T-Statistic	P-Value
Constant (low regime)	0.0347	0.0339	1.02	0.31
$\varphi_{low,1}$	-0.0882	0.2749	-0.32	0.75
$\varphi_{low,2}$	0.1786	0.2364	0.76	0.45
Constant (high regime)	0.0253	0.0293	0.87	0.39
$\varphi_{high,1}$	0.1519	0.3030	0.50	0.62
$\varphi_{high,2}$	-0.1246	0.1525	-0.82	0.42
Residuals Variance	0.007703		MAPE	100%

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

The number of significant coefficients, significant constants and the MAPE for each model is presented in Table 102 below. While some models do have significant coefficients, four out of the ten securities below have no significant coefficients or intercepts. This implies that the

return series cannot be modelled by a SETAR specification. Further, the MAPE for each model is high, implying a poor fit. Again, the small sample size needs to be considered.

Table 102 - SETAR model diagnostics for selected shares using quarterly returns

Share code	Number of significant coefficients	Number of significant intercepts	MAPE
BIL	0	2	100
MTN	0	0	100
SOL	2	1	100
FSR	2	1	100
SAB	0	0	100
NPN	0	0	100
AGL	2	1	100
J200	0	0	100
ALSI	0	0	100

The results of the SETAR model using semi-annual ALSI returns are presented in Table 103 below. The model does have a significant constant in the high regime as well as a significant coefficient in the low regime. However, given the high MAPE, this model is also not suitable to explain the return generating process.

Table 103 - SETAR model for semi-annual ALSI returns

ALSI	Estimate	Standard Error	T-Statistic	P-Value
Constant (low regime)	-0.1295	0.1192	-1.09	0.29
$\varphi_{low,1}$	-1.2480	0.5962	-2.09	0.05
$\varphi_{low,2}$	-0.3569	0.4885	-0.73	0.47
Constant (high regime)	0.1164	0.0510	2.28	0.03
$\varphi_{high,1}$	-0.1601	0.2992	-0.53	0.60
$\varphi_{high,2}$	-0.0537	0.1685	-0.32	0.75
Residuals Variance	0.01023		MAPE	100%

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

The number of significant coefficients, significant constants and the MAPE for each model is presented in Table 104 below. Again, while the models do have either significant intercepts or significant coefficients, the high MAPE implies a poor fit. Here, the issue of sample size is particularly important, as there are only 38 observations.

Table 104 - SETAR model diagnostics for selected shares using semi-annual returns

Share code	Number of significant coefficients	Number of significant intercepts	MAPE
BIL	0	1	100
MTN	1	0	100
SOL	0	1	100
FSR	0	0	100
SAB	1	1	100
NPN	1	0	100
AGL	0	0	100
J200	1	0	100
ALSI	1	1	100

In summary, it was found that all five frequency SETAR models were a poor fit to the data. While some models had statistically significant constants, indicating that additional factors are required to model the returns generating process. As the SETAR model is an improvement on a simple ARIMA model, it is still insufficient to capture the complexities of the data used in this study. As such, it is inappropriate for modelling that particular returns process, as it is not governed by a dynamic non-linear functional form. Indeed, the results of Seetharam and Britten (forthcoming) show that an improvement of the SETAR model, a STAR model, does provide a better fit to the data, but there are still significant intercept terms. The next improvement to these time series models would be to use a functional form that is not specified in advance. All returns series are now evaluated using models that are not specified *a priori*, namely neural networks.

4.6.4 Neural network modelling results

To compare the results of the SETAR model, and to provide a case for the inclusion of exogenous variables, a non-linear autoregressive (NAR) neural network model was run. Prior to this, a simple feed-forward network was considered, but discarded due to the poor fit of the model from a conceptual and empirical standpoint. The NAR model is conceptually similar to

a SETAR model, with the advantages of neural network models discussed previously (namely the non-specification of a functional form).

All of the neural networks in this study are trained in Matlab using the Levenberg-Marquadt algorithm and the performance of the network was based on the mean-squared error, with six iterations used to determine this performance (and therefore terminate training). This algorithm is a more sophisticated version of the non-linear least squares method used in regression analysis. The results of the ALSI are discussed in detail, with the result of the remaining networks being in the Appendix. The results of the NAR model using daily ALSI returns are shown below. The network converged after 15 iterations (Figure 37), with the best performance at the 9th iteration. In other words, the network took 15 attempts to model the daily ALSI return generating process. While there is no autocorrelation in the residuals (Figure 38), the R^2 of the network is only 12.23% according to Figure 38. This is an improvement over the ARIMA and SETAR models, yet it is still not adequate to use a NAR to explain the return process. Conceptually, the higher R^2 can be attributed to a function form that was not specified *a priori*. The next evolution of this model would be to include exogenous inputs.

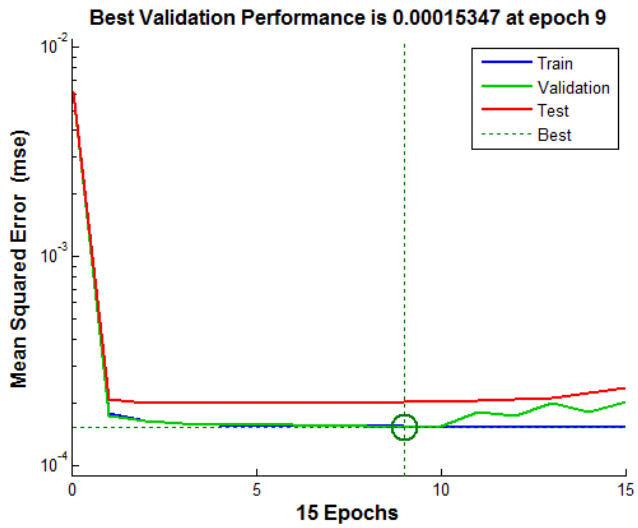


Figure 37 - Performance of the NAR network using daily returns

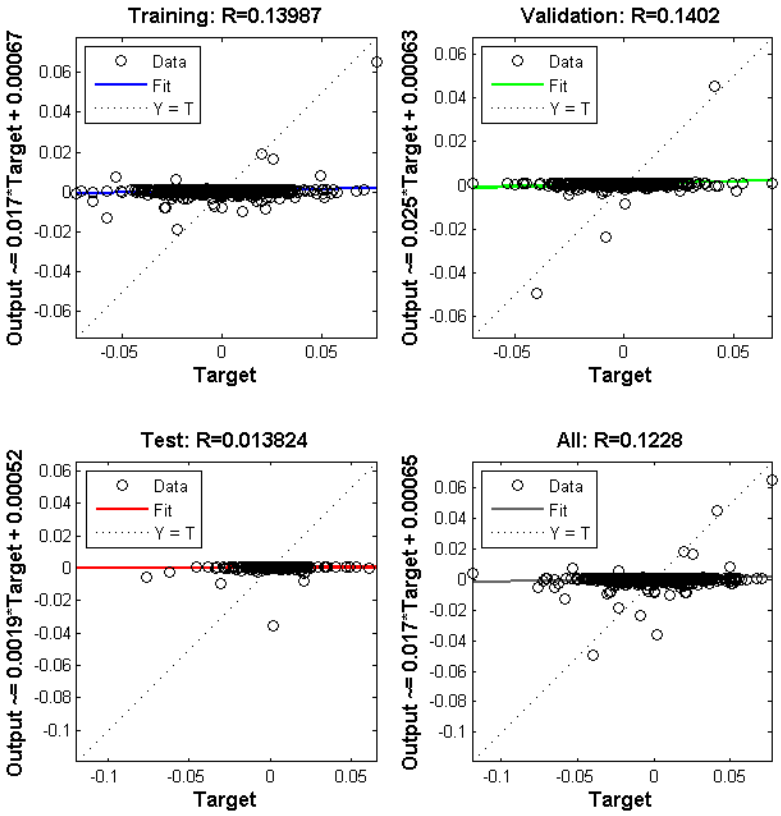


Figure 38 - R^2 of the NAR network using daily returns

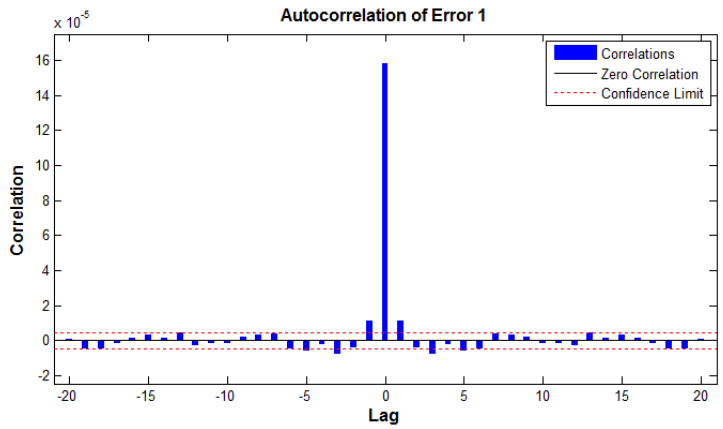


Figure 39 - Autocorrelation of the error terms of the NAR network using daily returns

Table 105 below provides a summary of the NAR model for the remaining daily series. The variability in the hidden nodes parameter points towards some series having more complex data generating processes than others. For example, BIL and MTN have the highest number of hidden nodes, indicating the most complexity in these return series (recall that Basheer (1998) suggests that one hidden layer is sufficient to approximate continuous functions, whereas Masters (1994) suggests two hidden layers for discontinuous functions). However,

the overall R^2 are low, indicating that while the NAR NN is better than the SETAR model, it still does not model the returns series well.

Table 105 - NAR model diagnostics for selected shares using daily returns

Share Code	Iterations	Best Iteration	Hidden Nodes; Delay Parameter (h,d)	Test R^2 (%)	Overall R^2 (%)	Autocorrelation in error terms
BIL	14	8	4,1	29.45	10.92	No
MTN	15	9	4,1	10.22	10.06	No
SOL	21	15	2,1	10.33	12.08	No
FSR	90	84	2,1	14.80	7.01	No
SAB	34	28	3,1	20.11	10.32	
NPN	12	6	1,1	12.73	9.07	
AGL	10	4	2,1	5.07	5.81	
J200	18	12	3,1	2.05	13.07	
ALSI	15	9	2,2	1.38	12.28	No

The results of the NAR model using weekly ALSI returns are shown below. The network converged after 10 iterations (Figure 40), with the best performance at the fourth iteration. In other words, the network took 10 attempts to model the daily ALSI return generating process. While there is no autocorrelation in the residuals (Figure 42), the R^2 of the network is only 17.32% according to Figure 41. This is an improvement over the ARIMA and SETAR models, yet it is still not adequate to use a NAR to explain the return process. Conceptually, the higher R^2 can be attributed to a function form that was not specified *a priori*. The next evolution of this model would be to include exogenous inputs.

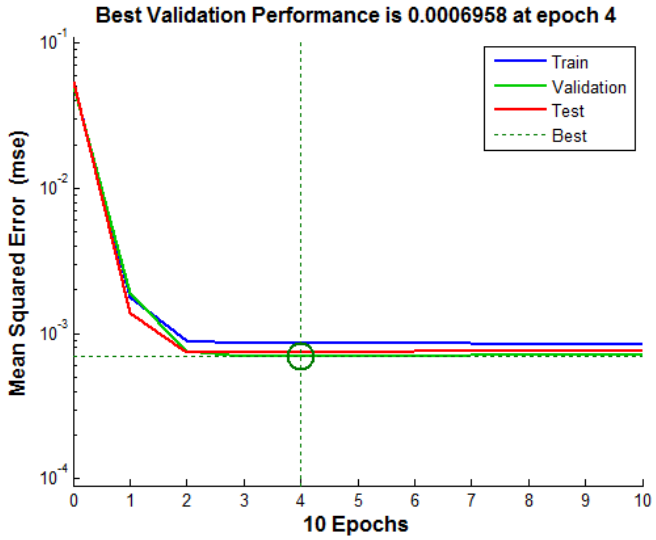


Figure 40 - Performance of the NAR network using weekly returns

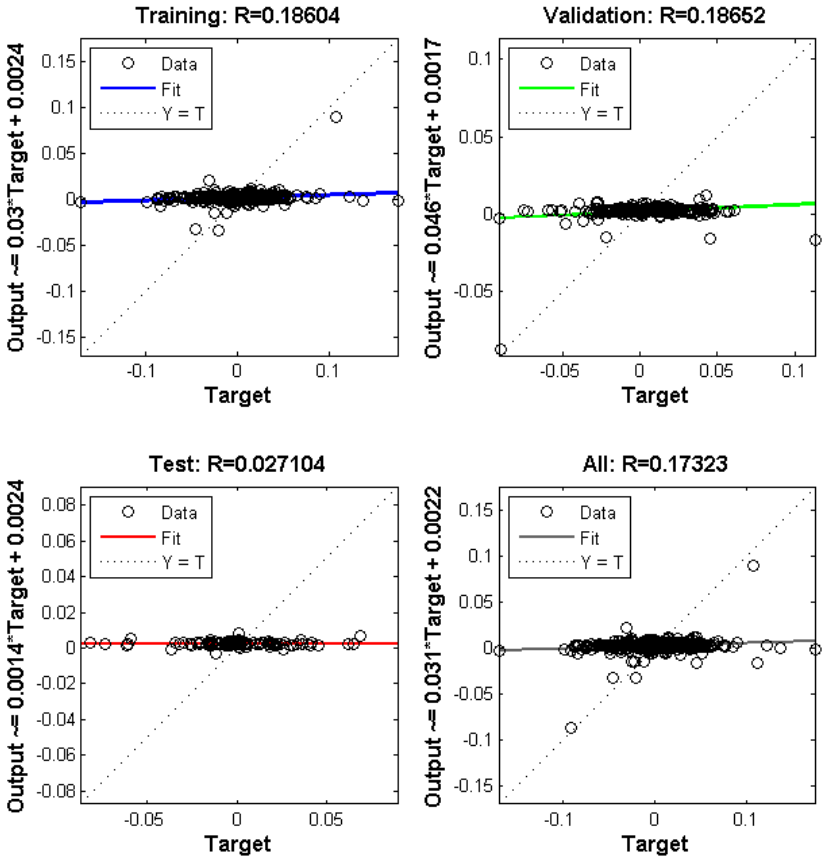


Figure 41 - R^2 of the NAR network using weekly returns

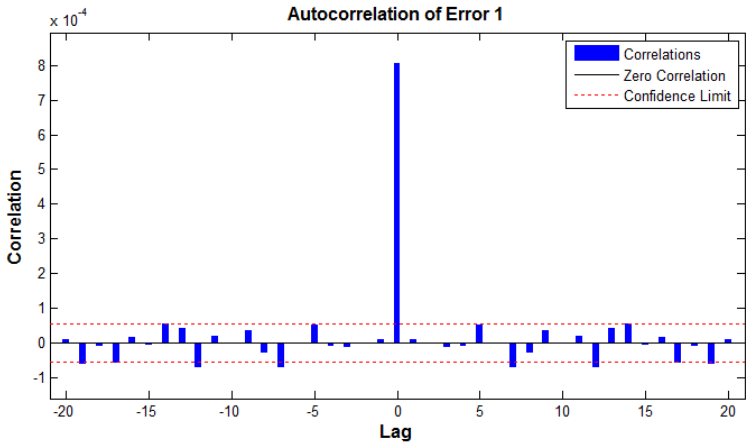


Figure 42 - Autocorrelation of the error terms of the NAR network using weekly returns

Table 106 below provides a summary of the NAR model for the remaining weekly series. In contrast to the trend found using daily returns, the number of iterations for the ALSI is relatively low compared to some of other series. Further, the complexity of FSR, WBO and the ALSI is higher than that of the remaining series as they have the highest number of hidden nodes. As with the daily return series, the overall R^2 for each share return is still low, indicating possible additional factors that could influence returns.

Table 106 - NAR model diagnostics for selected shares using weekly returns

Share Code	Iterations	Best Iteration	Hidden Nodes; Delay Parameter (h,d)	Test R^2 (%)	Overall R^2 (%)	Autocorrelation in error terms
BIL	9	3	1,1	12.36	22.51	No
MTN	11	5	2.1	5.11	17.16	No
SOL	8	2	1,1	28.59	16.10	No
FSR	10	4	3,1	37.44	22.90	No
SAB	12	6	2,1	14.02	10.92	No
NPN	19	13	10,1	2.30	30.77	No
AGL	7	1	2,1	25.21	10.59	No
J200	9	3	10,1	11.36	27.61	No
ALSI	10	4	3,1	2.71	17.32	No

The results of the NAR model using monthly ALSI returns are shown below. The network converged after nine iterations (Figure 43), with the best performance at the third iteration. In

other words, the network took nine attempts to model the daily ALSI return generating process. While there is no autocorrelation in the residuals (Figure 45), the R^2 of the network is only 23.87% according to Figure 44. This is an improvement over the ARIMA and SETAR models, yet it is still not adequate to use a NAR to explain the return process. Conceptually, the higher R^2 can be attributed to a function form that was not specified *a priori*. The next evolution of this model would be to include exogenous inputs.

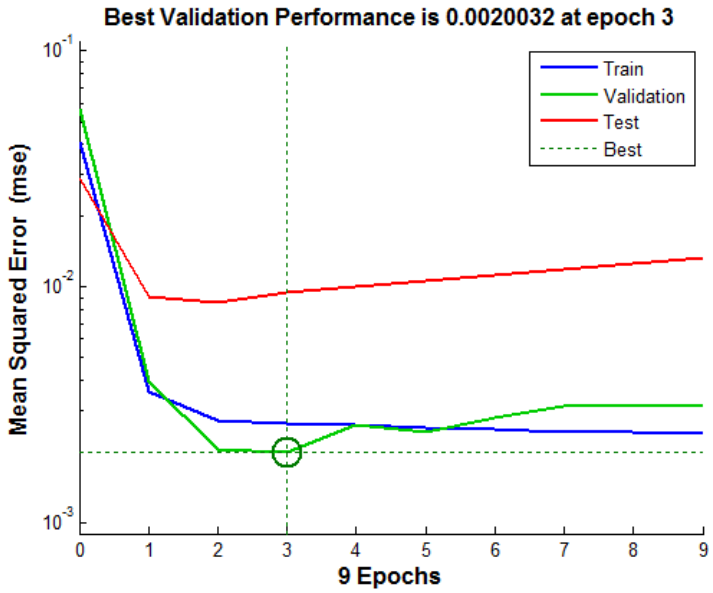


Figure 43 - Performance of the NAR network using monthly returns

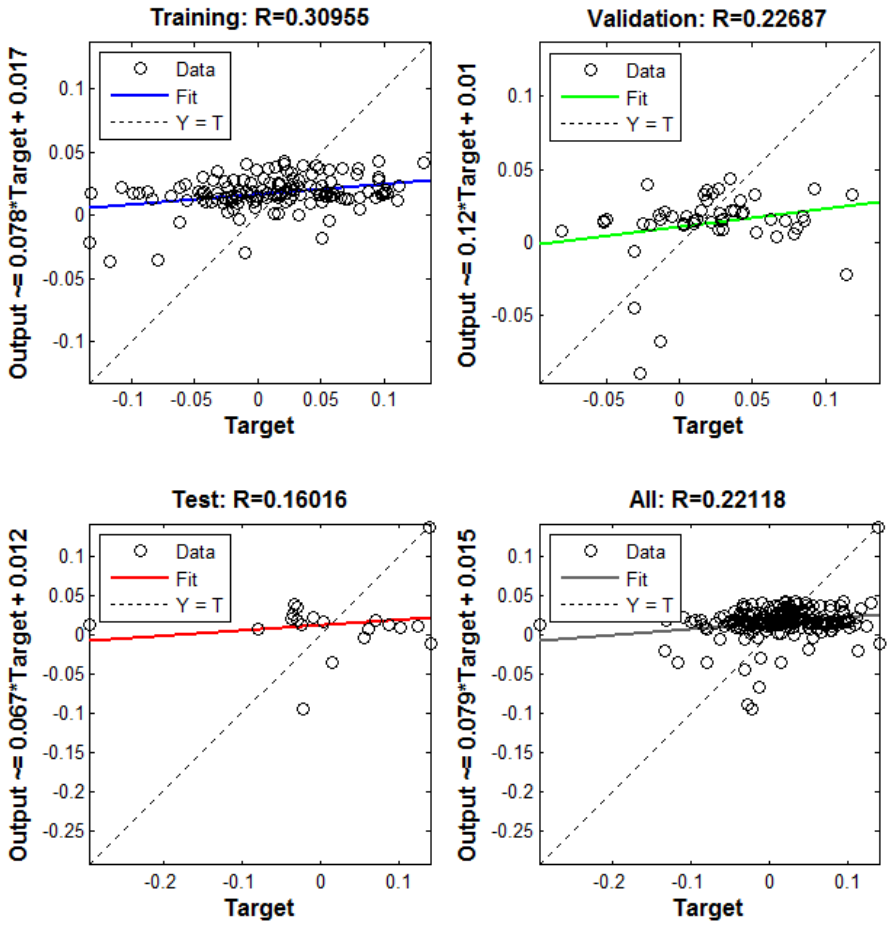


Figure 44 - R^2 of the NAR network using monthly returns

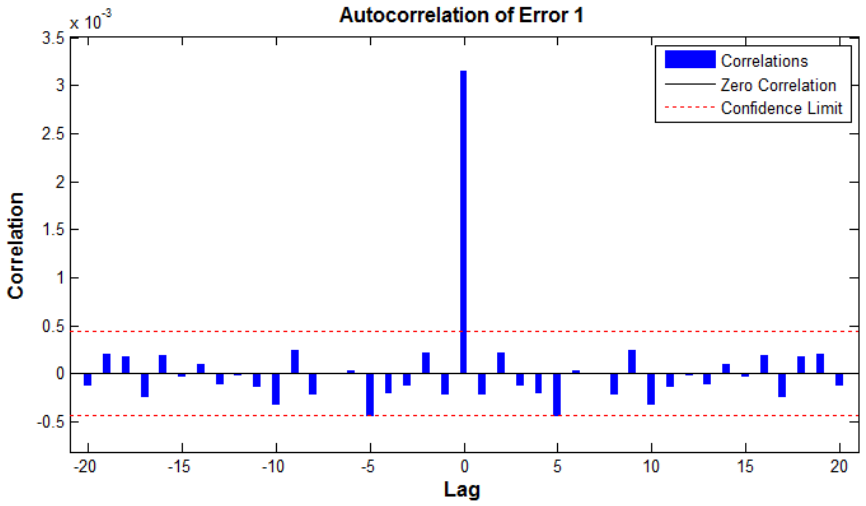


Figure 45 - Autocorrelation of the error terms of the NAR network using monthly returns

Table 107 below provides a summary of the NAR model for the remaining monthly series. Similar to the trend found using weekly returns, the number of iterations for the ALSI is relatively low compared to other series. Further, the complexity of the ALSI is higher than that of the remaining series as it has the highest number of hidden nodes. It is interesting to note that the R^2 for the monthly series is higher compared to their weekly and daily counterparts. This implies that lower frequency data is easier to model that higher frequency data.

Table 107 - NAR model diagnostics for selected shares using monthly returns

Share Code	Iterations	Best Iteration	Hidden Nodes; Delay Parameter (h,d)	Test R^2 (%)	Overall R^2 (%)	Autocorrelation in error terms
BIL	13	7	2,1	26.71	37.60	No
MTN	8	2	1,1	31.24	25.98	No
SOL	7	1	2,1	22.18	22.22	No
FSR	14	8	2,1	19.19	45.69	No
SAB	7	1	10,1	19.42	16.71	No
NPN	8	2	10,1	30.23	33.15	No
AGL	7	1	2,1	25.21	10.59	No
J200	9	3	10,1	11.36	27.61	No
ALSI	9	3	3,1	16.02	22.12	No

In summary, using a NAR network to specify the return generating process of the daily, weekly and monthly ALSI return series did result in a model that was no better (at times poorer) than the SETAR counterparts. The one improvement however was in the monthly NAR NN. Previously, a SETAR model was a poor fit to the monthly return data, indicative of either an alternate model being required. Thus, the existence of a NAR NN using monthly return data does point towards the return process being more complex than the models used thus far. The results of the NAR using quarterly and semi-annual data are not presented here.

While the SETAR model is generally good at modelling regime changes, given the data used in this study, either the SETAR model is ill-equipped to the data, or the data is too complex to be modelled by a simple regime changing model. This avenue of inquiry is explored by modelling the return series using artificial intelligence and is discussed later. Given the objective of this thesis – to present a case for cyclical efficiency – it is necessary to examine the power of the above autoregressive models under a different sampling method. As such, the autoregressive models are now presented under the sub-samples of the ALSI.

4.6.5 ALSI Sub-sample results

As the SETAR model methodology effectively first tests for the fit of a linear AR model, for each sub-sample, an attempt was made to fit an ARIMA model to the daily return series. The summary results of the ARIMA model for each sub-sample return series is shown in Table 108 below. Both the MAPE and R^2 indicate a poor fit to the data. An interesting observation should be noted in that for the latter half of the sub-samples, no ARIMA coefficients or an intercept appears to be significant, pointing towards a randomly generated series. There is no discernible relationship between a particular sub-sample being found to be randomly (or not randomly) generated and the lack of an ARIMA model fit to the data.

Table 108 - ARIMA model for each sub-sample

Sub-sample	ARIMA (p,q,r)	Intercept	R^2 (%)	MAPE (%)
#1	0,0,2	No	0.00	100
#2	2,0,2	No	1.13	100
#3	0,0,1	No	1.26	100
#4	1,0,2	No	1.04	100
#5	0,0,0	Yes	0.91	100
#6	0,0,0	Yes	1.25	100
#7	0,0,0	No	2.04	100
#8	0,0,0	No	1.09	100
#9	0,0,0	Yes	0.92	100
#10	0,0,0	No	0.84	100

The SETAR test results for each sub-sample are displayed in Table 109 below. The first test examines a linear AR model against a SETAR model with one threshold, whereas the second test examines a linear AR model against a SETAR model with two thresholds. The results show that a SETAR model is favoured over a linear AR model for all sub-samples except that of sub-samples 4 and 7. The SETAR parameters and their significance will therefore lead to picking either a one or two threshold model, with the exception of the two sub-samples mentioned. While sub-sample 4 did have an ARIMA model with lags on both the autoregressive and moving average terms, the model was a poor fit. Indeed, sub-sample 7 did could not be fit to any ARIMA model. These results, in conjunction with the SETAR tests, point towards either volatility models being used or towards a non-specified linear AR model being used.

Table 109 - SETAR test on each sub-sample

#1	Test Statistic	P-value
1vs2	52.3963	0.00***
1vs3	84.5385	0.00***

#2	Test Statistic	P-value
1vs2	52.0156	0.00***
1vs3	84.1592	0.00***

#3	Test Statistic	P-value
1vs2	26.0985	0.00***
1vs3	60.6853	0.00***

#4	Test Statistic	P-value
1vs2	21.3012	0.4
1vs3	46.3230	0.6

#5	Test Statistic	P-value
1vs2	39.4198	0.00***
1vs3	82.7977	0.00***

#6	Test Statistic	P-value
1vs2	49.2332	0.00***
1vs3	83.9205	0.00***

#7	Test Statistic	P-value
1vs2	29.3066	0.2
1vs3	61.0855	0.2

#8	Test Statistic	P-value
1vs2	37.6693	0.00***
1vs3	66.9051	0.00***

#9	Test Statistic	P-value
1vs2	43.4419	0.00***
1vs3	71.8188	0.00***

#10	Test Statistic	P-value
1vs2	37.9706	0.00***
1vs3	63.8646	0.00***

Note: * denotes a 10% level of statistical significance, ** denotes a 5% level of statistical significance and *** denotes a 1% level of statistical significance.

The number of significant coefficients, significant constants and the MAPE for each model is presented in Table 110 below. In the case of the SETAR models under daily returns, it is found that all models have considerably high MAPEs, along with few significant coefficients.

Indeed, there are cases where the intercept terms are significant, indicative of further unknown factors that may play a role in explaining that particular return generating process. In contrast, a model with no significant coefficients or intercept terms points towards an alternate model form that is required.

Table 110 - SETAR model diagnostics for each sub-sample

Sub-sample	Significant coefficients	Intercepts	MAPE
#1	2	1	100
#2	3	1	100
#3	1	1	100
#4	N/A	N/A	N/A
#5	1	1	100
#6	2	2	100
#7	N/A	N/A	N/A
#8	0	0	100
#9	1	0	100
#10	1	1	100

Table 111 below provides a summary of the NAR model for the sub-samples. The number of hidden and delay nodes seems to fluctuate, with the former indicating changing complexity and the latter indicating some form of memory. Similarly, the overall R^2 seems to fluctuate over each sub-sample, but still remains relatively low. This implies that additional factors should be considered apart from lagged returns.

Table 111 - NAR model diagnostics for each sub-sample

Sub-sample	Time period of sample	Hidden, Delay	Number of epochs	Training R^2 (%)	Validation R^2 (%)	Testing R^2 (%)	Overall R^2 (%)
#1	Sep 1997 to May 1999	2,2	11	19.99	14.31	2.32	16.32
#2	May 1999 to Feb 2001	1,1	8	12.39	3.98	27.25	12.17
#3	Feb 2001 to Oct 2002	1,1	15	10.11	25.49	22.12	14.49
#4	Oct 2002 to Jul 2004	1,2	10	13.91	13.01	37.13	16.30
#5	Jul 2004 to Mar 2006	1,1	17	7.26	15.95	9.41	9.72
#6	Apr 2006 to Dec 2007	2,1	17	22.72	2.31	21.94	18.85
#7	Dec 2007 to Sep 2009	2,1	11	8.15	19.41	19.22	10.54
#8	Sep 2009 to May 2011	1,1	13	21.59	7.12	17.48	18.81
#9	May 2011 to Feb 2013	2,2	9	8.40	7.55	16.52	8.65
#10	Feb 2013 to Oct 2014	2,2	12	15.66	11.67	3.71	14.03

4.6.6 Summary of results for modelling returns without additional factors

In applying the SETAR model to daily, weekly and monthly returns data, it was found that the fit of this family of models differs heavily based on the frequency of data used. The monthly SETAR model was the model to have a significant coefficient as well as a significant constant. The latter point towards additional factors that need to be included in explaining the monthly ALSI returns process. Applying the SETAR methodology to each sub-sample, it was found that two out of the ten sub-samples could not have a SETAR model fitted to them (the returns process was not non-linear). The first of these sub-samples was found to be randomly generated under the Hurst exponent, which makes the lack of a SETAR model fit plausible. The other sub-sample, however, was found to not be randomly generated. That finding, in conjunction with the lack of a SETAR model fit, points towards some alternate model form that is required.

The results support the findings of van Gysen, Huang and Kruger (2013), in that non-linear models are more suitable to modelling South Africa equity returns. As stated earlier, the time period under examination for SETAR modelling has a heavy influence on the results. This time period (and thus the sample size) can be extended to include previous financial anomalies (such as the political shift in South Africa during 1992 to 1994). Forecasting can also be done via the recursive modelling procedure, which will assist in providing forecasts over a longer time horizon.

Smith and Dyakova (2014) show that several African markets experience good and poor periods of predictability, with the South African market (the JSE) being in the latter group. This result, while supportive of those in this thesis, focuses instead on predictability compared to explanation, the latter of which is the viewpoint taken in this thesis. The authors conclude that this widely varying degree of predictability provides evidence in favour of the Adaptive Market Hypothesis of Lo (2004, 2005). It could quite well be the case that a SETAR model, modified to include additional (lagged) factors may have better predictive ability. Overall, this sub-section has provided some evidence as to the non-linearity of share prices in South Africa and the results appear promising for future research. This sub-section has established that a particular non-linear model of daily returns of the ALSI is less than adequate in capturing non-linearities present in the data. However, it is believed that a more robust model can be used. Specifically, the NAR NN was barely able to obtain a good fit of the data, which is perhaps not surprising considering that more than one hidden layer was found in the optimal network (multiple hidden layers point towards a discontinuous function in the dependent variable). Thus, the proceeding sub-section attempts to model the ALSI using both endogenous and exogenous variables, similar to the Arbitrage Pricing Theory framework of Ross (1976).

4.7 Modelling the data generating process with additional factors

The NARX network was implemented in Matlab^{TM} and is represented in Figure 46 below. The input data are passed to a number of hidden layers, with the output also being passed backpropagated. The NARX network therefore allows for information to flow in both

directions before reaching the output layer. Also recall that the ALSI dataset (all three frequencies) is the only variable used in this section. As such, the results do not relate to any of the other shares or equity indices previously analysed.

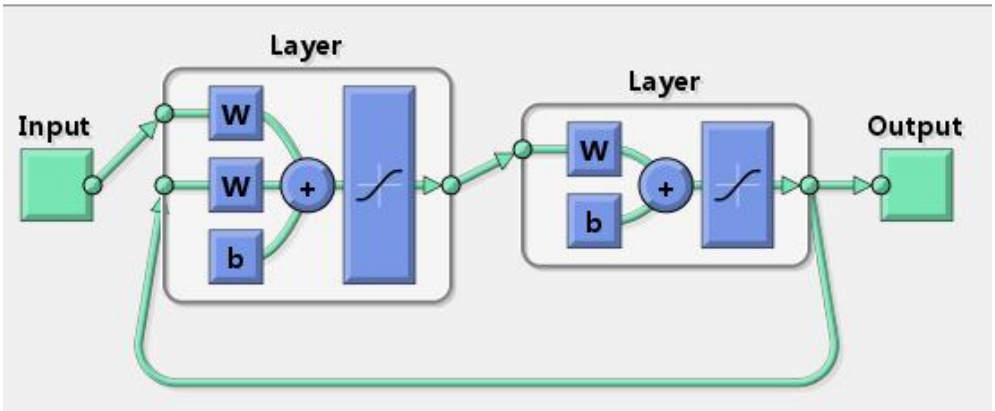


Figure 46 – Representation of the NARX network in Matlab

4.7.1 Daily return sample results

The input data is transformed using the tangens hyperbolicus (*tanh*) function and passed to each layer of the network (in this case, two layers) with possible feedback between the layers according the gradient descent method. Recall that this method finds the path of "least resistance" in that it selects the steepest descent at each iteration, finding either a minimum or infinitely decreasing path. In addition, the Levenberg-Marquadt algorithm is used to provide a means of generating the error term. This algorithm is a more sophisticated version of the non-linear least squares method used in regression analysis. A summary of the parameters or methods of the network are shown in Table 112 below.

Table 112 – Overview of methods used in training the daily NARX network

Parameter	Method or Value
Training	Levenberg-Marquadt algorithm
Learning	Gradient Descent method
Performance	Mean Squared Error criterion
Number of hidden layers	2
Number of neurons	2
Transfer Function	Tangens hyperbolicus (tanh)

The network was trained using 65% of the sample, validated on 25% and tested on 10% as per Looney (1996). Convergence of the training and validation MSEs were reached after 47 epochs. In other words, it took 17 attempts for the network to most accurately learn the relationship between the dependent and independent variables. The performance of the network is shown in Figure 47 below. The network converged after 11 epochs (training cycles) as the minimum mean squared error (MSE) was reached at this point. The figure below shows that the MSE begins to increase after 17 epochs, thus training was stopped to avoid overfitting.

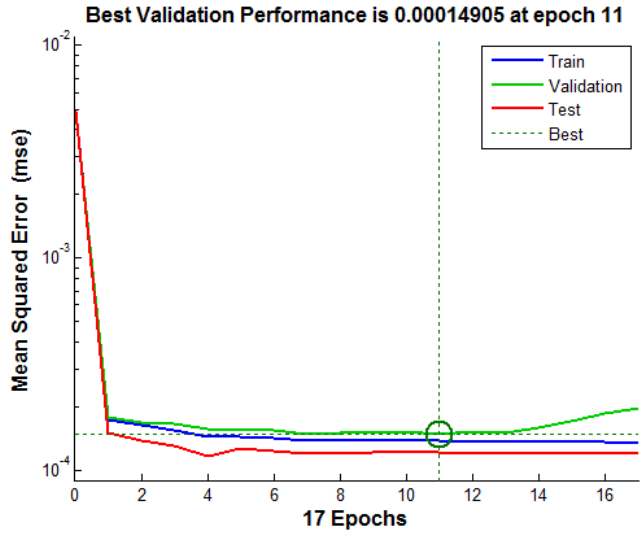


Figure 47 - Depiction of MSE over time, using daily ALSI data

The network was able to correctly fit 37% of the data, as given by the overall R^2 of the model in Figure 48. The training phase had the highest R^2 of 38%, followed by the validation phase of 37%. This is considerably higher than the SETAR and NAR models and can be attributed to both the time series nature of the network along with the additional inputs. Examining the autocorrelation of the error terms does not reveal any serial correlation between the error terms, which does not cause concern.

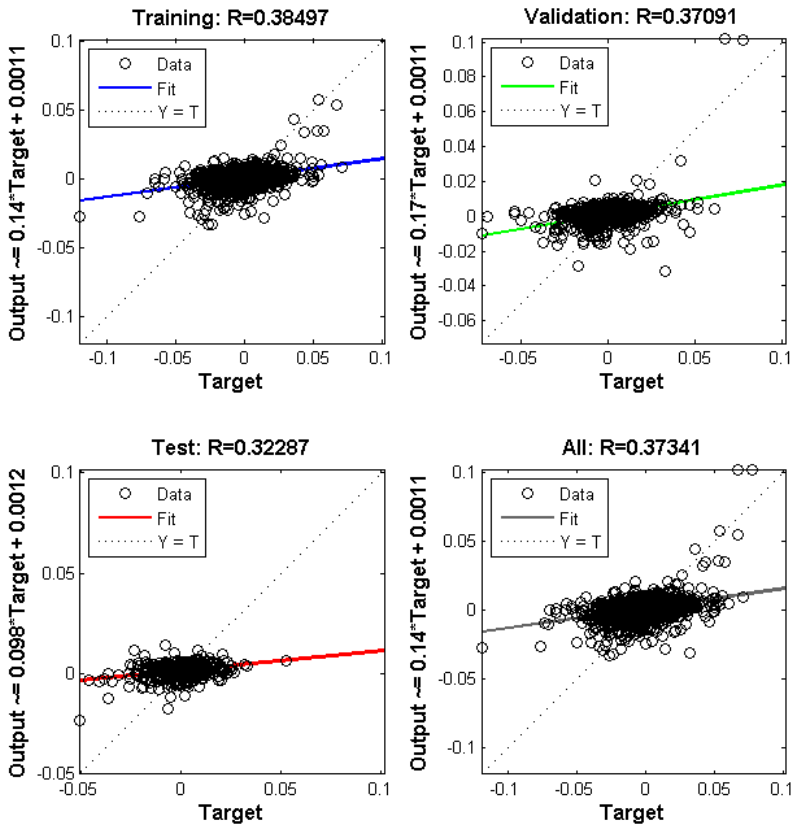


Figure 48 - Comparison of R^2 during training, validation and testing phases

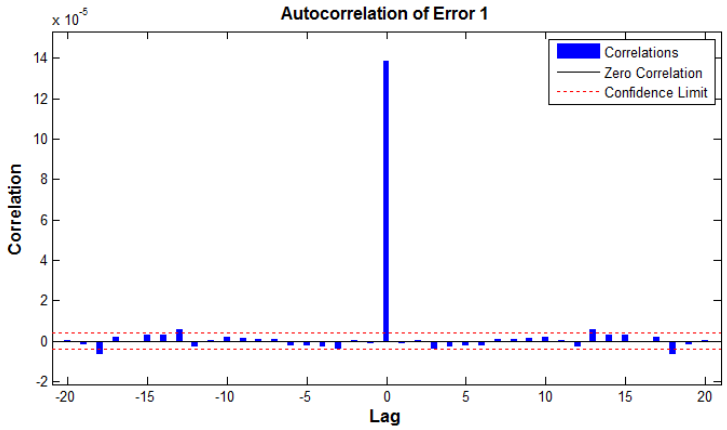


Figure 49 - Autocorrelation function over time

4.7.2 Weekly return sample results

The input data is transformed using the tangens hyperbolicus (tanh) function and passed to each layer of the network (in this case, three layers) with possible feedback between the layers according the gradient descent method. Recall that this method finds the path of "least resistance" in that it selects the steepest descent at each iteration, finding either a minimum or infinitely decreasing path. In addition, the Levenberg-Marquadt algorithm is used to provide a means of generating the error term. This algorithm is a more sophisticated version of the non-linear least squares method used in regression analysis. A summary of the parameters or methods of the network are shown in Table 113 below.

Table 113 – Overview of methods used in training the weekly NARX network

Parameter	Method or Value
Training	Levenberg-Marquadt algorithm
Learning	Gradient Descent method
Performance	Mean Squared Error criterion
Number of hidden layers	3
Number of neurons	3
Transfer Function	Tangens hyperbolicus (tanh)

The network was trained using 65% of the sample, validated on 25% and tested on 10% as per Looney (1996). Convergence of the training and validation MSEs were reached after five epochs. In other words, it took five attempts for the network to most accurately learn the relationship between the dependent and independent variables. The performance of the network is shown in Figure 50 below. The network converged after five epochs (training cycles) as the minimum mean squared error (MSE) was reached at this point. The figure below shows that the MSE begins to increase after three epochs, thus training was stopped to avoid overfitting.

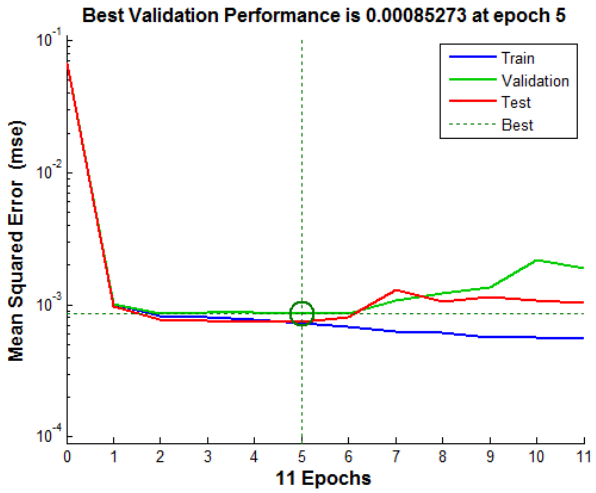


Figure 50 - Depiction of MSE over time, using weekly ALSI data

Examining the fit of the network during each phase of training, testing and validation, one finds that the overall fit of the model is 32% (the bottom right quadrant of Figure 51 below). This implies that the network was able to correctly explain 32% of weekly returns over the sample period, implying that there is a (small) degree of inefficiency in the ALSI. This is again seen to be higher than the previous SETAR and NAR models, primarily due to the inclusion of additional inputs in explaining the data generating process.

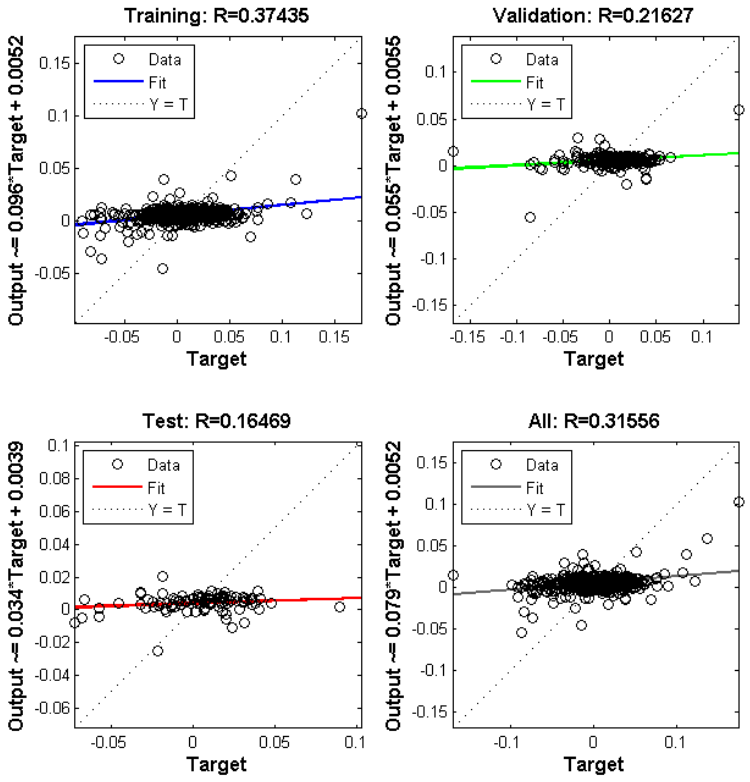


Figure 51 - Comparison of R^2 during training, validation and testing phases

An examination of the autocorrelation coefficients of the error term does not reveal any hindrances with the model. Indeed, the errors appear stationary.

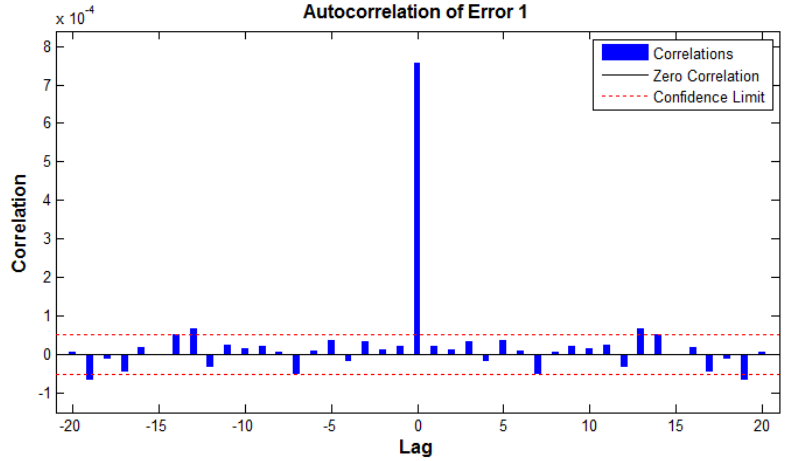


Figure 52 - Autocorrelation function over time

4.7.3 Monthly return sample results

The input data is transformed using the tangens hyperbolicus (tanh) function and passed to each layer of the network (in this case, two layers) with possible feedback between the layers according the gradient descent method. Recall that this method finds the path of "least resistance" in that it selects the steepest descent at each iteration, finding either a minimum or infinitely decreasing path. In addition, the Levenberg-Marquadt algorithm is used to provide a means of generating the error term. This algorithm is a more sophisticated version of the non-linear least squares method used in regression analysis. A summary of the parameters or methods of the network are shown in Table 114 below.

Table 114 – Overview of methods used in training the monthly NARX network

Parameter	Method or Value
Training	Levenberg-Marquadt algorithm
Learning	Gradient Descent method
Performance	Mean Squared Error criterion
Number of hidden layers	2
Number of neurons	2
Transfer Function	Tangens hyperbolicus (tanh)

The network was trained using 65% of the sample, validated on 25% and tested on 10% as per Looney (1996). Convergence of the training and validation MSEs were reached after six epochs. In other words, it took 11 attempts for the network to most accurately learn the relationship between the dependent and independent variables. The performance of the network is shown in Figure 53 below. The network converged after six epochs (training cycles) as the minimum mean squared error (MSE) was reached at this point. The figure below shows that the MSE begins to increase after six epochs, thus training was stopped to avoid overfitting.

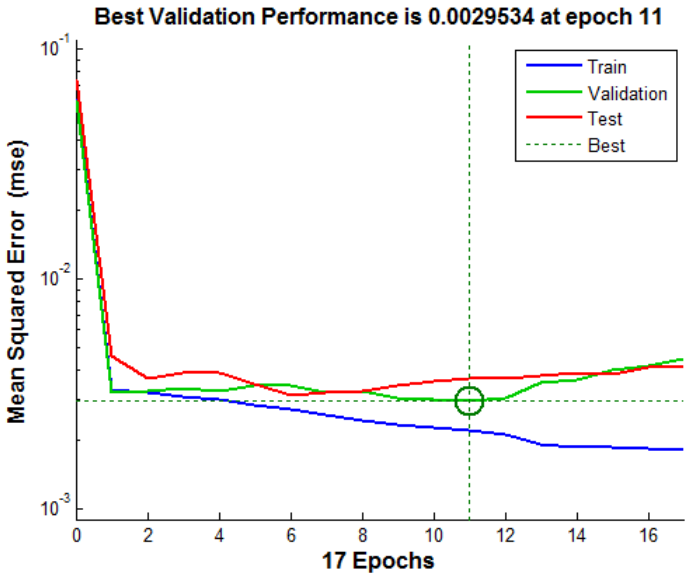


Figure 53 - Depiction of MSE over time, using monthly ALSI data

Examining the fit of the network during each phase of training, testing and validation, one finds that the overall fit of the model is 47% (the bottom right quadrant of Figure 54 below). This figure implies that the network was able to correctly explain 47% of daily returns over the sample period, implying that there is a degree of inefficiency in the ALSI. As was the case with the daily and weekly frequency data, the monthly data used in the NARX network provided a higher goodness of fit compared to the SETAR and NAR models.

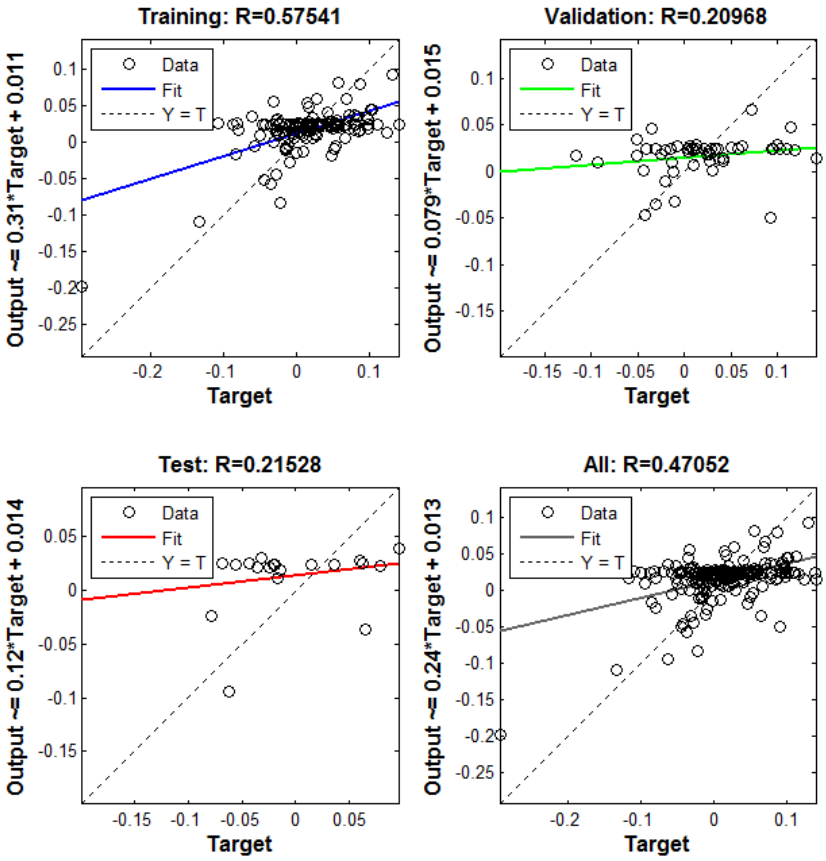


Figure 54 - Comparison of R^2 during training, validation and testing phases

An examination of the autocorrelation coefficients of the error term does not reveal any hindrances with the model. Indeed, the errors appear stationary.

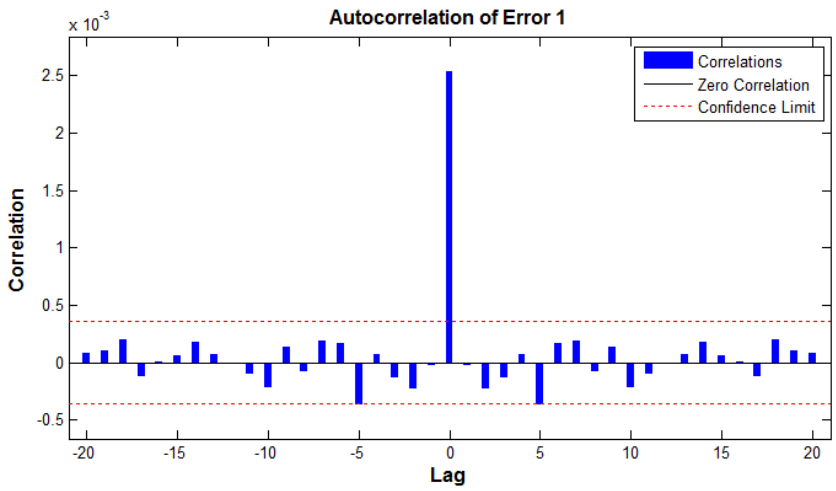


Figure 55 - Autocorrelation function over time

4.7.4 Quarterly return sample results

A summary of the parameters or methods of the network are shown in Table 115 below. The final network had two hidden layers as well as two neurons.

Table 115 – Overview of methods used in training the quarterly NARX network

Parameter	Method or Value
Training	Levenberg-Marquadt algorithm
Learning	Gradient Descent method
Performance	Mean Squared Error criterion
Number of hidden layers	2
Number of neurons	2
Transfer Function	Tangens hyperbolicus (tanh)

The network was trained using 65% of the sample, validated on 25% and tested on 10% as per Looney (1996). Convergence of the training and validation MSEs were reached after two epochs. In other words, it took eight attempts for the network to most accurately learn the relationship between the dependent and independent variables. The performance of the network is shown in Figure 56 below. The network converged after two epochs (training cycles) as the minimum mean squared error (MSE) was reached at this point. The figure below shows that the MSE begins to increase after two epochs, thus training was stopped to avoid overfitting.

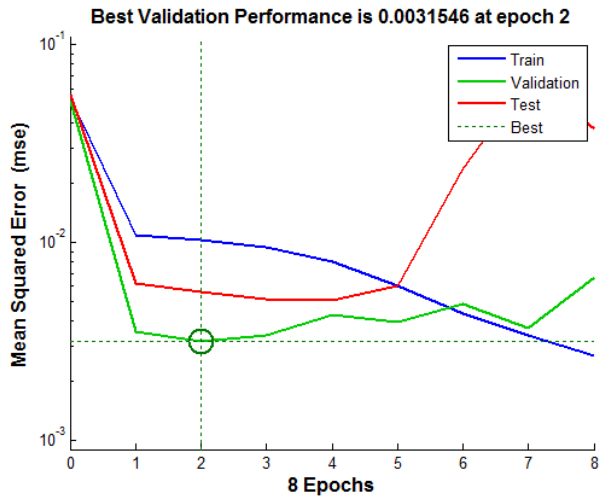


Figure 56 - Depiction of MSE over time, using quarterly ALSI data

Examining the fit of the network during each phase of training, testing and validation, one finds that the overall fit of the model is 37% (the bottom right quadrant of Figure 57 below). This figure implies that the network was able to correctly explain 37% of quarterly returns over the sample period, implying that there is a degree of inefficiency in the ALSI. This is lower than the corresponding monthly NARX network by 10%, implying that either the NARX model is a worse fit to the quarterly data, or that there are not enough observations.

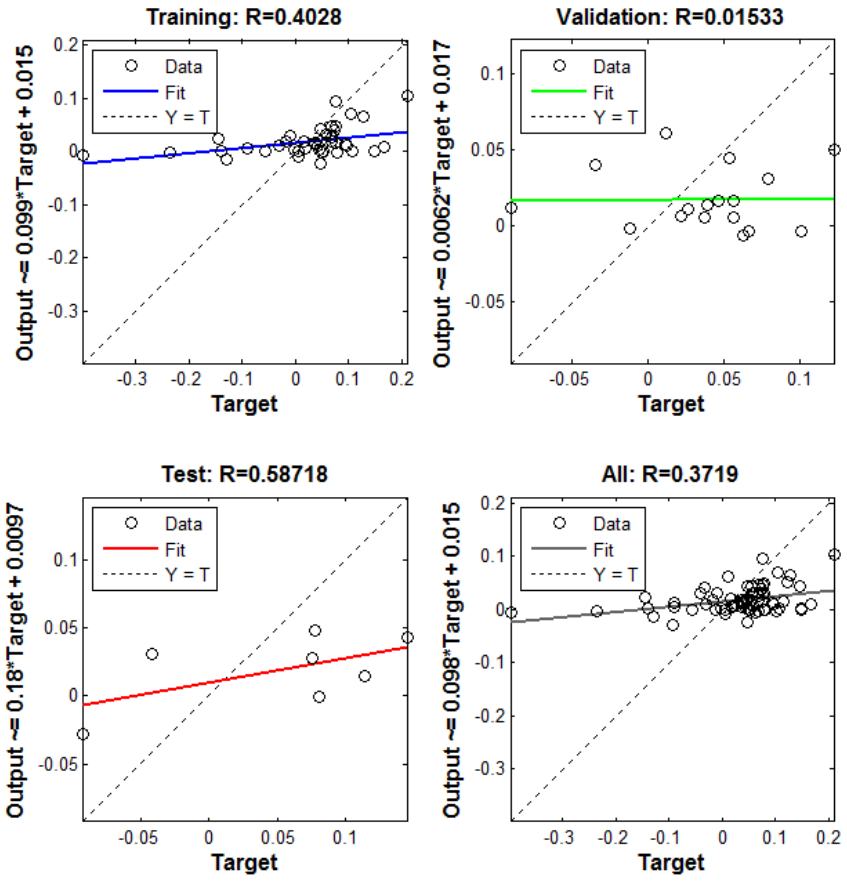


Figure 57 - Comparison of R^2 during training, validation and testing phases

An examination of the autocorrelation coefficients of the error term does not reveal any hindrances with the model. Indeed, the errors appear stationary.

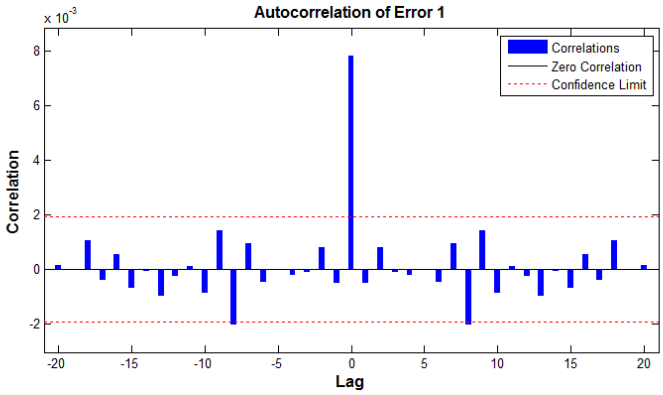


Figure 58 - Autocorrelation function over time

4.7.5 Semi-annual return sample results

A summary of the parameters or methods of the network are shown in Table 116 below. The network had three delay nodes and one hidden layer.

Table 116 – Overview of methods used in training the semi-annual NARX network

Parameter	Method or Value
Training	Levenberg-Marquadt algorithm
Learning	Gradient Descent method
Performance	Mean Squared Error criterion
Number of hidden layers	1
Number of neurons	3
Transfer Function	Tangens hyperbolicus (tanh)

The network was trained using 65% of the sample, validated on 25% and tested on 10% as per Looney (1996). Convergence of the training and validation MSEs were reached after four epochs. In other words, it took 10 attempts for the network to most accurately learn the relationship between the dependent and independent variables. The performance of the network is shown in Figure 59 below. The network converged after four epochs (training cycles) as the minimum mean squared error (MSE) was reached at this point. The figure below shows that the MSE begins to increase after six epochs, thus training was stopped to avoid overfitting.

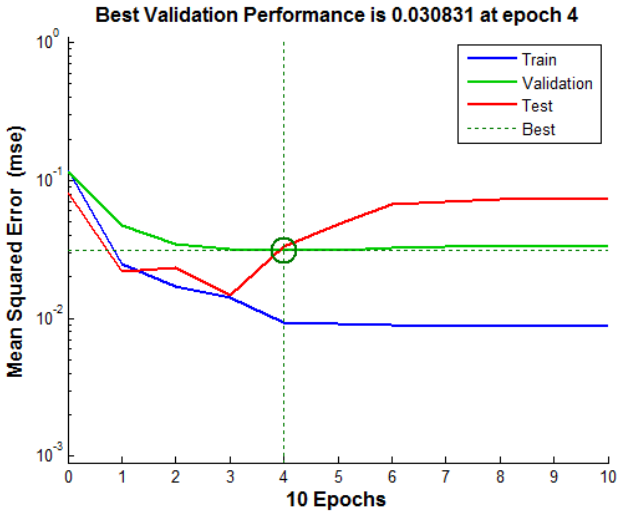


Figure 59 - Depiction of MSE over time, using semi-annual ALSI data

Examining the fit of the network during each phase of training, testing and validation, one finds that the overall fit of the model is 66% (the bottom right quadrant of Figure 60 below). This figure implies that the network was able to correctly explain 66% of daily returns over the sample period, implying that there is a larger degree of inefficiency in the ALSI. This goodness of fit is quite high in relation to the other models, which is either due to the ability of the NARX to fit the data better, or conversely, overfit the data due to the small sample size.

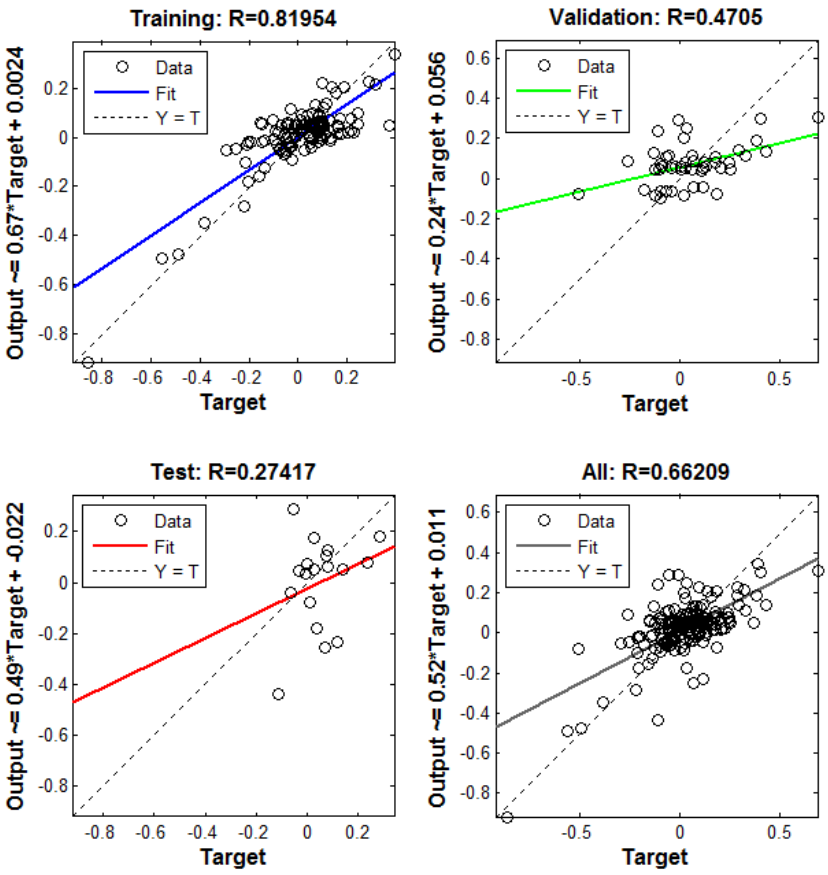


Figure 60 - Comparison of R^2 during training, validation and testing phases

An examination of the autocorrelation coefficients of the error term reveals the suspicion of overfitting above – the error terms at longer lags are non-stationary. This is expected to occur with lower frequency data, implying that more observations are needed, along with possible further differencing.

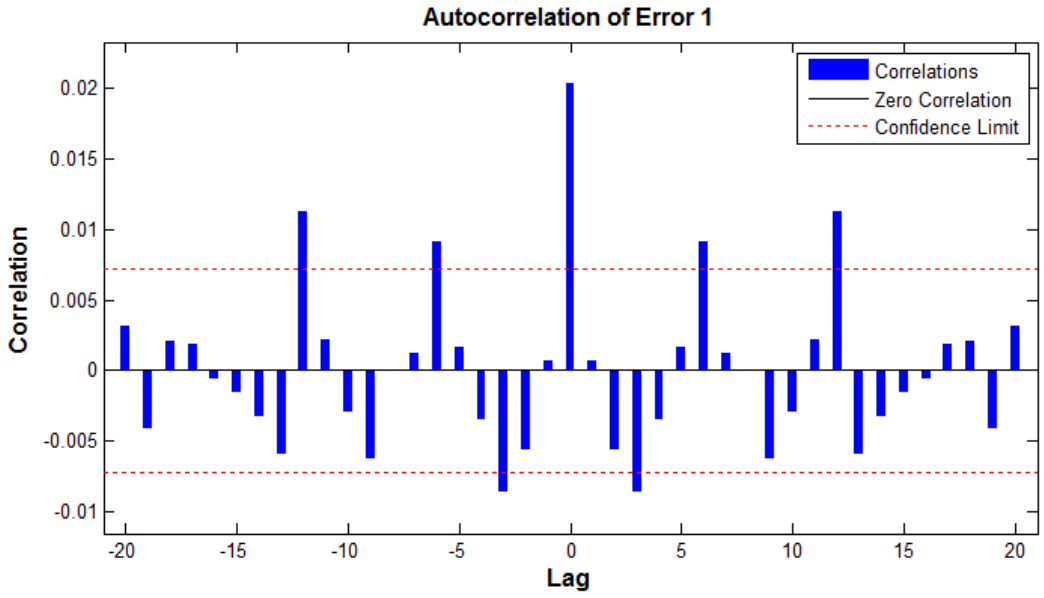


Figure 61 - Autocorrelation function over time

In summary, it was found that a NARX network is suitable to model the returns generating process of all five frequencies of data of the ALSI, however, as the frequency of data lowers, the small number of observations as well as the issue of non-stationarity arises. Of particular interest is that the daily return NARX NN provided the best goodness of fit statistic (as the results of the semi-annual NARX network is rejected due to non-stationarity) to the data. This is in contrast to the results of the random walk hypothesis in that the daily return series was found to be randomly generated. At this juncture, two avenues can be explored as conclusions. The first reasoning points towards the inadequacy of the Hurst exponent (in relation to the NARX NN) in capturing random walk behaviour. The second reasoning points towards frugality on the part of the researcher - if a test result indicated the random walk nature of the daily ALSI return series, then the daily NARX NN should not have been run in the first place. As the number of training cycles was the largest compared to the weekly and monthly frequencies, it is plausible that non-random behaviour can exist yet be complex to detect. However, as the Hurst exponent is a particularly sophisticated measure of randomness in a series, it is plausible that the pattern detected could be spurious. Thus, the overall sample is split into smaller sub-samples and the NARX network run on each sub-sample.

4.7.6 ALSI Sub-sample results

The NARX NN was trained and evaluated using each of the 10 sub-samples to investigate if any deviation in the network output would occur for particular phases of the business cycle as well as to deal with non-stationarity that may be apparent in smaller time intervals. From the 10 sub-samples used, the results of each NARX NN are provided in Table 117 below. For each sub-sample, the optimal number of hidden nodes and delay steps are found. This network is then trained, validated and tested to produce an overall goodness of fit metric. The number of epochs until convergence across all sub-samples remains fairly low, with a spike in sub-sample 5. The number of hidden and delay parameters is relatively stable, with the exception of the delay nodes in sub-sample 6. The overall goodness of fit measures, given by R^2, for each sub-sample is reasonably good and fluctuates over time. This last statement is explored graphically below.

Table 117 – Split sample results of the NARX network

Sub-sample	Time period of sample	Hidden, Delay	Number of epochs	Training R^2 (%)	Validation R^2 (%)	Testing R^2 (%)	Overall R^2 (%)
#1	Sep 1997 to May 1999	1,1	12	52.11	34.56	35.41	46.61
#2	May 1999 to Feb 2001	2,1	11	58.02	34.57	39.45	51.48
#3	Feb 2001 to Oct 2002	2,1	15	38.04	43.25	25.61	38.04
#4	Oct 2002 to Jul 2004	1,1	16	41.47	33.04	39.06	38.74
#5	Jul 2004 to Mar 2006	2,1	12	37.82	34.61	40.01	36.62
#6	Apr 2006 to Dec 2007	1,3	14	46.67	31.58	33.49	41.45
#7	Dec 2007 to Sep 2009	1,1	10	38.34	32.13	45.67	37.85
#8	Sep 2009 to May 2011	2,1	16	41.86	32.56	36.02	38.24
#9	May 2011 to Feb 2013	1,2	13	35.99	38.19	30.88	35.84
#10	Feb 2013 to Oct 2014	1,1	12	34.00	31.39	52.30	35.51

Exploring the plots of hidden nodes and delay parameters in Figure 62, one observes that the number of hidden nodes remains fairly constant throughout the sample period. However, the

number of delay nodes seems to fluctuate at the onset of the global recession during 2006 to 2007. This could possibly indicate that with increased volatility, the current share price is generated by historic share prices. Indeed, the ability of a NARX to capture long term dependencies in the data is superior to that of other network architectures.

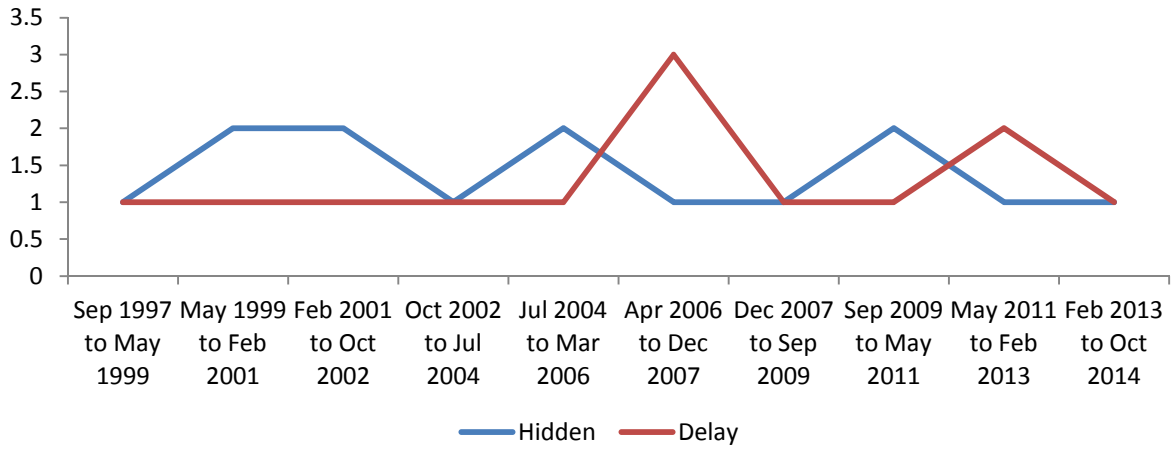


Figure 62 - Comparison of hidden and delay nodes for each sub-sample

Examining the overall goodness of fit R^2 and the testing (out-of-sample forecasting) R^2 in Figure 63, one sees a cyclical pattern that is difficult to correspond to a business cycle. The network appears to perform well at times and poorly at others; but upon closer inspection, performance seems to increase during a recession and decrease during times of prosperity. Indeed, the highest overall fit is during Dec 2007 to Sep 2009, which is during the financial recession. Examining times of prosperity, for example the technology bubble during 2000 to 2001 and the recovery from the recession during 2012 to 2013, one finds that the network performance during these times is below the average overall R^2 of 32%.

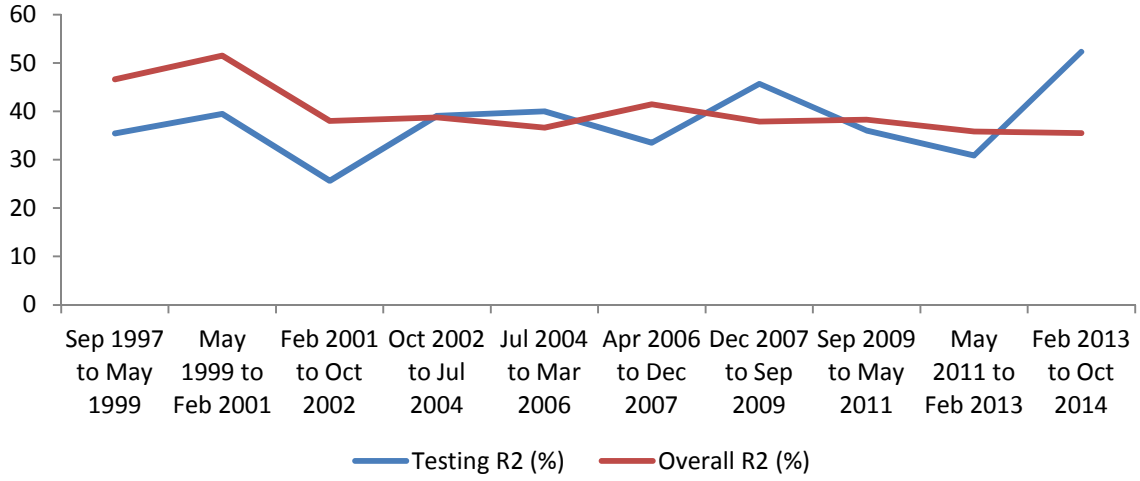


Figure 63 - Comparison of training and overall R^2 for each sub-sample

4.7.7 Summary of results for modelling returns using additional factors

This sub-section set out to model the returns on the ALSI as a function of macroeconomic, microeconomic and behavioural type variables, using an APT framework. The modelling procedure employed artificial intelligence techniques, in particular a NARX neural network.

Recurrent Neural Networks (RNNs) can be used to model time-varying problems, recognise patterns or for forecasting purposes. These networks can model non-linear chaotic, dynamic systems and in principle, should be able to predict future values of the output variable. The results showed that a NARX network was somewhat adequate at explaining returns over the sample period, with the accuracy of the network improving at lower frequency data (the monthly NARX NN performed marginally worse than the daily NARX NN). While the out of sample forecasts were not error prone, the errors were of small values, implying that the use of a neural network can aide in explaining the returns generating process. Further, the results of the network did show that over certain periods in the sample, the errors of the output increased dramatically – the network was not able to consistently explain well over the entire sample period. The NARX NN methodology employed appeared to perform better at fitting the data during recessionary times and poor during times of prosperity. This finding is particularly interesting as it implies that these time varying, "non-specified" models can explain return relationships better when an economy is experiencing a recession and worse when an economy is prospering. This result can be linked to that of Seetharam and Britten (2013) in that during recessionary times *only*, investors on the JSE (or at least those that have an investment position on the ALSI) exhibit herding behaviour and mimic the actions of other investors. Further, as each of the networks constructed had a positive number of hidden nodes, indirect evidence on the non-linearity of share returns was found. Recall that if a problem is non-linearly separable, a multi-layer network is considered best to model such a problem. Last, recall that Basheer (1998) suggests that one hidden layer is sufficient to approximate continuous functions, whereas Masters (1994) suggests two hidden layers for discontinuous functions. As most of the NNs based on ALSI data had more than one hidden layer, this pointed towards that particular sample of data being approximated by a discontinuous function - an indicator of complexity.

Conceptually, one can view the results of the NARX model that fluctuated over time as a form of evidence in favour of the AMH. It has been established that the model is indeed worthy of inspection as the forecasts were quite close to the actual values of the daily returns on the ALSI, however, this was not always the case over the entire sample period. Even with the adaptive capability of neural network to change its parameters over time, prediction errors fluctuated to levels that were acceptable, to levels that were not acceptable. This implies that an adaptive model, while good, is not perfect at explaining the daily returns generating process of the ALSI. It is thus considered evidence in favour of the AMH as the differing goodness of fit statistics over the sub-samples can be seen as the conditions (or rather efficiency) of the market changing over time.

A short digression to the testing of a portfolio strategy is now presented, to assist in a holistic view of testing market efficiency.

4.8 A practical test of market efficiency

The practical question of whether markets are efficient or not can be examined by comparing the returns of a passive (buy and hold) trading strategy against an active trading strategy. To determine the returns of the buy and hold strategy, a monetary amount of R10 000 was invested in each share at the beginning of the sample period, 1 September 1997, and the total return was calculated based on the ending share price as at 31 October 2014. To account for transaction costs, a fixed amount of 1% is levied when the share is bought and 1% when the share is sold on 31 October 2014. Further, the returns are inclusive of any cash dividends (only) paid during the sample period.

Second, a technical analysis trading rule is employed as a proxy for the active trading strategy. Specifically, a 50 day and 200 day moving average crossover rule is used to determine when the investor would buy (or sell) a particular share. These particular moving averages are chosen due to their popularity. This hypothetical investor will either invest money in the stock market or in the corresponding risk free asset (proxied by the 91 day treasury bill rate). Again, dividends accrued while the share is held are taken into account, along with transaction costs as described previously.

Third, a portfolio strategy utilising the output of the neural network is employed, with the logic outlined above. In other words, when the network "forecasts" an upswing on the ALSI, relative to the current ALSI value, then the investor would purchase the ALSI (and vice versa). Transaction costs of 1% per purchase and 1% per sale of the index is also levied, with the investor beginning with an amount of R10 000.

The annualised returns over the entire sample period are shown in Table 118 and Table 119 below. Out of the 44 equities and six indices invested in, there are 18 examples (17 equities and one index) of an active strategy outperforming a passive strategy, net of costs. The neural network trading strategy performs worst of the three (albeit this conclusion is limited to the ALSI). While there are instances where the difference in returns between the two strategies is economically significant, there are also cases where this difference is negligible. Given that the 1% transaction cost per purchase and sale is a proxy for actual transaction costs, it is plausible that following the active trading strategy might produce results which are not in favour of market efficiency. Further, there are a number of shares that produce relatively high returns (around 30%). Last, it is interesting to note that one of the indices (the J150) produced higher returns under the active strategy compared to the passive strategy. One possible explanation for this outperformance would be the volatility in the underlying asset (gold, in this example). While the gold price (in U.S. Dollars) has steadily increased over the sample period, the somewhat recent "gold price bubble", along with corporate social responsibility issues with gold suppliers, would have caused more volatility and thus potentially greater returns if an investor were to time the purchase and sale of this commodity. It is also worth noting that the difference in returns for the indices is somewhat marginal, implying that whether one were to follow a passive or active strategy on an index, the outcome can be considered the same. However, given the higher prevalence of transaction costs in the active strategy, the rational investor would opt for the buy and hold strategy over the active strategy for indices.

As the passive strategy outperforms the active strategy in more than half of equities examined, one can be tempted to conclude that the market (proxied by these 44 equities and six indices) is weak form efficient. However, there are instances where this is not the case.

The philosophical question then arises as to whether these 18 instances are enough to conclude that the market is not weak form efficient. In terms of the implications of the AMH, that of cyclical efficiency and cyclical profitability, one would need to examine the trading rule over different sub-samples. However, given the criticisms of testing the universe of trading rules outlined in Chapter 2, this avenue is practically impossible, even with the assistance of artificial intelligence. In other words, while AI might assist in automating the trading rule process, one cannot test for market efficiency as there are a theoretically infinite amount of rules in existence. Thus the rejection of market efficiency by any one trading rule is not considered absolute proof (and the converse is similarly true).

Table 118 - Comparison of returns across active and passive strategies (1)

	B&H Return	MA Return	MA strategy outperforms B&H strategy	NN Return	B&H End Value	MA End Value
SAB	16.56%	9.54%			R 138 736	R 37 814
BIL	21.28%	19.16%			R 274 250	R 192 584
NPN	22.11%	18.50%			R 308 388	R 174 178
MTN	25.95%	25.12%			R 524 740	R 458 365
SOL	17.92%	16.55%			R 169 478	R 128 553
AGL	10.93%	8.80%			R 59 375	R 32 550
FSR	17.74%	8.09%			R 164 969	R 27 998
SBK	15.97%	13.16%			R 127 149	R 73 475
APN	50.61%	12.70%			R 11 299 244	R 67 868
BGA	15.83%	7.97%			R 124 661	R 27 292
RMH	16.82%	10.39%			R 144 290	R 44 535
MDC	24.32%	21.49%			R 419 878	R 272 787
GRT	13.75%	13.60%			R 91 248	R 79 265
INP	10.79%	8.01%			R 58 056	R 27 553
MPC	31.30%	24.21%			R 1 071 435	R 403 246
IMP	21.17%	9.80%			R 270 199	R 39 809
NTC	22.12%	24.68%	Y		R 308 782	R 430 821
MMI	13.09%	8.98%			R 82 584	R 33 744
ANG	1.59%	3.55%	Y		R 13 105	R 8 196
IPL	11.82%	12.72%	Y		R 68 042	R 68 077
NPK	9.84%	10.71%	Y		R 50 073	R 47 371
GFI	3.10%	7.12%	Y		R 16 888	R 22 554
ASR	30.33%	30.92%	Y		R 943 667	R 1 010 601
INL	10.84%	8.16%			R 58 473	R 28 452
PIK	17.01%	7.95%			R 148 373	R 27 181

Table 119 - Comparison of returns across active and passive strategies (2)

	B&H Return	MA Return	MA strategy outperforms B&H strategy	NN Return	B&H End Value	MA End Value
TFG	17.91%	21.92%	Y		R 169 154	R 290 469
SNT	17.91%	18.42%	Y		R 169 147	R 172 253
HYP	28.61%	19.65%			R 751 183	R 207 544
SAP	1.78%	8.64%	Y		R 13 531	R 31 507
CLS	19.59%	12.78%			R 215 645	R 68 835
GND	29.69%	26.77%			R 867 889	R 576 869
PPC	13.77%	13.47%			R 91 546	R 77 581
AFE	16.26%	18.38%	Y		R 132 720	R 171 214
RCL	12.53%	17.32%	Y		R 75 927	R 145 111
SUI	10.34%	11.19%	Y		R 54 141	R 51 809
ILV	11.49%	9.12%			R 64 657	R 34 733
RLO	15.98%	17.43%	Y		R 127 347	R 147 769
FBR	31.26%	31.79%	Y		R 1 066 004	R 1 133 267
MUR	6.48%	14.62%	Y		R 29 363	R 94 139
SPG	-5.64%	4.44%	Y		R 3 690	R 11 078
FPT	8.62%	8.68%	Y		R 41 327	R 31 761
SAC	16.92%	13.62%			R 146 367	R 79 543
OCE	23.04%	21.06%			R 351 470	R 255 816
WBO	23.84%	21.68%			R 392 800	R 280 370
J150	1.30%	6.11%	Y		R 12 489	R 17 688
J200	12.41%	10.91%			R 74 443	R 49 119
J203	12.59%	11.15%		-1.00%	R 76 544	R 51 415
J211	13.05%	12.21%			R 82 153	R 62 279
J213	11.62%	10.11%			R 66 050	R 42 258
J177	11.26%	10.14%			R 62 468	R 42 510

4.9 Overall summary of results

To examine cyclical efficiency on the South African equities market, a range of tests and modelling was conducted on 50 returns series - 44 equities and six indices - using daily, weekly and monthly frequency data to evaluate three secondary null hypotheses[31]. This

[31] **$H_{0,A}$:** Share price behaviour, in the South African market, does not follow a random walk.
$H_{0,B}$: Share price behaviour, in the South African market, cannot be modelled by an autoregressive function *with no* exogenous inputs.

provided a holistic view of market efficiency at both an individual share level and aggregated index level. It was found that 18 of the 44 shares had returns that were randomly generated under daily and weekly data, but not randomly generated under monthly data. Further, as the frequency of data lowered (from daily to monthly), more shares appeared to not follow a random walk. These two results indicate that the aggregated share market appears to not follow a random walk under monthly data, thus one would fail to reject the null hypothesis labelled $H_{0,A}$. Before examining the consecutive hypotheses, a summary and comparison of the tests run for all equities and indices is provided. Given that the majority of tests examine significance at the 95% level of confidence, it is quite possible that the results could be manifested by chance, 5% of the time. This section discusses the results of each test across frequencies and sectors of the shares analysed, with emphasis on the results that are noteworthy of discussion. In other words, the results of all tests across all frequencies of all 50 securities is not presented here, but rather those results that are "anomalies".

While four shares (BIL, ANG, GFI and GND) along with the J150 and J177 indices have normal distributions under monthly data according to the JB test, nois normally distributed under the KS test. Given the non-parametric nature of the K-S test, as well as the reasoning outlined previously, the results of non-normality across all securities holds.

From Tables 120, 121 and 122 below, fourteen shares (AFE, ASR, SAB, PIK, CLS, MPC, SUI, RMH, INP, INL, SAC, PPC, MUR and SPG) and three indices (J211, J213, J177) had distributions that followed both a linear and non-linear pattern (recall that the distributions could be linear at values below the mean and non-linear at values above the mean); with five shares (AGL, ANG, ILV, BGA and IPL) and one index (the J150) following a strictly linear distribution under the BDS test under monthly data. These results are equivalent to 54% of securities following a strictly non-linear distribution. The financial and consumer services industries show up strongly, which is corroborated by the corresponding index (J213) also not following a linear distribution. In the example of mining, it is interesting to see the mining index (J177) in the results but not many mining shares. This could imply that the trading activity, which is a function of the share's distribution, differs for individual shares and

---

$H_{0,C}$: Share price behaviour, in the South African market, cannot be modelled by an autoregressive function *with* exogenous inputs.

indices. Under quarterly data, six shares (AGL, ILV, IPL, NPK, HYP and GRT) followed a linear distribution, 26 securities followed both a linear and nonlinear distribution. Last, under semi-annual data, five shares (NTC, IPL, GFI, SAP, SUI) and the J150 followed a strictly linear distribution, whereas 36 securities followed both a linear and non-linear distribution. Given that the majority of shares follow a mix of a linear and non-linear distribution, there is little insight gained from industries. In most instances, the distribution remains the same under monthly, quarterly and semi-annual data. There are two cases (GFI and ASR), where the distributions became linear under semi-annual data but were linear and non-linear under quarterly data. From a trading perspective, the non-linear nature of the return distributions points towards some form of mean reversion in share prices, which implies that if one can time the market, it is possible to consistently earn abnormal profits.

Table 120 – BDS test results for all shares (1)

Share	Sector	BDS		
		Monthly	Quarterly	Semi-Annual
AFE	Basic materials	Linear and Non-linear	Linear and Non-linear	
SAP	Basic materials			Linear
BIL	Mining			Linear and Non-linear
AGL	Mining	Linear	Linear	Linear and Non-linear
IMP	Mining		Linear and Non-linear	Linear and Non-linear
ANG	Mining	Linear	Linear and Non-linear	Linear and Non-linear
GFI	Mining		Linear and Non-linear	Linear
ASR	Mining	Linear and Non-linear	Linear and Non-linear	Linear
SAB	Consumer goods	Linear and Non-linear	Linear and Non-linear	Linear and Non-linear
RCL	Consumer goods			Linear and Non-linear
ILV	Consumer goods	Linear	Linear	Linear and Non-linear
OCE	Consumer goods		Linear and Non-linear	Linear and Non-linear
GRT	Consumer goods		Linear	Linear and Non-linear
FBR	Consumer goods		Linear and Non-linear	
PIK	Consumer services	Linear and Non-linear	Linear and Non-linear	Linear and Non-linear

Table 121– BDS test results for all shares (2)

Share	Sector	Frequency		
		Monthly	Quarterly	Semi-Annual
CLS	Consumer services	Linear and Non-linear	Linear and Non-linear	Linear and Non-linear
MPC	Consumer services	Linear and Non-linear	Linear and Non-linear	Linear and Non-linear
TFG	Consumer services		Linear and Non-linear	Linear and Non-linear
NPN	Consumer services			Linear and Non-linear
SUI	Consumer services	Linear and Non-linear		Linear and Non-linear
FSR	Financials		Linear and Non-linear	Linear and Non-linear
SBK	Financials		Linear and Non-linear	Linear and Non-linear
BGA	Financials	Linear	Linear and Non-linear	Linear and Non-linear
RMH	Financials	Linear and Non-linear	Linear and Non-linear	
INP	Financials	Linear and Non-linear	Linear and Non-linear	Linear and Non-linear
INL	Financials	Linear and Non-linear	Linear and Non-linear	Linear and Non-linear
MMI	Financials			Linear and Non-linear
SNT	Financials		Linear and Non-linear	Linear and Non-linear
HYP	Financials		Linear	Linear and Non-linear
FPT	Financials			Linear and Non-linear
SAC	Financials	Linear and Non-linear		Linear and Non-linear
MDC	Healthcare			Linear and Non-linear
NTC	Healthcare			Linear
APN	Healthcare		Linear and Non-linear	Linear
PPC	Industrials	Linear and Non-linear	Linear and Non-linear	Linear and Non-linear
MUR	Industrials	Linear and Non-linear		Linear and Non-linear
WBO	Industrials		Linear and Non-linear	Linear and Non-linear
NPK	Industrials		Linear	Linear and Non-linear
IPL	Industrials	Linear	Linear	Linear
GND	Industrials		Linear and Non-linear	Linear

Table 122– BDS test results for all shares (3)

Share	Sector	Frequency		
		Monthly	Quarterly	Semi-Annual
SPG	Industrials	Linear and Non-linear	linear	Linear and Non-linear
SOL	Oil		Linear and Non-linear	Linear
MTN	Teleco		Linear and Non-linear	Linear
J150	JSE Gold Mining Index	Linear		Linear
J200	JSE Top 40			Linear and Non-linear
J203	JSE All Share Index (ALSI)			Linear and Non-linear
J211	JSE Industrial 25	Linear and Non-linear	Linear	Linear and Non-linear
J213	JSE Financial and Industrial 30	Linear and Non-linear	Linear	Linear and Non-linear
J177	JSE Mining Index	Linear and Non-linear	Linear	Linear and Non-linear

When examining the results of the Runs test in Tables 123 and 124 below, 38 shares and 6 indices were non-randomly generated under daily data. There is no particular industry which stands out, apart from two financial shares (SBK and RMH), which are randomly generated. These shares had runs that were less than expected by pure chance, which infers that there were trends in these shares' returns over the sample period on a daily frequency. Nine shares (ASR, GRT, PIL, SBK, FPT, SAC, WBO, NPK and SOL) were non-randomly generated under weekly data. These shares came from various sectors in the market – mining, consumer goods, consumer services, financials and industrials and their runs are a mix of greater than and less than expected by pure chance. However, those shares in the consumer goods (services) sector as well as the financial sector had runs that were greater than expected by pure chance, implying higher trading activity. When examined in relation to the trading strategy performance, there was little overlap between the share's outperformance of a buy and hold strategy and being non-random. There was a single share, ASR, that had a consistent lower number of runs under these three frequencies. Indeed the trading strategy for ASR outperformed the market, confirming the logic outlined above. ASR also appears to have low liquidity, which points towards the outperformance being generated not from "share effects" as opposed to market microstructure effects. Last, when comparing those shares that had a greater number of runs, they were more likely to have a non-linear distribution under the BDS test – implying that they can be modelled. The shares that did outperform (where the

active strategy outperformed the buy and hold strategy) were ASR, FPT and NPK. As there is no consistent industry which outperformed, an alternative reason is that perhaps the liquidity of these shares caused the outperformance. Six shares (AFE, BIL, ASR, TFG, NTC and SOL) were non-randomly generated under monthly data. The mining industry features strongly again in the monthly data. The runs are predominantly greater than that expected by chance, implying higher volatility, which is generated by higher trading volumes. When viewed alongside the results from the JB test, it was found that mining shares are typically normally distributed. If returns are non-random according to the Runs test, yet are normally distributed, it is plausible that no consistent abnormal profits are made, especially so when the frequency of data is monthly. The results of the Runs test are equivalent to 12%, 82% and 88% of securities being randomly generated from daily, weekly and monthly data respectively. Last, under quarterly data, 11 shares are non-randomly generated and five are non-randomly generated under semi-annual data. The Healthcare and industrials sectors show up in both quarterly and semi annual frequencies.

Table 123 – Runs test results for all shares (1)

Share	Sector	Frequency				
		Daily	Weekly	Monthly	Quarterly	Semi-Annual
AFE	Basic materials	Less		Less	Less	Greater
SAP	Basic materials	Less				
BIL	Mining			Greater		
AGL	Mining	Less				
IMP	Mining	Less			Less	
ANG	Mining	Less				
GFI	Mining	Less				
ASR	Mining	Less	Less	Less		
RCL	Consumer goods	Less				
ILV	Consumer goods	Less				
OCE	Consumer goods	Less				
GRT	Consumer goods	Less	Greater			
FBR	Consumer goods	Less			Less	Less
PIK	Consumer services	Less	Greater	Greater	Greater	
CLS	Consumer services	Less			Less	
MPC	Consumer services	Less				
TFG	Consumer services	Less		Less		
NPN	Consumer services	Less				
SUI	Consumer services	Less				
FSR	Financials	Less				

Table 124 – Runs test results for all shares (2)

Share	Sector	Frequency				
		Daily	Weekly	Monthly	Quarterly	Semi-Annual
SBK	Financials		Greater			
BGA	Financials	Less				
INP	Financials	Less				
INL	Financials	Less				
MMI	Financials	Less				
SNT	Financials	Less		Greater	Less	
HYP	Financials	Less				
FPT	Financials	Less	Greater		Greater	
SAC	Financials	Less	Greater			
MDC	Healthcare	Less				Greater
NTC	Healthcare	Less		Greater		Less
APN	Healthcare	Less			Less	
PPC	Industrials				Greater	
MUR	Industrials	Less				
WBO	Industrials	Less	Less			
RLO	Industrials	Less				
NPK	Industrials		Greater		Greater	
IPL	Industrials	Less				
GND	Industrials	Less			Less	Less
SPG	Industrials	Less				
SOL	Oil	Less	Greater	Greater	Less	
J150	JSE Gold Mining Index	Less				
J200	JSE Top 40	Less				
J203	JSE All Share Index (ALSI)	Less				
J211	JSE Industrial 25	Less				
J213	JSE Financial and Industrial 30	Less				
J177	JSE Mining Index	Less				

From the Chow Denning test results in Tables 125, 126 and 127 below, 43 shares were non-randomly generated using daily data, 20 shares and one index (the J213) were non-randomly generated under weekly data. The most heavily represented sectors of mining, consumer goods and services, financials and industrials appear to have non-randomly generated returns. Most of these results are in line with the Runs test, yet there are some shares which were found to be non-random under the CD test, but random under the Runs test. Given that the CD test examines variances, as opposed to the Runs test which examines "level" data, it is possible that heteroscedasticity is the cause of non-randomness in these shares. Three shares

(PIK, TFG and MDC) were non-randomly generated under monthly data and given the small sample; there is no discernible pattern per industry. This is equivalent to 14%, 56% and 94% being randomly generated from daily, weekly and monthly data respectively. Using quarterly data, eight securities are non-randomly generated (BGA, RMH, MDC, IMP, PIK, SNT, GND and OCE) and three under semi-annual data (FSR, RMH and AFR). The financial sector shows up strongly in both frequencies (indeed under all frequencies).

Table 125 – Chow Denning test results for all shares (1)

Share	Sector	Frequency				
		Daily	Weekly	Monthly	Quarterly	Semi-Annual
AFE	Basic materials	Non-Random				Non-Random
SAP	Basic materials	Non-Random				
BIL	Mining	Non-Random	Non-Random			
AGL	Mining	Non-Random				
IMP	Mining	Non-Random	Non-Random		Non-Random	
ANG	Mining	Non-Random	Non-Random			
GFI	Mining	Non-Random	Non-Random			
ASR	Mining	Non-Random				
SAB	Consumer goods	Non-Random	Non-Random			
RCL	Consumer goods	Non-Random				
ILV	Consumer goods	Non-Random	Non-Random			
OCE	Consumer goods	Non-Random	Non-Random		Non-Random	
GRT	Consumer goods	Non-Random				
FBR	Consumer goods	Non-Random	Non-Random			
PIK	Consumer services			Non-Random	Non-Random	
CLS	Consumer services	Non-Random	Non-Random			
MPC	Consumer services	Non-Random	Non-Random			
TFG	Consumer services	Non-Random		Non-Random		
NPN	Consumer services	Non-Random				

Table 126 – Chow Denning test results for all shares (2)

Share	Sector	Frequency				
		Daily	Weekly	Monthly	Quarterly	Semi-Annual
SUI	Consumer services	Non-Random				
FSR	Financials	Non-Random	Non-Random			Non-Random
SBK	Financials	Non-Random	Non-Random			
BGA	Financials	Non-Random	Non-Random		Non-Random	
RMH	Financials	Non-Random			Non-Random	Non-Random
INP	Financials	Non-Random				
INL	Financials	Non-Random				
MMI	Financials	Non-Random				
SNT	Financials				Non-Random	
HYP	Financials					
FPT	Financials	Non-Random	Non-Random			
SAC	Financials	Non-Random	Non-Random			
MDC	Healthcare	Non-Random		Non-Random	Non-Random	
NTC	Healthcare					
APN	Healthcare	Non-Random				
PPC	Industrials					
MUR	Industrials	Non-Random				
WBO	Industrials	Non-Random				
RLO	Industrials					
NPK	Industrials	Non-Random	Non-Random			
IPL	Industrials	Non-Random				
GND	Industrials	Non-Random	Non-Random		Non-Random	
SPG	Industrials	Non-Random	Non-Random			
SOL	Oil	Non-Random	Non-Random			
MTN	Teleco	Non-Random	Non-Random			
J150	JSE Gold Mining Index	Non-Random				

Table 127 – Chow Denning test results for all shares (3)

Share	Sector	Frequency				
		Daily	Weekly	Monthly	Quarterly	Semi-Annual
J200	JSE Top 40	Non-Random				
J203	JSE All Share Index (ALSI)	Non-Random				
J211	JSE Industrial 25	Non-Random				
J213	JSE Financial and Industrial 30	Non-Random	Non-Random			
J177	JSE Mining Index	Non-Random				

The results of the Wright test in Tables 128, 129 and 130 below show that all securities were non-randomly generated under daily data, 28 shares and four indices were non-randomly generated under weekly data; and 15 shares (SAP, IMP, ASR, SAB, FBR, CLS, TFG, MDC, NTC, MUR, WBO, GND, SPG, SOL and MTN) and two indices (the J200 and J203) were non-randomly generated under monthly data. Consumer goods and services, as well as industrials stand out as being non-randomly generated from the monthly data. In contrast, financials show up as non-randomly generated under weekly data but randomly generated under monthly data. This is roughly in line with the CD test and is equivalent to 0%, 36% and 66% of securities being randomly generated under daily, weekly and monthly data respectively. Using quarterly data, two shares (IMP and GND) are non-randomly generated, whereas under semi-annual data, none of the shares are strictly non-randomly generated (in other words, there is not enough statistical evidence to conclude non-randomness).

Table 128 – Wright test results for all shares (1)

Share	Sector	Frequency			
		Daily	Weekly	Monthly	Quarterly
AFE	Basic materials	Non-Random	Non-Random		
SAP	Basic materials	Non-Random		Non-Random	
BIL	Mining	Non-Random	Non-Random		
AGL	Mining	Non-Random			
IMP	Mining	Non-Random	Non-Random	Non-Random	Non-Random
ANG	Mining	Non-Random	Non-Random		

Table 129 – Wright test results for all shares (2)

Share	Sector	Frequency			
		Daily	Weekly	Monthly	Quarterly
GFI	Mining	Non-Random	Non-Random		
ASR	Mining	Non-Random	Non-Random	Non-Random	
SAB	Consumer goods	Non-Random	Non-Random	Non-Random	
RCL	Consumer goods	Non-Random	Non-Random		
ILV	Consumer goods	Non-Random			
OCE	Consumer goods	Non-Random	Non-Random		
GRT	Consumer goods	Non-Random	Non-Random		
FBR	Consumer goods	Non-Random	Non-Random	Non-Random	
PIK	Consumer services	Non-Random	Non-Random		
CLS	Consumer services	Non-Random	Non-Random	Non-Random	
MPC	Consumer services	Non-Random	Non-Random		
TFG	Consumer services	Non-Random	Non-Random	Non-Random	
NPN	Consumer services	Non-Random			
SUI	Consumer services	Non-Random			
FSR	Financials	Non-Random	Non-Random		
SBK	Financials	Non-Random	Non-Random		
BGA	Financials	Non-Random	Non-Random		
RMH	Financials	Non-Random			
INP	Financials	Non-Random			
INL	Financials	Non-Random	Non-Random		
MMI	Financials	Non-Random			
SNT	Financials	Non-Random			
HYP	Financials	Non-Random			
FPT	Financials	Non-Random	Non-Random		

Table 130 – Wright test results for all shares (3)

Share	Sector	Frequency			
		Daily	Weekly	Monthly	Quarterly
SAC	Financials	Non-Random	Non-Random		
MDC	Healthcare	Non-Random		Non-Random	
NTC	Healthcare	Non-Random	Non-Random	Non-Random	
APN	Healthcare	Non-Random			
PPC	Industrials	Non-Random			
MUR	Industrials	Non-Random		Non-Random	
WBO	Industrials	Non-Random	Non-Random	Non-Random	
RLO	Industrials	Non-Random			
NPK	Industrials	Non-Random	Non-Random		
RLO	Industrials	Non-Random			
NPK	Industrials	Non-Random	Non-Random		
IPL	Industrials	Non-Random	Non-Random		
GND	Industrials	Non-Random	Non-Random	Non-Random	Non-Random
SPG	Industrials	Non-Random		Non-Random	
SOL	Oil	Non-Random	Non-Random	Non-Random	
MTN	Teleco	Non-Random	Non-Random	Non-Random	
J150	JSE Gold Mining Index	Non-Random	Non-Random		
J200	JSE Top 40	Non-Random		Non-Random	
J203	JSE All Share Index (ALSI)	Non-Random	Non-Random	Non-Random	
J211	JSE Industrial 25	Non-Random	Non-Random		
J213	JSE Financial and Industrial 30	Non-Random	Non-Random		
J177	JSE Mining Index	Non-Random			

When examining the variance decomposition test results in Tables 131, 132 and 133 below, 35 shares and six indices were non-randomly generated under daily data, 23 shares were non-randomly generated under weekly data and one share (MUR) was non-randomly generated under monthly data. Consumer goods and services appear to be randomly generated under

both daily and weekly data, along with shares in the healthcare and industrial sectors. Financials are randomly generated when examining Variance Decompositions, but do appear to be randomly generated under more sophisticated tests. This result points towards a possibly complex return generating process in those shares, which could also be coupled with peculiarities in their trading compared to other sectors. For example, with spikes in trading volumes for financial shares, it is possible that multiple variances (the CD test) would be detected as opposed to variances at a particular lag (the variance decomposition test). In summary, the results are equivalent to 18%, 46% and 98% being randomly generated under daily, weekly and monthly data respectively. Last, under quarterly data, two shares (IMP and GND) are non-randomly generated, whereas all shares are randomly generated under semi-annual data.

Table 131 – Variance Decomposition test results for all shares (1)

Share	Sector	Frequency			
		Daily	Weekly	Monthly	Quarterly
AFE	Basic materials	Non-Random	Non-Random		
SAP	Basic materials	Non-Random	Non-Random		
BIL	Mining	Non-Random	Non-Random		
AGL	Mining	Non-Random			
IMP	Mining	Non-Random	Non-Random		Non-Random
ANG	Mining	Non-Random	Non-Random		
GFI	Mining		Non-random		
ASR	Mining	Non-Random			
SAB	Consumer goods	Non-Random	Non-Random		
RCL	Consumer goods	Non-Random			
ILV	Consumer goods	Non-Random	Non-Random		
OCE	Consumer goods		Non-random		
GRT	Consumer goods	Non-Random	Non-Random		
FBR	Consumer goods	Non-Random	Non-Random		

Table 132 – Variance Decomposition test results for all shares (2)

Share	Sector	Frequency			
		Daily	Weekly	Monthly	Quarterly
CLS	Consumer services	Non-Random	Non-Random		
MPC	Consumer services	Non-Random			
TFG	Consumer services	Non-Random			
NPN	Consumer services	Non-Random			
SUI	Consumer services	Non-Random			
FSR	Financials	Non-Random	Non-Random		
SBK	Financials	Non-Random	Non-Random		
BGA	Financials	Non-Random	Non-Random		
RMH	Financials	Non-Random			
INP	Financials	Non-Random			
INL	Financials	Non-Random	Non-Random		
MMI	Financials	Non-Random			
FPT	Financials	Non-Random	Non-Random		
SAC	Financials	Non-Random	Non-Random		
APN	Healthcare	Non-Random	Non-Random		
MUR	Industrials	Non-Random		Non-Random	
WBO	Industrials	Non-Random			
NPK	Industrials	Non-Random			
IPL	Industrials	Non-Random			
GND	Industrials	Non-Random	Non-Random		Non-Random
SPG	Industrials	Non-Random	Non-Random		
SOL	Oil	Non-Random	Non-Random		
MTN	Teleco	Non-Random	Non-Random		
J150	JSE Gold Mining Index	Non-Random			

Table 133– Variance Decomposition test results for all shares (3)

Share	Sector	Frequency			
		Daily	Weekly	Monthly	Quarterly
J200	JSE Top 40	Non-Random			
J203	JSE All Share Index (ALSI)	Non-Random			
J211	JSE Industrial 25	Non-Random			
J213	JSE Financial and Industrial 30	Non-Random			
J177	JSE Mining Index	Non-Random			

Lastly, from the Hurst exponent results in Tables 134 and 135 below, 8 shares (AGL, FBR, SBK, RMH, FPT, SAC, GND and SOL) were non-randomly generated under daily data. This is in line with previous test results where the mining, consumer goods, financials and industrial sectors show strongly. No shares were non-randomly generated under weekly data, which is largely due to the confidence intervals given by the Whittle test estimate; and 22 shares and 4 indices (the J200, J203, J211 and J213) were non-randomly generated under monthly data. This is equivalent to 84%, 100% and 56% being randomly generated under daily, weekly and monthly data respectively. From the trend, it would be presumable to say that the number of shares randomly generated under weekly data should lie between the number of daily and yearly shares. Further, the industrials and financial sector shares show up strongly under monthly data to be non-randomly generated. In general, if one were to examine whether the non-random trend is persistent or anti-persistent (mean reverting), the majority of shares under these three frequencies appear to be mean reverting; with industrial shares under monthly data being persistent in their trend (non-mean reverting). Under quarterly data, 33 shares and 6 indices are non-randomly generated, whereas under semi-annual data, 32 shares and 6 indices are non-randomly generated. In both of these frequencies, the majority of financial shares are non-randomly generated, along with healthcare and industrials. All of the indices show up as non-random. Further, of the 39 shares that are non-randomly generated under quarterly data, 10 are mean reverting. This trend is particularly strong in the financial sector. Similarly, of the 38 shares under semi-annual data, 12 are mean reverting. However, there is no particular trend across industries for this frequency of data. Looking across all five frequencies, all of the shares that are non-randomly generated are mean reverting under daily data, with some also mean reverting

under monthly data – particularly those in the financial sector. As the frequency lowers, more shares are found to be non-randomly generated, and more appear to not be mean reverting. However, there are certain shares that remain mean reverting even under semi annual data (such as AFE, IMP, ASR, GRT, MPC, TFG, MMI, PPC, WBO and GND), yet there is no discernible industry pattern across the frequencies. This points towards share specific effects that affect the results as opposed to market effects.

Table 134 – Hurst exponent results for all shares (1)

Share	Sector	Frequency				
		Daily	Weekly	Monthly	Quarterly	Semi-Annual
AFE	Basic materials				Anti-Persistent	Anti-Persistent
SAP	Basic materials				Persistent	
BIL	Mining				Persistent	Persistent
AGL	Mining	Anti-Persistent			Persistent	Persistent
IMP	Mining					Anti-Persistent
ANG	Mining				Persistent	
GFI	Mining				Persistent	Persistent
ASR	Mining			Persistent		Anti-Persistent
SAB	Consumer goods				Persistent	Persistent
RCL	Consumer goods				Anti-Persistent	Persistent
OCE	Consumer goods			Persistent		
GRT	Consumer goods					Anti-Persistent
FBR	Consumer goods	Anti-Persistent			Persistent	
PIK	Consumer services				Anti-Persistent	Persistent
CLS	Consumer services					Persistent
MPC	Consumer services			Anti-Persistent	Persistent	Anti-Persistent
TFG	Consumer services					Anti-Persistent
NPN	Consumer services				Persistent	
SUI	Consumer services			Anti-Persistent		Persistent
FSR	Financials			Anti-Persistent	Anti-Persistent	
SBK	Financials	Anti-Persistent		Anti-Persistent	Anti-Persistent	

Table 135– Hurst exponent results for all shares (2)

Share	Sector	Frequency				
		Daily	Weekly	Monthly	Quarterly	Semi-Annual
BGA	Financials			Anti-Persistent	Anti-Persistent	
RMH	Financials	Anti-Persistent			Anti-Persistent	
INP	Financials			Anti-Persistent	Anti-Persistent	Persistent
INL	Financials			Anti-Persistent	Anti-Persistent	Persistent
MMI	Financials				Anti-Persistent	Anti-Persistent
SNT	Financials				Persistent	Persistent
HYP	Financials				Persistent	Persistent
FPT	Financials	Anti-Persistent		Persistent	Persistent	Persistent
SAC	Financials	Anti-Persistent		Persistent		
MDC	Healthcare				Persistent	Persistent
NTC	Healthcare				Persistent	Persistent
APN	Healthcare			Persistent	Persistent	Persistent
PPC	Industrials			Persistent	Persistent	Anti-Persistent
MUR	Industrials			Persistent	Persistent	
WBO	Industrials				Persistent	Anti-Persistent
RLO	Industrials			Persistent	Persistent	Persistent
NPK	Industrials				Persistent	Anti-Persistent
IPL	Industrials					Persistent
GND	Industrials	Anti-Persistent			Persistent	Persistent
SPG	Industrials			Anti-Persistent		Anti-Persistent
SOL	Oil	Anti-Persistent		Persistent	Persistent	Persistent
MTN	Teleco			Persistent	Persistent	Persistent
J150	JSE Gold Mining Index				Persistent	Anti-Persistent
J200	JSE Top 40			Anti-Persistent	Persistent	Persistent
J203	JSE All Share Index (ALSI)			Anti-Persistent	Persistent	Persistent
J211	JSE Industrial 25			Anti-Persistent	Persistent	Persistent
J213	JSE Financial and Industrial 30			Anti-Persistent	Persistent	Persistent
J177	JSE Mining Index				Persistent	Persistent

From the summary of test results above, it can be inferred that the results found are not due to "pure chance", as the majority of results point toward a significantly high (low) number of securities that reject the null hypothesis of each test. There is however, one instance where 94% of securities were found to be randomly generated under monthly data from the Chow Denning test. Attention is now drawn to focus on the J203 (ALSI).

In determining whether the three frequencies of return data of the ALSI follow a random walk, various parametric and non-parametric tests were performed. These results, for the overall sample, are provided in Table 136 below. All tests for normality concluded that the daily, weekly and monthly ALSI return series are non-normal. Similarly, the BDS test showed that all three frequencies of data exhibit non-linear behaviour and are stationary according to the ADF and KPSS tests. A simple measure of rolling autocorrelation depicted a cyclical trend in all of the return series, indicating that the ALSI was efficient over periods of time, but also inefficient over other periods of time.

In examining the Random Walk Hypothesis, five tests were conducted. Two of these tests focus on the mean of the series, whereas the remaining three focus on the variance (standard deviation).

Examining the results of the Runs test and Hurst exponent, the daily ALSI series appears to not be randomly generated under the Runs test, whereas it appears to be randomly generated under the Hurst exponent. Similarly, the monthly ALSI series appears to be randomly generated under the Runs test and not randomly generated under the Hurst exponent. These contradicting results can be explained however. Recall that the Runs test examines randomness at a lag order of 1 only, in contrast to the Hurst exponent which examines randomness along a rolling window. Therefore, the Hurst exponent may detect randomness or non-randomness over smaller intervals that may not be apparent over the entire sample period. As such, the results of the Hurst exponent are considered more reliable than that of the Runs test.

Attention is now turned to the results of the variance tests - the Chow Denning test, Wright test and variance decomposition plot. The Chow Denning method of the variance ratio test pointed towards only the weekly and monthly series following a random walk; whereas the Wright modification for the variance ratio test led to all three series following a random walk. As the Chow and Denning modification is seen as a superior method compared to the Wright modification (the former tests for multiple variances compared to the latter), the results of the Chow and Denning modification hold more significance. Further, the variance decomposition plot offers the same conclusion as the Chow Denning test. However, this result conflicts with those of the Hurst exponent.

Summarising the above two paragraphs produces conflicting results when comparing the Hurst exponent test to the Chow and Denning test. In other words, according to the Hurst exponent, the ALSI appears randomly generated using daily and weekly data and not randomly generated using monthly data. In contrast, the Chow Denning test shows that the ALSI appears not randomly generated under daily data and is randomly generated under weekly and monthly data. Reconciling these results is arguably quite straight forward. Recall that the Hurst exponent examines random walk behaviour using a rolling window approach; whereas the variance ratio tests used in this study used a "fixed" window approach. Thus, the Hurst exponent can be considered superior to that of the variance ratio test. In order for the results to be truly comparable, one would need to use a rolling window approach when implementing the variance ratio test.

In fitting the SETAR model, it was found that all three models had a poor fit to the data, indicative of either a different functional form or the inclusion of additional variables. While this enabled one to reject the null hypothesis labelled $H_{0,B}$, the evidence is weak and the rejection (or failure to reject) is a question of semantics. In other words, share price behaviour can be modelled by an autoregressive function with no exogenous inputs; however the resulting model is a poor fit. Further, in creating NAR and NARX networks, it was found that the NAR network performed marginally better than the SETAR counterpart; whereas the NARX network performed better than the SETAR and NAR models. Again, this leads one to reject the null hypothesis labelled $H_{0,C}$, however, the evidence is not definitive.

Table 136 - Summary of results for the random walk hypothesis

		Daily	Weekly	Monthly	Quarterly	Semi-Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Cyclical	Cyclical
	Runs test	Not Random	Random	Random	Random	Random
	Chow Denning test	Not Random	Random	Random	Random	Random
	Wright test	Random	Random	Random	Random	Random
	Variance decomposition plot	Not Random	Random	Random	Random	Random
	Hurst exponent	Random	Random	Random	Not Random	Not Random
Modelling	SETAR MAPE	100%	100%	100%	100%	100%
	NAR R^2	12.28	12.28	12.28		
	NARX R^2	37.34	37.34	37.34	37.19	66.21

Similarly, the results for each sub-sample are summarised below in Table 137 and Table 138. The tests for normality all point towards each sub-sample not being drawn from a normal distribution as well as being stationary. The test for linearity produces mixed results throughout the sample period. There are sub-samples that are linear, sub-samples that are non-linear and sub-samples that are both linear and non-linear. This interesting result shows that while the overall daily sample is non-linear, there are components of the daily data that have both linear and non-linear behaviour. Within samples where there were both evidence of linear and non-linear behaviour, one can arguably divide these samples further. Further, as per Kaboudan (1999), if a series' data generating process is a combination of linear and non-linear or linear, non-linear and stochastic, then the predictability of the series decreases significantly.

As the results of the Chow and Denning and Wright variance ratio tests differ, the logic outlined in the methodology will be employed. Therefore, according to the Chow and Denning modification of the variance ratio test, all of the sub-samples are randomly generated. This shows that a non-randomly generated series can consist of series that are randomly generated.

Reconciling the difference in results between the Runs test, variance ratio plot and Hurst exponent is somewhat more difficult. It was previously mentioned that the Runs test examines randomness at a single lag order. Therefore, the results of Runs test are not considered superior to that of the variance ratio plot and Hurst exponent. The variance ratio plot examines the fraction of known and unknown variance against lag orders. Each of the plots was compiled to a maximum of 10 lags. Thus, any long term memory would theoretically not be captured in these plots. Therefore, according to the Hurst exponent, one of the sub-samples exhibit non-random behaviour, in contrast to the overall daily sample results in which the exponent showed random behaviour. This implies that a series that is randomly generated can consist of sub-series that are both random and non-randomly generated. Further, given that the Hurst exponent employs a sliding window approach, it is considered more sophisticated than the other tests employed in detecting randomness in a series. It is possible, however, that the window used in the Hurst exponent, which differed per frequency, could have affected the results. In other words, if a longer (shorter) sliding window was used, the test result could have differed. The interaction between this element along with the frequency of data used presents an interesting avenue to explore as future research – the impact of data frequency in determining the optimal sliding window.

In dividing the daily sample into sub-samples, it was found that no SETAR model could be fit to any of the sub-sample data. NARX networks were more successful in fitting the data in each sub-sample, performing better than their corresponding NAR network.

Table 137 - Summary of results for each sub-sample for the random walk hypothesis (1)

		S1	S2	S3	S4	S5
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Non-linear	Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling autocorrelation	Cyclical	Cyclical	Cyclical	Cyclical	Cyclical
	Runs test	Not Random	Not Random	Not Random	Not Random	Not Random
	Chow Denning test	Random	Random	Random	Random	Random
	Wright test	Not Random	Not Random	Not Random	Not Random	Not Random
	Variance decomposition plot	Not Random	Not Random	Not Random	Not Random	Not Random
	Hurst exponent	Random	Random	Random	Random	Random
Modelling	SETAR MAPE	100%	100%	100%	100%	100%
	NAR R2	16.32	16.32	16.32	16.32	16.32
	NARX R2	46.61	46.61	46.61	46.61	46.61

Table 138 - Summary of results for each sub-sample for the random walk hypothesis (2)

		S6	S7	S8	S9	S10
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Non-linear	Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling autocorrelation	Cyclical	Cyclical	Cyclical	Cyclical	Cyclical
	Runs test	Random	Random	Random	Random	Random
	Chow Denning test	Random	Random	Random	Random	Random
	Wright test	Random	Random	Random	Random	Random
	Variance decomposition plot	Random	Random	Random	Random	Random
	Hurst exponent	Random	Random	Random	Random	Random
Modelling	SETAR MAPE	100%	100%	100%	100%	100%
	NAR R2	18.85	18.85	18.85	18.85	18.85
	NARX R2	41.45	41.45	41.45	41.45	41.45

In application of the above tests and models, a framework has emerged to test a market (proxied by a market index) for cyclical efficiency. One first determines if the returns series follows a random walk or deterministic process, and then attempts to model the deterministic process using specified and "unspecified" models which include (exclude) exogenous factors. The primary feature of determining cyclical efficiency emerges from the use of non-overlapping sub-samples of the data. Further, an empirical result emerged in that the frequency of data has a significant role in determining whether markets can be considered efficient.

With reference to the primary and secondary hypotheses of this thesis, it was found that: (1) share price behaviour in the South African market, under a daily and weekly frequency does not follow a random walk; whereas under a monthly frequency do follow a random walk; (2) an autoregressive function with no exogenous inputs could model both daily and weekly returns data, but not monthly returns data; however (3) an autoregressive model with exogenous inputs provides a better fit to the daily and weekly data than its counterpart model with no exogenous inputs. The use of a variety of tests provided robustness to the results; and enabled one to both empirically and theoretically determine if market efficiency, as described by the AMH can be considered a reality. By employing various tests, the researcher becomes cognisant of the shortcomings of any one test, allowing the most sophisticated version to be used in further studies of market efficiency. While statistical techniques might aide in the discussion, a simple, practical test of market efficiency was also employed. The returns of a buy and hold strategy were compared to that of a trading rule and a neural network inspired rule. The results show that in all but one example, the buy and hold return outperformed the trading rule, after costs. While this is found to be in favour of market efficiency, the single example of the active strategy outperforming the passive strategy should not be ignored.

When these secondary hypotheses were subjected to smaller sub-samples, it emerged that the ability of the autoregressive model with exogenous inputs was able to perform better in some time periods and worse in others. Therefore, the primary hypothesis whether market efficiency is cyclical or not has some merit, but requires further empirical analysis. At a minimum, a procedure has been outlined for determining whether market efficiency, in South Africa, can be considered cyclical or not.

4.10 A combined view of market efficiency

Thus far, the results of the statistical tests of market efficiency and the practical test of market efficiency have been viewed in isolation. Table 139 and Table 140 view these results in parallel. The table summarises the results of the random walk test (with emphasis on the Hurst exponent test), showing whether the particular equity (index) had returns that followed a random walk under daily (D), weekly (W), monthly (M), quarterly (Q) or semi-annual (S) data. Further, the results of whether a passive strategy outperformed an active strategy are displayed again.

At first glance, one observes that if the active strategy outperformed the buy and hold strategy, the share's returns were found to follow a random walk in the majority of examples. Conversely, there are five examples where the active strategy outperformed the passive strategy, and the shares returns were not random on a daily basis. The former statement might lead to conclusion that either the test for the random walk is not robust, or the trading strategy is flawed. If a share's return does follow a random walk, this does not preclude the possibility of an active trading strategy outperforming a passive one as one needs to definitively test a variety of trading rules to reach a general conclusion. Further, the result of whether a share's returns follow a random walk or not hold over the entire sample period, yet as observed with sub-samples of the ALSI, this conclusion may not always hold true over all sub-samples. Indeed, while these results can be analysed on an individual share level, by viewing this from a market level, multiple investors following the same active trading strategy should theoretically eliminate any profits that can be made, in line with the AMH and its dual implications of cyclical efficiency and cyclical profitability.

Table 139 - Statistical and trading results (1)

Share Code	Share Name	Not Random	Outperform
SAB	South African Breweries	QS	
BIL	BHP Billiton	QS	
NPN	Naspers	Q	
MTN	MTN Group	MQS	
SOL	Sasol	DMQS	
AGL	Anglo American	DQS	
FSR	Firstrand Group	MQ	
SBK	Standard Bank Group	DMQ	

Table 140 - Statistical and trading results (2)

Share Code	Share Name	Not Random	Outperform
APN	Aspen Healthcare	MQS	
BGA	Barclay's Group Africa	MQ	
RMH	RMB Holdings Ltd	DQ	
MDC	Medi-Clinic Corp	QS	
SHF	Steinhoff International Holdings	S	
INP	Investec	MQS	
MPC	Mr Price Group	MQS	
IMP	Impala Platinum	S	
NTC	Network Healthcare	QS	Y
MMI	MMI Holdings	QS	
ANG	Anglogold	Q	Y
IPL	Imperial Holdings	S	Y
NPK	Nampak	QS	Y
GFI	Goldfields	QS	Y
ASR	Assore	MS	Y
INL	Investec Limited	MQS	
PIK	Pik N Pay Stores	QS	
TFG	The Foschini Group	S	Y
SNT	Santam	QS	Y
HYP	Hyprop Investments	QS	
SAP	Sappi	Q	Y
CLS	Clicks Group	S	
GND	Grindrod	DQS	
PPC	Pretoria Port Cement	MQS	
AFE	A E C I Ltd	QS	Y
RCL	RCL Foods	QS	Y
SUI	Sun International	MS	Y
ILV	Illovo Sugar		
RLO	Reunert	MQS	Y
FBR	Famous Brands	DQ	Y
MUR	Murray & Roberts	MQ	Y
SPG	Super Group	MS	Y
FPT	Fountainhead Property	DMQS	Y
SAC	SA Corporate Real estate Fund	DM	
OCE	Oceana Group	M	
WBO	Wilson Bayley Holmes Ovcon	QS	
J150	JSE Gold Mining Index	QS	Y
J200	JSE Top 40	MQS	
J203	JSE All Share Index (ALSI)	MQS	
J211	JSE Industrial 25	MQS	
J213	JSE Financial and Industrial 30	MQS	
J177	JSE Mining Index	QS	

5 Conclusion

While the results from this thesis aid in the debate on market efficiency, it is imperative to note that the viewpoint was not one of proving the EMH to be (in)correct but rather to provide evidence in favour of a more encompassing hypothesis in which it can be falsified. Mackay (1841) provides a history of financial errors in which the power of a group of individuals does not always produce the most efficient or effective outcome. This is perhaps the core differentiator between behavioural and traditional finance, namely that the majority does not always know better than the individual. In contrast, Surowiecki (2005) purports that organised crowds or institutions are more knowledgeable than any single individual. While there is evidence in favour of both arguments, there is no conclusive empirical viewpoint that has yet been settled on by finance academics. As mentioned previously, the EMH as it stands cannot be refuted as it is not falsifiable. Given that no progress can be made in that area, the alternative would be to develop a falsifiable theory which describes financial markets today. The AMH integrates psychology, sociology, behavioural finance and quantitative finance to produce a somewhat workable definition of efficiency. Recall that no formal means of testing cyclical efficiency has been established in the literature. Therefore, this is the first South African study to offer a comprehensive test of market efficiency, from both a statistical and economic perspective, with the results pointing towards a market whose efficiency changes over time. By examining 44 randomly selected equities and 6 indices, over five frequencies, as well as multiple tests of randomness, return generating processes and trading strategies, this thesis supports the cyclical efficiency implication of the AMH.

Chapter 2 outlined the literature on market efficiency, beginning with a qualitative exposition on how the concept of market efficiency emerged in finance academia. Simply, a market is considered efficient if one cannot use any means available to consistently earn abnormal returns, through the prediction of future stock prices. Market efficiency is not a new concept in the literature as the term has been used since the late 19th century. However, the concept became popularised by Fama (1970) in defining the Efficient Market Hypothesis, which stated that no abnormal profits may be made over time as prices reflect both private and public information. From the viewpoint of a market participant, studies have attempted to analyse the speed of adjustment of prices to new information; while others have taken the statistical definition of the EMH (that share prices follow a random walk) and have attempted

to test the hypothesis. However, irrespective of the viewpoint chosen, there is no consensus on whether markets are efficient according to the literature.

A foray to time series methods was thereafter discussed, to provide a foundation for the econometric and artificial intelligence methods used in this thesis. Various time series models, ranging from simple to complex, were presented as an "evolution" of the field to what led to models being developed in the field of computer science. This led to the application of models in that field to solve problems in finance, an application that is novel in the context of testing market efficiency in a South African context.

A discussion of asset pricing followed, where both considerations of investor rationality and the influence of exogenous factors were presented. Coupled with the foundation provided for time series methods, the discussion on asset pricing would then provide a background and motivation for the artificial intelligence models used, along with the inclusion of exogenous factors that could influence stock returns.

Lastly, some of the emerging (and perhaps esoteric) areas of finance research were discussed, providing a well-rounded view of how inter-disciplinary collaboration can provide solutions to long standing questions in finance.

This thesis sought to enhance the definition of cyclical efficiency by providing empirical evidence that examines whether the JSE equities market is efficient as defined by the AMH. In the journey towards cyclical efficiency, the random walk hypothesis was examined. The results of Chapter 4 confirmed that in the time period under investigation, the changes in the daily ALSI returns were random. An interesting result emerged in that by investigation of five frequencies of ALSI returns, the frequency chosen by the researcher has a significant impact on the results. In particular, it was found that lower frequency ALSI returns series did not follow a random walk, indicative of market inefficiency; whereas the daily and weekly ALSI return series did follow a random walk. Thus, the first of the secondary hypotheses (that of share returns following a random walk) can be rejected under lower frequency data,

but not rejected under daily and weekly frequency data, with respect to the ALSI. Extending that analysis to incorporate a sample of 50 securities for robustness, it was found that 86% of the shares and indices exhibited a random walk under daily data, 78% under weekly data, 56% under monthly data, 22% under quarterly data and 24% under semi-annual data. While there is a slight increase between the number of randomly generated securities in quarterly and semi-annual data, the overall trend points towards higher frequency data being randomly generated, and lower frequency data being non-randomly generated. This is in line with the ALSI specific results in that it appears that the JSE can be considered weak form efficient on a daily and weekly basis but not the remaining frequencies. This result highlights that concluding whether markets are efficient or not, according to the EMH, is a function of the data frequency chosen as well as the sample of assets used. Further, it is intuitive that as the frequency of data decreases from daily to monthly, a fewer number of shares exhibit random walk behaviour as the series have less "noise".

Having established that lower frequency returns do not follow a random walk, it implies that there exists some deterministic process that governs them. The first process would be that of an auto-regressive model, where the current return is only a function of past returns. To aide in this objective, a SETAR model was utilised to model the returns on the ALSI. The SETAR model can cater for non-linearities that may be present in the data and is quite appropriate for modelling cyclical behaviour. This model was run on all three frequencies of data, despite the outcome of tests for a random walk, as it is plausible that the random walk result can be decomposed into non-linear components, along with a noise term. The results of Chapter 4 indicated that while a SETAR model is appropriate, there did exist additional factors (perhaps exogenous) that influence the current daily return on the ALSI. Therefore, the second of the secondary hypotheses of whether returns can be modelled by an autoregressive process cannot be rejected under lower frequency data, but rejected under daily and weekly data.

Ultimately, the aim of Chapter 4 was to establish both the additional factors that influence the ALSI return as well as a suitable model for evaluation of historic patterns and possible prediction of future returns. Drawing from the field of computer science, neural networks were used as approximators to test market efficiency. The use of a neural network enables the researcher to simply specify the inputs to the model with no prior knowledge of the form of

the model itself. This is intuitively appealing when one considers the notion that while a multitude of variables may influence the ALSI, there is little guidance on how this influence actually occurs. This model was run on all five frequencies of data, despite the outcome of tests for a random walk, as it is plausible that the random walk result can be decomposed into non-linear components, along with a noise term, as evidenced by the BDS test. The results from Chapter 4 indicate that a NARX neural network shows potential in modelling the returns on the ALSI, with a number of exogenous factors being included in addition to lagged values of the ALSI itself. This implies that returns on the ALSI are influenced by both exogenous and endogenous (lagged) factors. The exogenous factors included oil returns, gold returns, change in ALSI dividend yield, change in ALSI earnings yield, S&P 500 returns, Hang Seng 100 returns, and FTSE 100 returns. Therefore, the final secondary hypothesis of whether returns can be modelled by an autoregressive model with exogenous inputs cannot be rejected under all data frequencies examined. However, this finding cannot be interpreted in isolation to those previously discussed. While an autoregressive model was not suited to monthly data, having found that a model not specified *a priori* was suited to monthly data indicates that the monthly returns generating process is more complex than initially perceived. In other words, a specified model may exist which can explain the lower frequency returns process, but the one used in this thesis was unable to do so. However, the advantage of neural networks as approximators compared to traditional econometric methods is most pronounced when the sample sizes decreased (quarterly and semi-annual data). There was no need to bin the data, as the overall goodness of fit increased marginally as the frequency decreased.

To enhance the statistical results as well as provide a comprehensive viewpoint on market efficiency, one needs to examine the economic significance of trading strategies to determine if one can outperform the market. This thesis adopted such an approach by determining the returns of a passive (buy and hold) strategy and an active (moving average crossover) strategy, net of costs. The results of a simple buy and hold strategy on the ALSI seemed to perform better than an active trading rule or even an active neural network inspired strategy; thus corroborating the results of the random walk test for daily ALSI data. Therefore, it can be said that the use of AI in developing and implementing a trading strategy, at least using South African data, is not warranted according to the results of this thesis.

It is interesting to note that when viewing the results of the random walk tests alongside the results of the trading strategy, an active strategy produced better returns than a passive strategy for eighteen shares, thirteen of which had returns of all three frequencies that followed a random walk. This implies that while the statistical test of market efficiency might provide one particular result, it needs to be viewed in conjunction with a practical test of market efficiency (as well as from the viewpoint of the AMH). Indeed, if one were to rely on the statistical result, then there would be no rational investor that would attempt a trading strategy on those shares. Given that scenario, in an adaptive market, there would exist at least one investor that would wish to attempt a trading strategy under the assumption that there exists an opportunity that none have capitalised upon (in other words, the investor would be the first to exploit this potentially lucrative opportunity). This result links back to the dual implications of the AMH - that of cyclical efficiency and cyclical profitability.

While neural networks are relatively new in the field of finance, their application in this thesis, indicates that they are less favourable to portfolio management problems and more favourable to asset pricing problems. A shortcoming in the implementation of the NARX neural network was that the accuracy of the network, while reasonably low, did experience periods where the network performed better at predicting the return on the ALSI. Empirically, this would prompt the researcher to re-evaluate the model over that time period, leading to perhaps a new specification. Conceptually, this can be linked to the AMH, as it can be concluded that the network traversed through periods of accurate prediction and inaccurate prediction. This cyclical pattern can indeed be seen as the efficiency of the market changing over time, where a variable that was once significant loses its significance over time (or *vice versa*). By dividing the sample period into smaller samples, prediction was improved, but at the cost of a lack of interpretation. In examination of sub-samples of the daily ALSI return series, it was found that most of the sub-samples exhibited long term memory; and that a NARX network was also suitable to model the returns process. However, given the existence of long term memory, it is advisable for future research to implement the process using overlapping samples as well as differing sample sizes. Further, the advice of Kendall (1953) should be heeded, in that when testing hypothesis, one should take caution to the results and model(s) used.

In summary, this thesis has attempted to show that markets go through cycles of efficiency. The contributions of the thesis are as follows. First, the use of artificial intelligence techniques in solving financial problems is a promising area of research. Indeed, a particular neural network model was found to be the best out of other models considered to explain the relationship between return data. The once considered disparate fields of finance and computer science can be merged to an extent, with techniques from one assisting in problem solving in the other. Second, the frequency of observation is significant in determining market efficiency. This result was found only by robust testing of the data under multiple tests. This brings about a renewed interest in examining data frequencies - from the lowest to the highest - but needs to be naturally tempered by pragmatic reasoning as to their use in testing a hypothesis. Third, a comprehensive study was undertaken to examine the concept of market efficiency at both a share and index level, showing that both the definition of a market (that of an index and of the interaction of agents) needs to be carefully considered. Last, a framework has emerged in which one can systematically examine market efficiency, according to the AMH. If one proceeds to simply begin modelling returns without preliminary tests on the data, spurious relationships can and will be found. The results point to a four stage framework.

1) Determine whether a return series is random or deterministic. If random walk behaviour is present, then the market is weak form efficient.

2) If the result of (1) points towards a deterministic process, determine if prices can be predicted by their lagged values only. If this is the case, then all public information is incorporated into prices – the market is semi-strong form efficient.

3) If the result of (2) shows significant intercept terms, then incorporate additional risk factors in the spirit of an APT framework. If prices can be predicted by both lagged values and exogenous factors, then there exists private information that is not incorporated into the share price.

4) Examine the results of (1), (2) and (3) across differing frequencies of data as well as non-overlapping samples. If the prediction of (3) does not vary over time, then the arrival and assimilation of new or private information is not instantaneous. If however, the prediction of (3) does vary over time, then one can argue that market efficiency is cyclical.

By comprehensively examining the behaviour of equity prices in South Africa, there are stylised facts that emerged. First, studies which examine the efficiency of markets are dependent on the frequency of data. If one were to only use a single frequency of data, one might obtain conflicting conclusions. Second, by binning data into smaller sub-samples (for example, splitting daily data into yearly sub-samples), one can obtain an interesting pattern of whether the equity market is efficient or not. Here, it is often the case that the sum of the parts is not equal to the whole – in other words, one might get a conclusion of, say, randomess, over the entire sample period of daily data, but there may be pockets of non-randomness with the daily data. Third, by running a battery of tests, one provides robustness to the results. This is a somewhat debateable issue as one could either run a variety of tests (each being an improvement over the other) or argue the theoretical merits of each test befoe selecting the more appropriate one. Fourth, analysis according to industries also adds to the result of efficiency, if markets have high concentration sectors (such as the JSE). One might be tempted to conclude that the entire JSE exhibits, say, randomness, where it could be driven by the resources sector as opposed to any other sector. Last, the use of neural networks as approximators is of benefit when examining data with less than ideal sample sizes. The NNs used for quarterly and semi-annual data did not suffer from overfitting in comparison to the more traditional econometric models.

As with all studies, one must be cognisant of generalising a result that held over a particular sample period to that over any sample period. As a natural extension of future research, one can apply the framework in this thesis to different sample periods, different countries and most importantly, to more frequencies of data - one might call it striving to be "comprehensive". Further, as returns for lower frequencies are calculated by using the beginning and end of that particular time period (such as the first and last day of the month), it would be interesting to employ the method in this thesis across an average of daily return data points, such that a weekly or monthly return represents the average daily return. This might present volatility not inherent during the first and last day of the observed prices. Similarly, the use of a particular test will always have proponents and opponents in the field. An attempt was made to circumvent this issue by employing, as far as possible, parametric, non-parametric and graphical versions of tests to ensure that the results obtained are consistent. There is slight favour towards non-parametric tests, as they do not rely on the underlying return distribution to be specified. Given the non-normality of returns found in

these five frequencies, this supports the notion of using non-parametric tests in studies of market efficiency. If one were to use a parametric test, one needs to first establish the distribution of that particular security's returns. When applied across multiple frequencies and multiple securities, this becomes cumbersome, along with the comparability of results being compromised. However, within each of the three categories, numerous tests do exist and it is quite possible that a superior test can be employed. Last, while a NARX network does capture the relationship between the ALSI return series better than other models considered, this relationship cannot be described as the network is considered a "black box". An avenue of future research would be to design and implement a means of describing such relationships, either within the realm of artificial intelligence, or by the use of agent-based models where agents have heterogeneous preferences. The use of Deterministic Finite State Automata appears quite promising in this regard and is left for future research. This recent addition to the field of artificial intelligence enables the researcher to map the pathways of the neural network symbolically, allowing the researcher to observe how the input(s) are transformed into an output. However, specifics such as the weighting of each input are considerably more difficult to obtain. The use of Automatic Encoders can assist in the variable selection problem, as they would select the best inputs for the neural network as opposed to being selected by the researcher. Another technique from the realm of Computer Science, would be the application of a Kalman filter in the neural network. The Kalman filter is used in the case of a multi-scaled data distribution (effectively a return series that is represented daily, weekly and monthly). Kalman filters can be promising in the field of asset pricing as they work well with "noise" in the data as well as being conceptually linked to Bayesian statistics.

While the early works in finance have set the groundwork for many financial fields today, subsequent works primarily focus on finding empirical evidence rather than conceptual or philosophical avenues. Practically, there does exist an abundance of data with many untold discoveries, but this does little to advance the field of finance in new directions. Indeed, the literature on which this thesis is founded presents an opportunity to warrant an investigation into the philosophy of finance. Further, after comprehensively examining market efficiency at both an individual share and aggregated market level, one can offer insights into how the South African market, and possibly other emerging markets, behave.

References

Alagidede, P., & Panagiotidis, T. (2009). Modelling stock returns in Africa's emerging equity markets. *International Review of Financial Analysis, 18*(1), 1-11.

Alexander, S. (1961). Price movements in speculative markets: Trends or random walks. *Industrial Management Review, 2*(2), 7-26.

Amenc, N., & Le Sourd, Y. (2003). *Portfolio Theory and performance analysis.* Chichester: John Wiley & Sons Ltd.

Amigo, J., Kocarev, L., & Szczepanski, J. (2006). Order patterns and chaos. *Physica Letters A, 355*, 27-31.

Angus, J. (1991). Criteria for choosing the best neural network: Part 1. *Naval Health Research Centre*, 1-26.

Antoniou, A., Garrett, I., & Priestly, R. (1998). Macroeconomic variables as common pervasive risk factors and the empirical content of the arbitrage pricing theory. *Journal of Empirical Finance, 5*(3), 221-240.

Anufriev, M., Hommes, C., & Philipse, R. (2013). Evolutionary selection of expectations in positive and negative feedback markets. *Journal of Evolutionary Economics, 23*(3), 663-688.

Arditti, F. (1967). Risk and the required return on equity. *Journal of Finance, 22*(1), 19-36.

Arrow, K. (1950). A difficulty in the concept of social welfare. *Journal of Political Economy*, 328-346.

Balkin, S., & Ord, J. (2000). Automatic neural network modeling for univariate time series. *International Journal of Forecasting, 16*, 509-515.

Barber, B., & Odean, T. (2001). Boys will be boys: Gender, overconfidence and common stock investment. *The Quarterly Journal of Economics, 116*(1), 261-292.

Barr, G. (1990). Macroeconomic identification of the pricing factors on the Johannesburg Stock Exchange. *South African Journal of Business Management, 21*(1), 17-26.

Basheer, I. (1998). Neuromechanistic-based modeling and simulation of constitutive behavior of fine-grained soils. *PhD Thesis, Kansas State University*.

Basheer, I., & Hajmeer, M. (2000). Artificial neural networks: fundamentals, computing, design, and application. *Journal of Microbiological Methods, 43*, 3-31.

Bernatzi, S., & Thaler, R. (1995). Myopic loss aversion and the equity premium puzzle. *The Quarterly Journal of Economics, 110*(1), 73-92.

Berry, M., Burmeister, E., & McElroy, M. (1988). Sorting out risks using known APT factors. *Financial Analysts Journal, 44*(2), 29-42.

Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics, 3*(1), 167-179.

Black, F. (1986). Noise. *Journal of Finance, 41*(3), 529-543.

Bodington, L. (2014). Gold in the South African market: A safe haven or hedge? *Unpublished Masters Research Report, University of the Witwatersrand, Johannesburg*.

Bonga-Bonga, L., & Makakabule, M. (2010). Modelling stock returns on the South African Stock Exchange: A nonlinear approach. *European Journal of Economics, Finance and Administrative Sciences, 19*, 168-177.

Box, G., & Jenkins, G. (1970). *Time series analysis: Forecasting and control.* San Francisco: Holden Day.

Brock, W., & Hommes, C. (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control, 22*(8), 1235-1274.

Brock, W., Dechert, W., & Scheinkman, J. (1987). A test for independence based on the correlation dimension. *Working Paper*.

Brockett, P., Cooper, W., Golden, L., & Pitaktong, U. (1994). A neural network method for obtaining an early warning of insurer insolvency. *Journal of Risk and Insurance, 61*(3), 402-424.

Brown, R. (1963). *Smoothing, forecasting and prediction of discrete time series.* Englewood Cliffs, New Jersey: Prentice-Hall.

Brown, S., & Warner, J. (1985). Using daily stock returns: The case of event studies. *Journal of Financial Economics, 14*(1), 3-31.

Butler, M., & Kazakov, D. (2012). Testing Implications of the Adaptive Market Hypothesis via Computational Intelligence. *IEEE Computational Intelligence Society*, (pp. 1-8). New Jersey.

Cajueiro, D., & Tabak, B. (2004). Ranking efficiency for emerging markets. *Chaos, Solitons and Fractals, 22*(2), 349-352.

Campbell, J., Lo, A., MacKinlay, A., & Whitelaw, R. (1998). The econometrics of financial markets. *Macroeconomic Dynamics, 2*(4), 559-562.

Casdagli, M. (1991). Nonlinear forecasting, chaos and statistics. In L. Lam, & V. Naroditsky, *Modeling complex phenomena.* New York: Springer-Verlag.

Chan, H., & Faff, R. (1998). The sensitivity of Australian industry equity returns to a gold price factor. *Accounting and Finance, 38*(2), 223-244.

Chan, K., Chen, N., & Hsieh, D. (1985). An exploratory investigation of the firm size effect. *Journal of Financial Economics, 14*(3), 451-471.

Chen, F. (1990). Back-propagation neural networks for nonlinear self-tuning adaptive control. *Control Systems Magazine, IEEE, 10*(3), 44-48.

Chen, N. (1983). Some empirical tests of the theory of arbitrage pricing. *Journal of Finance, 38*(5), 1393-1414.

Chen, N., Roll, R., & Ross, S. A. (1986). Economic forces and the stock market. *Journal of Business, 59*(3), 383-403.

Chen, S. (2008). Financial Applications: Stock Markets. In *Wiley Encyclopedia of Computer Science and Engineering.* New York: John Wiley and Sons.

Chen, S., Chang, C., & Du, Y. (2012). Agent-based economic models and econometrics. *The Knowledge Engineering Review, 27*(2), 187-219.

Cheng, A. (1996). Economic factors and stock markets: Empirical evidence from the UK and the US. *International Journal of Finance and Economics, 1*(4), 287-302.

Chow, K., & Denning, K. (1993). A simple multiple variance ratio test. *Journal of Econometrics, 58*, 385-401.

Clare, A., & Priestley, R. (1998). Risk factors in the Malaysian stock market. *Pacific-Basin Finance Journal, 6*(1), 103-114.

Clare, A., & Thomas, S. (1994). Macroeconomic factors, the APT and the UK stock market. *Journal of Business Finance and Accounting, 21*, 309-330.

Clements, M., Franses, P., & Swanson, N. (2004). Forecasting economic and financial time series with non-linear models. *International Journal of Forecasting, 20*, 169-183.

Conlisk, J. (1996). Why bounded rationality. *Journal of Economic Literature, 34*(2), 669-700.

Connor, G., & Korajczyk, R. (1988). Risk and return in an equilibrium APT: Application of a new test methodology. *Journal of Financial Economics, 21*(2), 255-289.

Conrad, J., & Kaul, G. (1998). Anatomy of trading strategies. *Review of Financial Studies, 11*, 489-519.

Constantinides, G. (2002). Rational asset prices. *The Journal of Finance, 57*(4), 1567-1594.

Cont, R. (2001). Empirical Properties of Asset Returns: Stylised Facts and Statistical Issues. *Quantitative Finance, 1*(1), 223-236.

Cowles III, A. (1933). Can stock market forecasters forecast? *Econometrica, 1*(3), 309-324.

Cowles III, A., & Jones, H. (1937). Some a posteriori probabilities in stock market action. *Econometrica, 5*(3), 280-294.

Cutler, D., Poterba, J., & Summers, L. (1989). What moves stock prices? *Journal of Portfolio Management, 15*(3), 4-12.

Daniel, K., & Titman, S. (1999). Market efficiency in an irrational world. *Financial Analysts Journal, 55*, 28-40.

Daniel, K., Hirshleifer, D., & Subrahmanyam, A. (1998). A theory of overconfidence, self attribution and security market under- and over- reactions. *The Journal of Finance, 53*, 1839-1885.

Davidson, S., Faff, R., & Hillier, D. (2003). Gold factor exposures in international asset pricing. *Journal of International Financial Markets, Institutions and Money, 13*(3), 271-289.

De Bondt, W., & Thaler, R. (1985). Does the stock market overreact? *Journal of Finance, 40*(3), 793-805.

De Gooijer, J. G., & Kumar, V. (1992). Some recent developments in non-linear time series modelling, testing and forecasting. *International Journal of Forecasting*, 135-156.

De Gooijer, J., & Hyndman, R. (2006). 25 years of time series forecasting. *International Journal of Forecasting, 22*(3), 443-473.

Demuth, H., & Beale, M. (1998). *Neural Network Toolbox: For use with MATLAB*. Natick: The Math Works Inc.

Desai, V., & Bharati, R. (1998). The Efficacy of Neural Networks in Predicting Returns on Stock and Bond Indices. *Decision Sciences, 29*(2), 405-423.

Di Narzo, A., Aznarte, J., & Stigler, M. (2009). *tsDyn: Nonlinear time series models with regime switching.* CRAN - R Archives.

Dickinson, J., & Muragu, K. (1994). Market efficiency in developing countries: a case study of the Nairobi Stock Exchange. *Journal of Business Finance and Accounting, 21*(1), 133-150.

Ding, Z., Granger, C. W., & Engle, R. F. (1993). A long memory property of stock market returns and a new model. *Journal of Empirical Finance, 1*(1), 83-106.

Durbin, J., & Watson, G. (1950). Testing for serial correlation in least squares regression I. *Biometrika, 37*(3-4), 409-428.

Elman, J. (1991). Distributed representations, simple recurrent networks and grammatical structure. *Machine Learning, 7*(2), 195-225.

Elster, J. (1998). Emotions and economic theory. *Journal of Economic Literature*, 47-74.

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

Engle, R. (2001). GARCH 101: The use of ARCH/GARCH models in applied econometrics. *The Journal of Economic Perspectives, 15*(4), 157-168.

Engle, R., & Ng, V. (1993). Measuring and testing the impact of news on volatility. *The Journal of Finance, 48*, 1749-1778.

Engle, R., & Patton, A. (2001). What good is a volatility model? *Quantitative Finance, 1*(2), 237-245.

Enke, D., & Thawornwong, S. (2004). The adaptive selection of!nancial and economic variables for use with arti!cial neural networks. *Neurocomputing, 56*, 205-232.

Fama, E. (1963). Mandelbrot and the stable paretian hypothesis. *Journal of Business, 36*(4), 420-429.

Fama, E. (1965b). The Behaviour of Stock-Market Prices. *Journal of Business, 38*(1), 34-105.

Fama, E. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance, 25*(2), 383-417.

Fama, E. (1976). *Foundations of finance: portfolio decisions and securities prices.* New York: Basic Books.

Fama, E. (1981). Stock returns, real activity, inflation and money. *American Economic Review, 71*(4), 545-565.

Fama, E. (1990). Stock returns, expected returns, and real activity. *Journal of Finance, 45*(4), 1089-1108.

Fama, E. (1998). Market efficiency, long term returns and behavioral finance. *Journal of Financial Economics, 49*(3), 283-306.

Fama, E., & French, K. (1989). Business conditions and expected returns on stocks and bonds. *Journal of Financial Economics, 25*(1), 23-49.

Fama, E., & French, K. (1992). The cross-section of expected stock returns. *Journal of Finance, 47*(2), 427-465.

Farmer, J., & Lo, A. (1999). Frontiers of finance: Evolution and efficient markets. *Proceedings of the National Academy of Sciences, 96*(18), 9991-9992.

French, K., Schwert, G., & Stambaugh, R. (1987). Expected stock returns and volatility. *Journal of Financial Economics, 19*, 3-29.

Friedman, B. (1979). Optimal expectations and the extreme information assumptions of "rational expectations" macromodels. *Journal of Monetary Economics, 5*, 23-41.

Gabaix, X., & Laibson, D. (2000). A boundedly rational decision algorithm. *The American Economic Review, 90*(2), 433-438.

Garth, A., Rollins, D., Zhu, J., & Chen, V. (1996). Evaluation of model discrimination techniques in artificial neural networks with application to grain drying. *Artificial Neural Networks In Engineering, 6*, 939-950.

Ghiassi, M., Saidane, H., & Zimbra, D. (2005). A dynamic artificial neural network model for forecasting series events. *International Journal of Forecasting, 21*, 341-362.

Giannopoulos, K. (2000). Measuring volatility. In *The Professional's Handbook of Financial Risk Management.* Oxford: Butterworth-Heinemann.

Gibbons, M. (1982). Multivariate tests of financial models: A new approach. *Journal of Financial Economics, 10*(1), 3-27.

Gibbons, M., & Hess, P. (1981). Day of the week effects and asset returns. *Journal of Business*, 579-596.

Gibson, G. (1889). *The Stock Markets of London, Paris and New York.* New York: G.P Putnam's Sons.

Giles, C., Lawrence, S., & Tsoi, A. (2001). Noisy time series prediction using recurrent neural networks and grammatical inference. *Machine Learning, 44*(1-2), 161-183.

Granger, C. (1991). Developments in nonlinear analysis of economic series. *Scandinavian Journal of Economics, 93*(2), 263-276.

Granger, C., & Morgenstern, O. (1963). Spectral analysis of New York stock market prices. *Kyklos, 16*(1), 1-27.

Granger, C., & Newbold, P. (1986). *Forecasting economic time series.* New York: Academic Press.

Granger, C., & Terasvirta, T. (1993). Modelling Nonlinear Economic Relationships. *Oxford University Press*.

Griffin, J., & Stulz, R. (2001). International competition and exchange rate shocks: A cross-country industry analysis of stock returns. *Review of Financial Studies, 14*(1), 215-241.

Groenewold, N. (1997). Share market efficiency: tests using daily data for Australia and New Zealand. *Applied Financial Economics, 7*(6), 645-657.

Grossman, S., & Stiglitz, J. (1980). On the impossibility of informationally efficient markets. *American Economic Review, 70*(3), 393-408.

Hamao, Y. (1988). An empirical examination of the Arbitrage Pricing Theory. *Japan and the World Economy, 1*(1), 45-61.

Hamilton, J. (1989). A New Approach to Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica, 57*, 357-384.

Hamilton, J. (1996). This is what happened to the oil price-macroeconomy relationship. *Journal of Monetary Economics, 38*(2), 215-220.

Hanson, S. (1995). Backpropagation: Some comments and variations. In D. Rumelhart, & C. Yves, *Backpropagation: Theory, architecture and applications* (pp. 237-271). New Jersey: Lawrence Erlbaum.

Hardin, R. (2002). Indeterminancy and Basic Rationality. In V. Jose, *Indeterminancy: The Mapped, The Navigable and The Uncharted.*

Harrison, J., & Kreps, D. (1978). Speculative investor behavior in a stock market with heterogeneous expectations. *The Quarterly Journal of Economics, 92*(2), 323-336.

Hassoun, M. (1995). *Fundamentals of Artificial Neural Networks.* Cambridge: MIT Press.

Haugen, R., Talmor, E., & Torous, W. (1991). The effect of volatilty changes on the level of stock prices and subsequent expected returns. *The Journal of Finance, 46*(3), 985-1007.

Haykin, S. (1994). *Neural Networks: A Comprehensive Foundation.* New York: Macmillan.

Hebb, D. (1949). *The Organisation of Behaviour.* New York: Wiley.

Hecht-Nielsen, R. (1988). Applications of counterpropagation networks. *Neural Networks, 1*, 131-139.

Hill, T., Marquez, L., O'Connor, M., & Remus, W. (1994). Artificial neural network models for forecasting and decision making. *International Journal of Forecasting, 10*(1), 5-15.

Holt, C. (1958). Forecasting seasonals and trends by exponentially weighted averages. *International Journal of Forecasting, 20*, 5-13.

Huang, R., Masulis, R., & Stoll, H. (1996). Energy shocks and financial markets. *Journal of Futures Markets, 16*(1), 1-27.

Hunt, R., & Ellis, H. (1999). *Fundamental of Cognitive Psychology.* McGraw Hill College 6th ed.

Hurst, H. (1951). Long-term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineering, 116*(1), 770-808.

Hutchinson, J., Lo, A., & Poggio, T. (1994). A nonparametric approach to pricing and hedging derivative securities via learning networks. *The Journal of Finance, 49*(3), 851-889.

Hyndman, R., & Koehler, A. (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting, 22*(4), 679-688.

Hyndman, R., Koehler, A., Ord, J., & Snyder, R. (2005). Prediction intervals for exponential smoothing state space models. *Journal of Forecasting, 24*, 17-37.

Hyndman, R., Koehler, A., Snyder, R., & Grose, S. (2002). A state space framework for automatic forecasting using exponential smoothing methods. *International Journal of Forecasting, 18*, 439-454.

Jacobsen, B., & Dannenburg, D. (2003). Volatility clustering in monthly stock returns. *Journal of Empirical Finance, 10*(4), 479-503.

Jadid, M., & Fairbairn, D. (1996). Predicting moment-curvature parameters from experimental data. *Engineering Applications of Artificial Intelligence, 9*(3), 309-319.

Jegadeesh, N., & Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance, 48*(1), 65-91.

Jensen, M. (1978). Some anomalous evidence regarding market efficiency. *Journal of Financial Economics, 6*(2), 95-101.

Johnson, N., Jefferies, P., & Hui, P. (2003). *Financial Market Complexity.* New York: Oxford University Press.

Jung, J., & Shiller, R. (2005). Samuelson's dictum and the stock market. *Economic Inquiry, 43*, 221-228.

Kaboudan, M. (1999). A mesaure of time series predictability using genetic programming applied to stock returns. *Journal of Forecasting, 18*, 345-357.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 263-291.

Kailath, T. (1980). *Linear Systems.* New Jersey: Prentice-Hall.

Kampouridis, M., Chen, S., & Tsang, E. (2012). Market fraction hypothesis: A proposed test. *International Review of Financial Analysis, 23*, 41-54.

Kanas, A. (2001). Neural network linear forecasts for stock returns. *International Journal of Finance and Economics, 6*(3), 245-254.

Kanzler, L. (1999). Very fast and correctly sized estimation of the BDS statistic. *Unpublished*.

Kendall, M. (1953). The analysis of economic time series - Part I: Prices. *Journal of the Royal Statistical Society, Series A (General), 116*(1), 11-34.

Kerber, R. (1992). ChiMerge: Discretization of numeric attributes. In *AAAI-92 Proceedings of the 9th National Conference on AI* (pp. 123-128). Cambridge: MIT Press.

Keynes, J. (1923). Some aspects of commodity markets. *Manchester Guardian Commercial: European Reconstruction Series*, 784-786.

Klimasauskas, K. (1993). *Neural Computing: A Technology Handbook.* Pittsburgh: NeuralWare.

Kolmogorov, A. (1941). Stationary sequences in Hilbert space. *Bulletin of Mathematics, University of Moscow, 2*(6), 1-40.

Krippendorf, K. (2002). Four (In)Determinabilities - not one. In V. Jose, *Indeterminancy: The Mapped, The Navigable and The Uncharted.*

Kuan, C., & Liu, T. (1995). Forecasting exchange rates using feedforward and recurrent neural networks. *Journal of Applied Econometrics, 10*(4), 347-364.

Kuhnen, C., & Knutson, B. (2005). The neural basis for financial risk taking. *Neuron, 47*, 763-770.

Kuo, C., & Reitsch, A. (1995). Neural networks vs conventional methods of forecasting. *The Journal of Business Forecasting*, 17-22.

Lakonishok, J., Shleifer, A., & Vishny, R. (1994). Contrarian investment, extrapolation and risk. *The Journal of Finance, 44*, 1541-1578.

LeRoy, S. (1989). Efficient capital markets and martingales. *Journal of Economic Literature, 27*(4), 1583-1621.

Lewellen, J., & Shanken, J. (2002). Learning, asset pricing tests and market efficiency. *The Journal of Finance, 57*(3), 1113-1145.

Li, E. (1994). Artificial neural networks and their business applications. *Information Management*.

Lim, K. (2007). Ranking market efficiency for stock markets: A nonlinear perspective. *Physica A: Statistical Mechanics and its Applications, 376*, 445-454.

Lin, T., Horne, B., Tino, P., & Giles, S. (1996). Learning long-term dependencies in NARX recurrent neural networks. *IEEE Transactions on Neural Networks, 7*(6), 1329-1338.

Lintner, J. (1964). Optimal dividends and corporate growth under uncertainty. *The Quarterly Journal of Economics, 78*(1), 49-95.

Litinova, J., & Ou-Yang, H. (2003). Endogenous information acquisition: a revisit of the Grossman-Stiglitz model. *Unpublished*.

Lo, A. (2004). The Adaptive Markets Hypothesis: Market Efficiency from an Evolutionary Perspective. *Journal of Portfolio Management, 30*(5), 15-29.

Lo, A. (2005). Reconciling Efficient Markets with Behavioral Finance: The Adaptive Markets Hypothesis. *Journal of Investment Consulting, 7*(2), 21-44.

Lo, A., & MacKinlay, A. (1988). Stock market prices do not follow random walks: Evidence from a simple specification test. *Review of Financial Studies, 1*(1), 41-66.

Lo, A., & Repin, D. (2002). The psychophysiology of real-time financial risk processing. *Journal of Cognitive Neuroscience, 14*, 323-339.

Loewenstein, G. (2000). Emotions in economic theory and economic behavior. *The American Economic Review, 90*(2), 426-432.

Looney, C. (1996). Advances in feedforward neural networks demystifying knowledge acquiring black boxes. *IEEE Transactions on Knowledge and Data Engineering, 8*(2), 211-226.

Maasoumi, E., Khotanzed, A., & Abaye, A. (1994). Artificial neural networks for some macroeconomic series: a first report. *Econometric Reviews, 13*(1), 105-122.

Mackay, C. (1841). *Extraordinary popular delusions and the madness of crowds.* New York: Wiley.

Mandelbrot, B. (1963). The Variation of Certain Stock Prices. *Journal of Business, 36*(4), 394-419.

Mandelbrot, B. (1966). Forecasts of future prices, unbiased markets, and "martingale" Models. *Journal of Business, 39*(S1), 242-255.

Mandelbrot, B. (1971). When can price be arbitraged efficiently? A limit to the validity of the random walk and martingale models. *The Review of Economics and Statistics, 53*(3), 225-236.

Mandelbrot, B. (1977). *Fractals: form, chance and dimension.* San Francisco: WH Freeman.

Masters, T. (1994). *Practical Neural Network Recipes in C++.* Boston: Academic Press.

McCulloch, W., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics, 5*(4), 115-133.

McElroy, M., Burmeister, E., & Wall, K. (1985). Two estimators for the APT model when factors are measured. *Economics Letters, 19*(3), 271-275.

McMillan, D. (2005). Non-linear dynamics in international stock market returns. *Review of Financial Economics, 14*(1), 81-91.

McMillan, D., & Ruiz, I. (2009). Volatilty persistence, long memory and time-varying conditional mean: Evidence from 10 equity indices. *The Quarterly Review of Economics and Finance, 49*(2), 578-595.

Meade, N., & Smith, I. (1985). ARARMA vs ARIMA - a study of the benefits of a new approach to forecasting. *Omega, 13*, 519-534.

Me'lard, G., & Pasteels, J. (2000). Automatic ARIMA modeling including interventions, using time series expert software. *International Journal of Forecasting, 16*, 497-508.

Miller, J. (1986). A genetic model of adaptive economic behaviour. *Working Paper, University of Michigan*.

Miller, J., & Ratti, R. (2009). Crude oil and stock markets: Stability, instability, and bubbles. *Energy Economics, 31*(4), 559-568.

Mitchell, W. (1915). The making and using of index numbers. *Bulletin of the US Bureau of Labor Statistics*.

Moody, J. (1995). Economic forecasting: Challenges and neural network solutions. *Proceedings of the international symposium on artificial neural networks.*

Moody, J., & Yarvin, N. (1992). Networks with learned unit response functions. In J. Moody, *Advances in Neural Information Processing Systems* (pp. 1048-1055). San Mateo: Morgan Kaufman.

Muth, J. (1961). Rational expectations and the theory of price movements. *Econometrica*, 315-335.

Neely, C., Weller, P., & Ulrich, J. (2009). The Adaptive Markets Hypothesis: Evidence from the Foreign Exchange Market. *Journal of Financial and Quantitative Analysis, 44*, 467-488.

Newbold, P., & Bos, T. (1989). On exponential smoothing and the assumption of deterministic trend plus white noise data generating models. *International Journal of Forecasting, 5*, 523-527.

Newbold, P., Agiakloglou, C., & Miller, J. (1994). Adventues with ARIMA software. *International Journal of Forecasting, 10*, 573-581.

Officer, R. (1972). The distribution of stock returns. *Journal of the Americal Statistical Association, 67*(340), 807-812.

Osborne, M. (1959). Brownian Motion in the Stock Market. *Operations Research, 7*(2), 145-173.

Parzen, E. (1982). ARARMA models for time series analysis and forecasting. *Journal of Forecasting, 1*, 67-87.

Pearson, K. (1905). The Problem of the Random Walk. *Nature, 295*.

Pegels, C. (1969). Exponential smoothing: Some new variations. *Management Science, 15*(5), 311-315.

Penrose, R. (1989). *The Emperor's New Mind.* London: Oxforrd University Press.

Poon, S. (2005). *A practical guide to forecasting financial market volatility.* Wiley.

Poskitt, D., & Tremayne, A. (1986). The selection and use of linear and bilinear time series models. *International Journal of Forecasting, 2*, 101-114.

Poterba, J., & Summers, L. (1988). Mean reversion in stock prices: Evidence and implications. *Journal of Financial Economics, 22*(1), 27-59.

Praetz, P. (1972). The distribution of share price changes. *Journal of Business*, 49-55.

Priestly, R. (1996). The arbitrage pricing theory, macroeconomic and financial factors, and expectations generating processes. *Journal of Banking & Finance, 20*(5), 869-890.

Qi, M. (2001). Predicting US recessions with leading indicators via neural networks. *International Journal of Forecasting, 17*, 383-401.

Rasheed, B., & Qian, B. (2004). Hurst exponent and financial market predictability. *IASTED conference on Financial Engineering and Applications* , (pp. 203-209).

Regnault, J. (1863). *Calcul des Chances et Philosophie de la Bourse.* Paris: Mallet-Bachelier et Castel.

Roberts, H. (1959). Stock-market "patterns" and financial analysis: Methodological suggestions. *Journal of Finance, 14*(1), 1-10.

Roberts, H. (1967). Statistical versus clinical prediction of the stock market. *Unpublished Manuscript*.

Roll, R., & Ross, S. A. (1980). An empirical investigation of the Arbitrage Pricing Theory. *Journal of Finance, 35*(5), 1073-1103.

Ross, S. (1976). The Arbitrage Theory of Capital Asset Pricing. *Journal of Economic Theory, 13*(3), 341-360.

Ross, S. (1989). Information and volatility: The no-arbitrage martingale approach to timing and resolution irrelevancy. *Journal of Finance, 44*, 1-17.

Rubinstein, M. (2001). Rational markets: yes or no? The affirmative case. *Financial Analysts Journal*, 15-29.

Ruck, D., Rogers, S., Kabrisky, M., Oxley, H., & Suter, B. (1990). The multilayer perceptron as an approximator to a Bayes optimal discriminant function. *IEEE Transactions on neural networks, 1*, 296-298.

Rumelhart, D., Hinton, G., & Williams, R. (1986). Learning internal representation by error propagation. In D. E. Rumelhart, & J. L. McClelland, *Parallel Distributed Processing Exploration in the Microstructure of Cognition* (p. Chapter 8). Cambridge: MIT Press.

Sabbatini, M., & Linton, O. (1998). A GARCH model of the implied volatility of the Swiss market index from option prices. *International Journal of Forecasting, 14*, 199-213.

Samuelson, P. (1965). Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review, 6*(2), 41-49.

Savage, L. (1954). *The Foundations of Statistics.* New York: Wiley.

Schalkoff, R. (1997). *Artificial Neural Networks.* New York: McGraw-Hill.

Schwartz, T. (1995). Mr. Spock has arrived. *Futures: The magazine of Commodities and Options, 24*(1), 46-48.

Seddighi, H., & Nian, W. (2004). The Chinese stock exchange market: operations and efficiency. *Applied Financial Economics, 14*(11), 787-797.

Seetharam, Y., & Britten, J. (2013). An analysis of herding behaviour during market cycles in South Africa. *Journal of Economics and Behavioral Studies, 5*(2), 89-98.

Seetharam, Y., & Britten, J. (Forthcoming). Nonlinear modelling of market cycles in South Africa. *International Journal of Emerging Markets, 10*(4).

Shachmurove, Y., & Witkowska, D. (2001). Dynamic interrelation among major work stock market: A neural network analysis. *International Journal of Business Management, 6*(1), 1-22.

Sharma, J., & Kennedy, R. (1977). A comparative analysis of stock price behaviour on the Bombay, London and New York Stock Exchanges. *Journal of Financial and Quantitative Analysis, September*, 391-413.

Sharpe, W. (1963). A simplified model for portfolio analysis. *Management Science, 9*(2), 277-293.

Sharpe, W. (1964). Capital Asset Prices: A theory of market equilibrium under conditions of risk. *Journal of Finance, 19*(3), 425-431.

Shefrin, H. (2002). Behavioral decision making: forecasting, game theory and role play. *International Journal of Forecasting, 18*(3), 375-382.

Shiller, R. (1981). The use of volatilty measures in assessing market efficiency. *Journal of Finance, 36*(2), 291-304.

Shiller, R. (2003). From efficient markets theory to behavioral finance. *The Journal of Economic Perspectives, 17*(1), 83-104.

Shleifer, A. (2000). *Inefficient markets.* Oxford: Oxford University Press.

Simon, H. (1982). *Models of bounded rationality: Empirically grounded economic reason.* MIT Press.

Simon, T. (1997). Reconceptualizing the origins of number knowledge: A “non-numerical” account. *Cognitive Development, 12*(3), 349-372.

Skalin, J., & Terasvirta, T. (1999). Another look at Swedish business cycles. *Journal of Applied Econometrics, 14*, 359-378.

Skaradzinski, D. (2003). The nonlinear behavior of stock prices: The impact of firm size, seasonality and trading frequency. *Unpublished doctoral thesis, Virginia Polytechnic Institute and State University*.

Smith, G., & Dyakova, A. (2014). African stock markets: efficiency and relative predictability. *South African Journal of Economics, 82*(2), 258-275.

Snyder, R. D. (1985). Recursive estimation of dynamic linear statistical models. *Journal of the Royal Statistical Society (B), 47*, 272-276.

Solnik, B. (1983). International arbitrage pricing theory. *Journal of Finance, 38*(2), 449-457.

Souza, L., & Smith, J. (2002). Bias in the memory for different sampling rates. *International Journal of Forecasting, 18*, 299-313.

Sullivan, R., Timmermann, A., & White, H. (1999). Data snooping, technical trading rule performance and the bootstrap. *The Journal of Finance, 54*, 1647-1691.

Surowiecki, J. (2005). *The wisdom of crowds.* London: Random House.

Swales, G., & Yoon, Y. (1992). Applying artificial neural networks to investment analysis. *Financial Analysts Journal, 48*, 78-80.

Swanson, N., & White, H. (1997). Forecasting economic time series using flexible versus fixed specification and linear versus non-linear econometric models. *International Journal of Forecasting, 13*, 439-461.

Swingler, K. (1996). *Applying Neural Networks: A Practical Guide.* New York: Academic Press.

Takens, F. (1981). Detecting strange attractors in turbulence. In *Dynamical systems and turbulence* (pp. 366-381). Heidelberg: Springer Berlin.

Taqqu, M., Teverovsky, V., & Willinger, W. (1995). Estimators for long-range dependence: an empirical study. *Fractals, 3*(4), 785-798.

Taylor, A. (1997). On the practical problems of computing seasonal unit root tests. *International Journal of Forecasting, 13*, 307-318.

Thorbecke, W. (1997). On stock market returns and monetary policy. *Journal of Finance, 52*(2), 635-654.

Tino, P., Horne, B., & Giles, C. (2001). Attractive periodic sets in discrete-time recurrent networks (with emphasis on fixed-point stability and bifurcations in two-neuron networks). *Neural Computation, 13*(6), 1379-1414.

Tkacz, G. (2001). Neural network forecasting of Canadian GDP growth. *International Journal of Forecasting, 17*, 57-69.

Todea, A., Ulici, M., & Silaghi, S. (2009). Adaptive Markets Hypothesis: Evidence from Asia-Pacific Financial Markets. *Review of Finance and Banking, 1*(1), 07-13.

Tong, H. (1983). *Threshold models in non-linear time series analysis.* New York: Springer-Verlag.

Tsang, E. (2008). Computational intelligence determines effective rationality. *Journal of Automation and Computing, 5*(1), 63-66.

Tseng, K. (2006). Behavioral finance, bounded rationality, neuro-finance and traditional finance. *Investment Management and Financial Innovations, 3*(4), 7-18.

van Gysen, M., Huang, C., & Kruger, R. (2013). The Performance Of Linear Versus Non-Linear Models In Forecasting Returns On The Johannesburg Stock Exchange. *International Business & Economics Research Journal*, 985-994.

van Rensburg, P. (1996). Macroeconomic identification of the priced APT factors on the Johannesburg Stock Exchange. *South African Journal of Business Management, 27*(4), 104-112.

van Rensburg, P. (2000). Macroeconomic variables and the cross-section of Johannesburg Stock Exchange returns. *South African Journal of Business Management, 31*(1), 31-43.

Volterra, V. (1930). *Theory of functionals and of integro-differential equations.* New York: Dover.

Wei, K., & Wong, K. (1992). Tests of inflation and industry portfolio stock returns. *Journal of Economics and Business, 44*(1), 77-94.

Weiss, S., & Kulikowski, C. (1991). *Computer systems that learn.* San Mateo: Morgan Kauffman.

Weron, R. (2002). Estimating long range dependence: finite sample properties and confidence intervals. *Physica A: Statistical Mechanics and its Applications, 325*(1), 285-299.

Wilson, E. (1975). *Sociobiology: The new synthesis.* Cambridge: Belk-nap/Harvard University Press.

Wright, J. (2000). Alternative variance-ratio tests using ranks and signs. *Journal of Business & Economic Statistics, 18*(1), 1-9.

Wythoff, B. (1993). Backpropagation neural networks: A tutorial. *Chemometrics and Intelligent Laboratory Systems, 18*(2), 115-155.

Yule, G. (1927). On the method of investigating periodicities in disturbed series, with special reference to Wolfer's sunspot numbers. *Philosophical Transactions of the Royal Society London, Series A, 226*, 267-298.

Zunino, L., Zanin, M., Tabak, B., Prez, D., & Rosso, O. (2009). Forbidden patterns, permutation entropy and stock market inefficiency. *Physica A: Statistical Mechanics and its Applications, 388*(14), 2584-2684.

Zupan, J., & Gasteiger, J. (1991). Neural networks: A new method for solving chemical problems or just a passing phase? *Analytica Chimica Acta, 248*(1), 1-30.

Appendix A

The BDS Test

The intuition behind the test is as follows. Let $X_t$be a univariate time series, independent and identically distributed from some distribution. Let P_A represent the probability that two points are within a distance ε of each other. In other words

$$P_A = P(|X_t - X_S| < \varepsilon) \quad \{A1\}$$

Further define

$$P_B = P(|X_t - X_S| < \varepsilon, |X_{t-1} - X_{s-1}| < \varepsilon) \quad \{A2\}$$

as the probability of a history of two points being within a distance ε of each other. Under independence of X_t, the two events contained in event B are independent implying that $P_B = P_A^2$. Therefore, it is possible to estimate P_A, $P_B$and $P_B - P_A^2$ which has an expected value of zero under the null hypothesis. To estimate the probability that two m length vectors are within ε, define

$$c_{m,n}(\varepsilon) = \frac{2}{(n-m+1)(n-m)} \sum_{s=m}^{n} \sum_{t=s+1}^{n} \prod_{j=0}^{m-1} I_\varepsilon(X_{s-j}, X_{t-j}) \quad \{A3\}$$

where

$$I_\varepsilon(X_{s-j}, X_{t-j}) = \begin{cases} 1 \; if \; |X_{s-j} - X_{t-j}| < \varepsilon \\ 0 \; otherwise \end{cases} \quad \{A4\}$$

n is the sample size and *m* is the embedding dimension. Under the null hypothesis of an independent and identical distribution,

$$E\left(c_{m,n}(\varepsilon)\right) = (E(c_{1,n}(\varepsilon)))^m \quad \{A5\}$$

Brock, Dechert and Scheikman (1987) show that given an embedding dimension, m, and the value of the radius, ε, the BDS statistic

$$w_{m,n}(\varepsilon) = \sqrt{n-m+1}\frac{c_{m,n}(\varepsilon) - c_{1,n-m+1}^{m}(\varepsilon)}{\sigma_{m,n}(\varepsilon)} \quad \{A6\}$$

is asymptotically distributed following N ~ (0,1), The consistent estimator, $\sigma_{m,n}^2(\varepsilon)$ is given by

$$\sigma_{m,n}^2(\varepsilon) = 4[k^m + 2\sum_{j=1}^{m-1} k^{m-j}\, c^{2j} + (m-1)^2 c^{2m} + m^2 k c^{2m-2} \quad \{A7\}$$

where

$$c = c_{1,n}(\varepsilon) \quad \{A8\}$$

$$k = k_n(\varepsilon) = \frac{6}{n(n-1)(n-2)}\sum_{t=1}^{n}\sum_{s=t+1}^{n}\sum_{r=S+1}^{n} h_\varepsilon(X_t, X_s, X_r) \quad \{A9\}$$

$$h_\varepsilon(i,j,k) = \frac{1}{3}[I_\varepsilon(i,j)I_\varepsilon(j,k) + I_\varepsilon(i,k)I_\varepsilon(k,j) + I_\varepsilon(i,j)I_\varepsilon(i,k) \quad \{A10\}$$

Kanzler (1999) shows that the consistent estimators $c_{1,n}(\varepsilon)$ and $k_n(\varepsilon)$ are in the class of U statistics and are the most efficient estimates of c and k respectively. The BDS test statistic is a two-sided test, with the null hypothesis of independent and identical distributions being rejected at the 5% level of statistical significance.

Appendix B

Table B 1 - Summary of results for SAB

SAB	Consumer Goods - Beverage	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Random	Random	Random	Random	Random
	Chow Denning test	Not IID	Not IID	IID	IID	IID
		Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	Not IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Not Random	Random	Random	Random
	Hurst exponent	Not Random	Not Random	Random	Non-Random	Non-Random
Modelling	SETAR MAPE	144.20%	139.70%	166.80%	237.50%	131.60%
	NAR R2	16.48	10.05	14.48		

Table B 2 - Summary of results for BIL

BIL	Basic Materials - Mining	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Normal	Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Random	Random	Non-Random	Random	Random
	Chow Denning test	Not IID	Not IID	IID	IID	IID
		Correlated and Heteroscedasticity	Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Not Random	Random	Random	Random
	Hurst exponent	Not Random	Not Random	Not Random	Non-Random	Non-Random
Modelling	SETAR MAPE	139.70%	142.90%	122.40%	994.50%	355.20%
	NAR R2	11.85	12.19	15.46		

Table B 3 - Summary of results for NPN

NPN	Consumer Services - Media	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	Not IID	IID	IID	IID	IID
		Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	IID	IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Random	Random	Random	Random
	Hurst exponent	Random	Random	Not Random	Non-Random	Random
Modelling	SETAR MAPE	273.70%	267.70%	341.70%	201.50%	172.60%
	NAR R2	9.41	13.42	13.7		

Table B 4 - Summary of results for MTN

MTN	Telecommunications - Mobile Telecommunicatons	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Random	Random	Random	Random	Random
	Chow Denning test	Not IID Uncorrelated with No Heteroscedasticity	Not IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	Not IID	Not IID	IID
	Variance decomposition plot	Not Random	Not Random	Random	Random	Random
	Hurst exponent	Random	Not Random	Not Random	Non-Random	Non-Random
Modelling	SETAR MAPE	138.50%	143.40%	126.70%	261.50%	646.80%
	NAR R2	12.98	15.64	19.49		

Table B 5 - Summary of results for SOL

SOL	Oil and Gas - Oil and Gas Producers	Daily	Weekly	Monthly	Quarterly	Semi Annual
	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Normality	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear
	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
Stationary	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Non-Random	Non-Random	Non-Random	Random
		Not IID	Not IID	IID	IID	IID
	Chow Denning test	Correlated and Heteroscedasticity	Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	Not IID	Not IID	IID
	Variance decomposition plot	Not Random	Not Random	Random	Random	Random
Random walk	Hurst exponent	Not Random	Not Random	Not Random	Non-Random	Non-Random
	SETAR MAPE	213.50%	228.90%	232.10%	200.10%	249.70%
Modelling	NAR R2	14.04	12.59	12.62		

Table B 6 - Summary of results for AGL

AGL	Basic Materials - Mining	Daily	Weekly	Monthly	Quarterly	Semi Annual
	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Normality	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear	Linear and Non-linear	Linear and Non-linear
	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
Stationary	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	Not IID	IID	IID	IID	IID
		Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	IID	IID	IID	IID
	Variance decomposition plot	Not Random	Random	Random	Random	Random
Random walk	Hurst exponent	Random	Random	Random	Non-Random	Non-Random
	SETAR MAPE	118.70%	116.10%	120.70%	144.30%	122.60%
Modelling	NAR R2	6.57	10.59	17.17		

Table B 7 - Summary of results for FSR

FSR	Financials - Banks	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	NON Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Random	Random	Random	Random	Random
	Chow Denning test	Not IID	Not IID	IID	IID	Not IID
		Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Correlated and Heteroscedasticity
	Wright test	Not IID	Not IID	IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Not Random	Random	Random	Random
	Hurst exponent	Not Random	Not Random	Not Random	Non-Random	Random
Modelling	SETAR MAPE	152.20%	141.60%	164.40%	175.10%	180.60%
	NAR R2	8.86	20.15	14.04		

Table B 8 - Summary of results for SBK

SBK	Financials - Banks	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Random	Non-Random	Random	Random	Random
	Chow Denning test	Not IID	Not IID	IID	IID	IID
		Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Not Random	Random	Random	Random
	Hurst exponent	Not Random	Not Random	Not Random	Non-Random	Random
Modelling	SETAR MAPE	145.50%	153.40%	145.60%	233.00%	250.10%
	NAR R2	5.95	20.51	26.11		

Table B 9 - Summary of results for APN

APN	Healthcare - Pharmaceuticals and Biotechnology	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Linear and Non-linear	Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Random	Random	Random	Non-Random	Random
	Chow Denning test	Not IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	IID	IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Not Random	Random	Random	Random
	Hurst exponent	Random	Random	Not Random	Non-Random	Non-Random
Modelling	SETAR MAPE	474.80%	833.30%	977.10%	245.00%	181.10%
	NAR R2	22.9	23.14	25.74		

Table B 10 - Summary of results for BGA

BGA	Financials - Banks	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear	Linear and Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	NON Stationary	NON Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	Not IID	Not IID	IID	Not IID	IID
		Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	IID	IID	IID
	Variance decomposition plot	Not Random	Not Random	Random	Random	Random
	Hurst exponent	Not Random	Not Random	Not Random	Non-Random	Random
Modelling	SETAR MAPE	121.60%	125.80%	133.60%	108.80%	1420.00%
	NAR R2	8.97	13.05	15.55		

Table B 11 - Summary of results for SHF

SHF	Consumer Goods - Household Goods and Home Construction	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	NON Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Non-Random	Random	Random	Random
	Chow Denning test	Not IID	IID	IID	IID	IID
		Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	IID	IID	IID
	Variance decomposition plot	Random	Not Random	Random	Random	Random
	Hurst exponent	Random	Not Random	Not Random	Random	Non-Random
Modelling	SETAR MAPE	108.10%	115.90%	112.90%	131.50%	386.70%
	NAR R2	10.54	11.32	16.36		

Table B 12 - Summary of results for RMH

RMH	Financials - Banks	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear	Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	Not IID	IID	IID	Not IID	Not IID
		Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Correlated and Heteroscedasticity	Correlated and Heteroscedasticity
	Wright test	Not IID	IID	IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Random	Random	Random	Random
	Hurst exponent	Not Random	Random	Not Random	Non-Random	Random
Modelling	SETAR MAPE	170.10%	170.50%	220.70%	118.00%	205.70%
	NAR R2	7.75	15.58	23.49		

Table B 13 - Summary of results for MDC

MDC	Health Care - Health Care Equipment and Services	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	NON Stationary	Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Random	Random	Random	Random	Non-Random
	Chow Denning test	IID	Not IID	IID	Not IID	IID
		Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	IID	Not IID	Not IID	Not IID
	Variance decomposition plot	Random	Random	Random	Random	Random
	Hurst exponent	Random	Random	Not Random	Non-Random	Non-Random
Modelling	SETAR MAPE	162.60%	163.20%	158%	434.60%	434.50%
	NAR R2	13.32	12.17	43.67		

Table B 14 - Summary of results for INP

INP	Financials - Financial Services	Daily	Weekly	Monthly	Quarterly	Semi Annual
	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Normality	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear	Linear and Non-linear
	ADF test	Stationary	Stationary	Stationary	NON Stationary	Stationary
Stationary	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	Not IID	IID	IID	IID	IID
		Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	IID	IID	IID	IID
	Variance decomposition plot	Not Random	Random	Random	Random	Random
Random walk	Hurst exponent	Not Random	Random	Random	Non-Random	Non-Random
	SETAR MAPE	116.30%	116.10%	106.40%	116.40%	1004.00%
Modelling	NAR R2	10.34	16.74	13.04		

Table B 15 - Summary of results for MPC

MPC	Consumer Services - General Retailers	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	Not IID	IID	IID	IID	IID
		Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Random	Random	Random	Random
	Hurst exponent	Random	Random	Random	Non-Random	Non-Random
Modelling	SETAR MAPE	180.30%	195.10%	194.10%	170.90%	100.30%
	NAR R2	11.01	10.03	7.29		

Table B 16 - Summary of results for IMP

IMP	Basic Materials - Mining	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	NON Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Non-Random	Random
	Chow Denning test	Not IID	Not IID	IID	Not IID	IID
		Correlated and Heteroscedasticity	Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	Not IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Not Random	Random	Not Random	Random
	Hurst exponent	Not Random	Not Random	Random	Random	Non-Random
Modelling	SETAR MAPE	141.30%	143.90%	155%	218.10%	151.10%
	NAR R2	10.9	11.37	14.73		

Table B 17 - Summary of results for NTC

NTC	Health Care - Health Care Equipment and Services	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Non-linear	Linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Random	Random	Non-Random	Random	Non-Random
	Chow Denning test	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	Not IID	Not IID	Not IID
	Variance decomposition plot	Random	Random	Random	Random	Random
	Hurst exponent	Random	Random	Not Random	Non-Random	Non-Random
Modelling	SETAR MAPE	130.50%	126.30%	123.90%	147.60%	177.10%
	NAR R2	8.43	15.52	19.85		

Table B 18 - Summary of results for MMI

MMI	Financials - Life Insurance	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Non-linear	Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	NON Stationary	NON Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Random	Random	Random	Random	Random
	Chow Denning test	Not IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	IID	IID	IID	Not IID
	Variance decomposition plot	Not Random	Random	Random	Random	Random
	Hurst exponent	Random	Random	Not Random	Non-Random	Non-Random
Modelling	SETAR MAPE	150.20%	164.30%	123.30%	111.70%	295.60%
	NAR R2	5.21	15.11	16.38		

Table B 19 - Summary of results for ANG

ANG	Basic Materials - Mining	Daily	Weekly	Monthly	Quarterly	Semi Annual
	Jarque Bera	Non-Normal	Non-Normal	Normal	Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Normality	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear	Linear and Non-linear	Linear and Non-linear
	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
Stationary	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	Not IID Uncorrelated with No Heteroscedasticity	Not IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	IID	IID	IID
	Variance decomposition plot	Not Random	Not Random	Random	Random	Random
Random walk	Hurst exponent	Not Random	Random	Not Random	Non-Random	Random
	SETAR MAPE	112.90%	103.30%	123.10%	264.80%	206.40%
Modelling	NAR R2	4.13	13.64	10.85		

Table B 20 - Summary of results for IPL

IPL	Industrials - Industrial Transportation	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear	Linear	Linear
Stationary	ADF test	Stationary	Stationary	Stationary	NON Stationary	Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Random	Random	Random	Random	Random
	Chow Denning test	Not IID	IID	IID	IID	IID
		Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	IID	IID	Not IID
	Variance decomposition plot	Not Random	Random	Random	Random	Random
	Hurst exponent	Not Random	Random	Not Random	Random	Non-Random
Modelling	SETAR MAPE	121.20%	122.70%	118.40%	157.80%	583.60%
	NAR R2	9.63	13.99	12.49		

Table B 21 - Summary of results for NPK

NPK	Industrials - General Industrials	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Non-Random	Random	Non-Random	Random
	Chow Denning test	Not IID Uncorrelated with No Heteroscedasticity	Not IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	IID	Not IID	IID
	Variance decomposition plot	Not Random	Random	Random	Random	Random
	Hurst exponent	Random	Random	Not Random	Non-Random	Non-Random
Modelling	SETAR MAPE	130%	146.10%	128.60%	154.50%	126.30%
	NAR R2	7.6	16.01	16.85		

Table B 22 - Summary of results for GFI

GFI	Basic Materials - Mining	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Linear and Non-linear	Linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	Not IID Uncorrelated with No Heteroscedasticity	Not IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	IID	IID	IID
	Variance decomposition plot	Random	Not Random	Random	Random	Random
	Hurst exponent	Not Random	Not Random	Random	Non-Random	Non-Random
Modelling	SETAR MAPE	15486%	9771%	13040%	134.90%	196.30%
	NAR R2	12.32	10.57	20.21		

Table B 23 - Summary of results for ASR

ASR	Basic Materials - Mining	Daily	Weekly	Monthly	Quarterly	Semi Annual
	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Normality	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear	Linear and Non-linear
	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
Stationary	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Random	Non-Random	Non-Random	Random	Random
	Chow Denning test	Not IID	IID	IID	IID	IID
		Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	Not IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Random	Random	Random	Random
Random walk	Hurst exponent	Random	Random	Not Random	Random	Non-Random
	SETAR MAPE	192.30%	173.90%	169.50%	659.70%	91.39%
Modelling	NAR R2	6.02	11.91	15.12		

Table B 24 - Summary of results for INL

INL	Financials - Financial Services	Daily	Weekly	Monthly	Quarterly	Semi Annual
	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Normality	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear	Linear and Non-linear
	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
Stationary	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	Not IID	IID	IID	IID	IID
		Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	IID	IID	IID
	Variance decomposition plot	Not Random	Not Random	Random	Random	Random
Random walk	Hurst exponent	Not Random	Random	Random	Non-Random	Non-Random
	SETAR MAPE	135.60%	126.10%	119.50%	110.30%	365.50%
Modelling	NAR R2	9.77	10.59	16.51		

Table B 25 - Summary of results for PIK

PIK	Consumer Services - Food and Drug Retailers	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Random	Non-Random	Random	Non-Random	Random
	Chow Denning test	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	Not IID Uncorrelated with No Heteroscedasticity	Not IID Correlated and Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	IID	Not IID	Not IID
	Variance decomposition plot	Random	Random	Not Random	Random	Random
	Hurst exponent	Random	Random	Not Random	Non-Random	Non-Random
Modelling	SETAR MAPE	127%	120.70%	125.60%	119.50%	193.60%
	NAR R2	9.13	7.06	46.44		

Table B 26 - Summary of results for TFG

TFG	Consumer Services - General Retailers	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	NON Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Non-Random	Random	Random
	Chow Denning test	Not IID	IID	Not IID	IID	IID
		Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	Not IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Random	Not Random	Random	Random
	Hurst exponent	Random	Random	Random	Random	Non-Random
Modelling	SETAR MAPE	129.50%	130.50%	130.40%	156.70%	211.90%
	NAR R2	10.27	20.72	15.1		

Table B 27 - Summary of results for SNT

SNT	Financials - Nonlife Insurance	Daily	Weekly	Monthly	Quarterly	Semi Annual
	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Normality	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear
	ADF test	Stationary	Stationary	Stationary	Stationary	NON Stationary
Stationary	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Non-Random	Random
	Chow Denning test	IID	IID	IID	Not IID	IID
		Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	IID	IID	Not IID	Not IID
	Variance decomposition plot	Random	Random	Random	Random	Random
Random walk	Hurst exponent	Random	Random	Not Random	Non-Random	Non-Random
	SETAR MAPE	148.40%	147.40%	226.50%	196.00%	202.50%
Modelling	NAR R2	16.98	12.9	16.67		

Table B 28 - Summary of results for HYP

HYP	Financials - Real Estate Investment Trusts	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	NON Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	IID	Not IID	Not IID	Not IID
	Variance decomposition plot	Random	Random	Random	Random	Random
	Hurst exponent	Random	Random	Not Random	Non-Random	Non-Random
Modelling	SETAR MAPE	6931%	7059%	7593%	202.50%	4078.00%
	NAR R2	13.31	12.31	23.43		

Table B 29 - Summary of results for SAP

SAP	Basic Materials - Forestry and Paper	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Non-linear	Linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	NON Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	Not IID	IID	IID	IID	IID
		Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	IID	IID	IID	IID
	Variance decomposition plot	Not Random	Not Random	Random	Random	Random
	Hurst exponent	Random	Random	Not Random	Non-Random	Random
Modelling	SETAR MAPE	115.30%	110.20%	115.60%	135.60%	265.60%
	NAR R2	7.33	10.26	11.83		

Table B 30 - Summary of results for CLS

CLS	Consumer Services - Food and Drug Retailers	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	NON Stationary	Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Non-Random	Random
	Chow Denning test	Not IID Correlated and Heteroscedasticity	Not IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	Not IID	IID	Not IID
	Variance decomposition plot	Not Random	Not Random	Random	Random	Random
	Hurst exponent	Random	Not Random	Not Random	Random	Non-Random
Modelling	SETAR MAPE	155.80%	156.60%	159%	163.50%	486.20%
	NAR R2	10.47	11.09	23.13		

Table B 31 - Summary of results for GND

GND	Industrials - Industrial Transportation	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	NON Stationary	NON Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Non-Random	Non-Random
	Chow Denning test	Not IID Uncorrelated with No Heteroscedasticity	Not IID Correlated and Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	Not IID Correlated and Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	Not IID	Non IID	Not IID
	Variance decomposition plot	Not Random	Not Random	Random	Not Random	Random
	Hurst exponent	Not Random	Not Random	Random	Non-Random	Non-Random
Modelling	SETAR MAPE	153.70%	151.80%	169.20%	157.20%	318.20%
	NAR R2	9.64	15.48	28.97		

Table B 32 - Summary of results for PPC

PPC	Industrials - Construction and Materials	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear	Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	NON Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Non-Random	Random
	Chow Denning test	IID	IID	IID	IID	IID
		Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	IID	IID	IID	Not IID
	Variance decomposition plot	Random	Random	Random	Random	Random
	Hurst exponent	Random	Random	Not Random	Non-Random	Non-Random
Modelling	SETAR MAPE	181.80%	177.80%	197.60%	220.50%	157.60%
	NAR R2	12.74	12.69	31.12		

Table B 33 - Summary of results for AFE

AFE	Basic Materials - Chemicals	Daily	Weekly	Monthly	Quarterly	Semi Annual
	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Normality	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear	Non-linear
	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
Stationary	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Non-Random
	Chow Denning test	Not IID	IID	IID	IID	Not IID
		Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Correlated and Heteroscedasticity
	Wright test	Not IID	Not IID	IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Not Random	Random	Random	Random
Random walk	Hurst exponent	Random	Random	Random	Non-Random	Non-Random
	SETAR MAPE	190.60%	277.70%	283.50%	187.10%	102.20%
Modelling	NAR R2	9.93	13.95	16.45		

Table B 34 - Summary of results for RCL

RCL	Consumer Goods - Food Producers	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Non-linear	Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	Not IID	IID	IID	IID	IID
		Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Random	Random	Random	Random
	Hurst exponent	Random	Random	Not Random	Non-Random	Non-Random
Modelling	SETAR MAPE	116.70%	130.70%	127.70%	324.40%	264.40%
	NAR R2	10.1	20.2	21.47		

Table B 35 - Summary of results for SUI

SUI	Consumer Services - Travel and Leisure	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear and Non-linear	Non-linear	Linear
Stationary	ADF test	Stationary	Stationary	Stationary	NON Stationary	NON Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Non-Random	Random
	Chow Denning test	Not IID	IID	IID	IID	IID
		Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	IID	IID	Not IID	IID
	Variance decomposition plot	Not Random	Random	Random	Random	Random
	Hurst exponent	Random	Random	Not Random	Random	Non-Random
Modelling	SETAR MAPE	139.80%	131.40%	168.60%	175.20%	176.80%
	NAR R2	10.89	14	28.02		

Table B 36 - Summary of results for ILV

ILV	Consumer Goods - Food Producers	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear	Linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	NON Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	Not IID	Not IID	IID	IID	IID
		Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	IID	IID	IID	IID
	Variance decomposition plot	Not Random	Not Random	Random	Random	Random
	Hurst exponent	Random	Not Random	Not Random	Random	Random
Modelling	SETAR MAPE	112.50%	110.90%	116.60%	112.60%	305.90%
	NAR R2	9.11	11.13	11.08		

Table B 37 - Summary of results for RLO

RLO	Industrials - Electronic and Electronic Equipment	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Non-linear	Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	NON Stationary	Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	IID	IID	Not IID	Not IID
	Variance decomposition plot	Random	Random	Random	Random	Random
	Hurst exponent	Random	Random	Not Random	Non-Random	Non-Random
Modelling	SETAR MAPE	122.40%	130.10%	134.50%	210.90%	127.40%
	NAR R2	8.47	21.33	16.07		

Table B 38 - Summary of results for FBR

FBR	Consumer Goods - Travel and Leisure	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Linear and Non-linear	Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	NON Stationary	NON Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Random	Random	Random	Random	Non-Random
	Chow Denning test	Not IID	Not IID	IID	IID	IID
		Correlated and Heteroscedasticity	Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	Not IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Not Random	Random	Random	Random
	Hurst exponent	Not Random	Not Random	Random	Non-Random	Random
Modelling	SETAR MAPE	200.40%	233.80%	270.80%	404.60%	709.20%
	NAR R2	12.48	12.44	10.77		

Table B 39 - Summary of results for MUR

MUR	Industrials - Construction and Materials	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear and Non-linear	Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	NON Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	Not IID	IID	IID	IID	IID
		Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	IID	Not IID	IID	Not IID
	Variance decomposition plot	Not Random	Random	Not Random	Random	Random
	Hurst exponent	Random	Random	Not Random	Non-Random	Random
Modelling	SETAR MAPE	141832%	359459%	224300%	102.30%	104.60%
	NAR R2	8.31	17.54	17.18		

Table B 40 - Summary of results for SPG

SPG	Industrials - Industrial Transportation	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	NON Stationary	NON Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	Not IID	Not IID	IID	IID	IID
		Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	IID	Not IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Not Random	Random	Random	Random
	Hurst exponent	Random	Random	Not Random	Random	Non-Random
Modelling	SETAR MAPE	110.40%	114%	111.90%	156.70%	132.00%
	NAR R2	6.26	16.45	19.58		

Table B 41 - Summary of results for FPT

FPT	Financials - Real Estate Investment Trusts	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	NON Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Non-Random	Random	Non-Random	Random
	Chow Denning test	Not IID	Not IID	IID	IID	IID
		Correlated and Heteroscedasticity	Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	IID	Not IID	IID
	Variance decomposition plot	Not Random	Not Random	Random	Random	Random
	Hurst exponent	Not Random	Not Random	Not Random	Non-Random	Non-Random
Modelling	SETAR MAPE	108.20%	109.10%	113.80%	148.50%	172.80%
	NAR R2	11.04	12.59	15.2		

Table B 42 - Summary of results for SAC

SAC	Financials - Real Estate Investment Trusts	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear and Non-linear	Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	NON Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Non-Random	Random	Random	Random
	Chow Denning test	Not IID	Not IID	IID	IID	IID
		Correlated and Heteroscedasticity	Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Not Random	Random	Random	Random
	Hurst exponent	Not Random	Not Random	Not Random	Random	Random
Modelling	SETAR MAPE	105.80%	111.20%	114%	129.30%	346.50%
	NAR R2	6.39	18.43	12.81		

Table B 43 - Summary of results for OCE

OCE	Consumer Goods - Food Producers	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	NON Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Random	Random	Random	Random	Random
	Chow Denning test	Not IID	Not IID	IID	Not IID	IID
		Correlated and Heteroscedasticity	Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	IID	Not IID	Not IID
	Variance decomposition plot	Random	Not Random	Random	Random	Random
	Hurst exponent	Not Random	Not Random	Not Random	Random	Random
Modelling	SETAR MAPE	131.40%	134%	133.50%	203.10%	200.50%
	NAR R2	7.53	15.21	18.63		

Table B 44 - Summary of results for WBO

WBO	Industrials - Construction and Materials	Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	NON Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Non-Random	Non-Random	Random	Random
	Chow Denning test	Not IID	IID	IID	IID	IID
		Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	Not IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Random	Random	Random	Random
	Hurst exponent	Random	Random	Not Random	Non-Random	Non-Random
Modelling	SETAR MAPE	258.60%	254.90%	253.30%	120.60%	154.70%
	NAR R2	21.27	12.87	34.69		

Table B 45 - Summary of results for the J150

J150		Daily	Weekly	Monthly	Quarterly	Semi Annual
	Jarque Bera	Non-Normal	Non-Normal	Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Normality	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear	Non-linear	Linear
	ADF test	Stationary	Stationary	Stationary	Stationary	Stationary
Stationary	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
		Not IID	IID	IID	IID	IID
	Chow Denning test	Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	IID	IID	IID
	Variance decomposition plot	Not Random	Random	Random	Random	Random
Random walk	Hurst exponent	Random	Random	Not Random	Non-Random	Non-Random
	SETAR MAPE	142.60%	115.70%	159.60%	133.00%	422.90%
Modelling	NAR R2	10.36	12.03	12.2		

Table B 46 - Summary of results for the J200

J200		Daily	Weekly	Monthly	Quarterly	Semi Annual
	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Normality	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Non-linear	Linear and Non-linear
	ADF test	Stationary	Stationary	Stationary	Stationary	NON Stationary
Stationary	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	Not IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	IID	Not IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Random	Random	Random	Random
Random walk	Hurst exponent	Random	Random	Not Random	Non-Random	Non-Random
	SETAR MAPE	123.60%	123.60%	125.80%	151.00%	186.40%
Modelling	NAR R2	11.28	12.16	9.62		

Table B 47 - Summary of results for ALSI

J203		Daily	Weekly	Monthly	Quarterly	Semi Annual
	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Normality	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Non-linear	Non-linear	Linear and Non-linear
	ADF test	Stationary	Stationary	Stationary	Stationary	NON Stationary
Stationary	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	Not IID	IID	IID	IID	IID
		Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	Not IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Random	Random	Random	Random
Random walk	Hurst exponent	Random	Random	Not Random	Non-Random	Non-Random
	SETAR MAPE	297.20%	283.20%	241.60%	118.30%	167.80%
	NAR R2	12.8	16.49	23.87		
Modelling	NARX R2	38.91	24.55	36.94		

Table B 48 - Summary of results for the J211

J211		Daily	Weekly	Monthly	Quarterly	Semi Annual
Normality	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear and Non-linear	Non-linear	Linear and Non-linear
Stationary	ADF test	Stationary	Stationary	Stationary	Stationary	NON Stationary
	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
Random walk	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	Not IID	IID	IID	IID	IID
		Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Random	Random	Random	Random
	Hurst exponent	Random	Random	Not Random	Non-Random	Non-Random
Modelling	SETAR MAPE	332.50%	370.20%	418.90%	180.70%	573.90%
	NAR R2	13.19	17.31	17.74		

Table B 49 - Summary of results for the J213

J213		Daily	Weekly	Monthly	Quarterly	Semi Annual
	Jarque Bera	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Normality	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear	Linear and Non-linear
	ADF test	Stationary	Stationary	Stationary	NON Stationary	Stationary
Stationary	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	Not IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity	IID Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	Not IID	IID	Not IID	Not IID
	Variance decomposition plot	Not Random	Random	Random	Random	Random
Random walk	Hurst exponent	Random	Random	Not Random	Non-Random	Non-Random
	SETAR MAPE	313.60%	215.70%	1006%	139.20%	82.06%
Modelling	NAR R2	12.73	15.72	14.13		

Table B 50 - Summary of results for the J177

J177		Daily	Weekly	Monthly	Quarterly	Semi Annual
	Jarque Bera	Non-Normal	Non-Normal	Normal	Normal	Normal
	Q-Q Plot	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Normality	K-S test	Non-Normal	Non-Normal	Non-Normal	Non-Normal	Non-Normal
Non-linearity	BDS test	Non-linear	Non-linear	Linear and Non-linear	Linear and Non-linear	Linear and Non-linear
	ADF test	Stationary	Stationary	Stationary	Stationary	NON Stationary
Stationary	KPSS test	Stationary	Stationary	Stationary	Stationary	Stationary
	Rolling Autocorrelation	Cyclical	Cyclical	Cyclical	Non-Cyclical	Non-Cyclical
	Runs test	Non-Random	Random	Random	Random	Random
	Chow Denning test	Not IID	IID	IID	IID	IID
		Correlated and Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity	Uncorrelated with No Heteroscedasticity
	Wright test	Not IID	IID	IID	IID	IID
	Variance decomposition plot	Not Random	Random	Random	Random	Random
Random walk	Hurst exponent	Random	Random	Random	Non-Random	Non-Random
	SETAR MAPE	113.60%	115.30%	115.30%	248.30%	157.10%
Modelling	NAR R2	7.14	12.15	15.53		

Plagiarism Report

