

***An Empirical Investigation of
Sector Correlation Forecasting
Techniques and the Potential
Benefits to Investors***

Stephen Fraser Metcalfe
Student No: 9810221D

Supervisor: Prof. R.J. Thomson
School of Statistics & Actuarial Science
University of the Witwatersrand

December 2007

Declaration

I, Stephen Fraser Metcalfe, declare that, to the extent usually and reasonably expected, professional guidance and supervision were received from my supervisor. As mentioned in the acknowledgements, certain assistance was received from other persons. In all other respects, this research report is my own work. It is being submitted in part fulfilment for the degree of Master of Science in Actuarial Science in the School of Statistics and Actuarial Science, Faculty of Science, University of the Witwatersrand, Johannesburg. It has not been submitted before for any degree or examination at this or any other University.

S F Metcalfe

This 5th day of August 2008

Acknowledgements

I would like to thank Professors Rob Thomson and Jackie Galpin for their professional guidance and supervision of this research; my colleagues Lethabo and Luke for their assistance in obtaining the data; and my parents Dr Barry Metcalfe and Professor Mary Ross for their moral, financial and editorial support to complete the MSc course work and research report.

Abstract

Financial modelling is of considerable value to portfolio management. The effectiveness of different methods of forecasting correlation between sub-sectors, as part of the sector-allocation stage of the portfolio-construction process, has not yet been investigated. This focus is useful since it is relatively practical to collect data pertaining to sector and sub-sector indices, and hence the calculation of figures necessary to determine their investment performance is simpler.

The aim of this research paper was to examine the performance of various correlation estimation techniques under two assessment criteria and to identify, if possible, the most suitable methods to employ in the sub-sector allocation stage of the 'top-down' approach to portfolio construction. Monthly total returns were calculated for each of the market indices, the sectors and their sub-sectors from the relevant total return indices as part of the analysis. The first assessment criterion was the statistical performance of the methods, which measured their ability to estimate future correlation coefficients between different sub-sectors by analysing the distributions of their absolute forecast errors. The second assessment criterion was the economic performance of the forecast methods. MPT was used to select the optimal portfolios for certain levels of expected return and the economic performance of the efficient sub-sector allocations, selected using the different correlation estimation techniques, was then evaluated.

The two models used to estimate correlation that stood out from the rest in terms of their overall performance were the full HCM model and the industry mean model. From the perspective of the statistical performance criterion, the industry mean model consistently performed the best and the full HCM model also performed well. The economic performance of all the models tested, with the exception of the overall mean model, outperformed the passive investment strategy of holding the market portfolio. The economic performance of the full HCM model was best overall and that of the industry mean model was also strong. Prior research has found that the industry mean model is useful in forecasting future correlation between individual shares. This research found that the industry mean model also has value in forecasting future correlation between sub-sectors. Furthermore, despite demonstration in prior research of the full HCM model's poor ability to estimate future correlation between individual shares, it was one of the most effective models at

forecasting correlation between sub-sectors. Both of these models therefore hold value to investors for the purposes of sub-sector allocation as part of a top-down approach to financial portfolio construction.

List of acronyms

ALSI	All-share index
CAPM	Capital asset pricing model
EMH	Efficient markets hypothesis
FINI	Financial index
FTSE	Financial Times Stock Exchange
FTSE100	Financial Times Stock Exchange Top 100 index
HCM	Historical correlation matrix
INDI	Industrial index
LSE	London Stock Exchange
MPT	Modern portfolio theory
PE	Price—Earnings
SIC	Standard industrial classification
SSI	Standard single-index
S&P	Standard & Poors
TRI	Total-return index

List of tables

Table 1	– Average absolute forecast error of correlation coefficients
Table 2	– Relative stochastic dominance of correlation coefficients estimates
Table 3	– Relative stochastic dominance in first three-year period (P1)
Table 4	– Relative stochastic dominance in second three-year period (P2)
Table 5	– Relative stochastic dominance in combined period 2 (CP2)
Table 6	– Economic performance using FTSE ALSI return as the target return
Table 7	– Economic performance using FTSE100 return as the target return

List of figures

Graph 1	– Cumulative distribution of absolute forecast errors in period 1 (P1)
Graph 2	– Cumulative distribution of absolute forecast errors in period 2 (P2)
Graph 3	– Cumulative distribution of absolute forecast errors in combined second three-year period (CP2)

TABLE OF CONTENTS

Declaration
Acknowledgements
Abstract
List of Acronyms
List of Tables
List of Figures

CHAPTER 1: INTRODUCTION AND BACKGROUND

- 1.1 Introduction
- 1.2 Theoretical overview
 - 1.2.1 Modern portfolio theory
 - 1.2.2 Modelling correlation structures
 - 1.2.3 Characteristics of models
 - 1.2.4 Investment benchmarks
 - 1.2.5 Passive investment management techniques
 - 1.2.6 Active versus passive investment management
- 1.3 Motivation, aim and objectives

CHAPTER 2: LITERATURE REVIEW

- 2.1 Overview of literature
- 2.2 Empirical evaluation of portfolio-selection models
- 2.3 Estimating the dependence structure of share prices
 - 2.3.1 Full historical correlation matrix model
 - 2.3.2 Index models
 - 2.3.3 Averaging (mean) models
- 2.4 Evaluating models
 - 2.4.1 Five-year results – statistical significance
 - 2.4.2 Five-year results – economic significance
 - 2.4.3 One-year results – statistical significance
 - 2.4.4 One-year results – economic significance
- 2.5 Selection of analytical model

CHAPTER 3: METHODS

- 3.1 Selection of methods
- 3.2 Data and sample selection
- 3.3 Method of analysis
- 3.4 Limitations

CHAPTER 4: RESULTS

- 4.1 Statistical performance
 - 4.1.1 Average absolute forecast error of correlation coefficients
 - 4.1.2 Stochastic dominance of models using absolute forecast errors
- 4.2 Economic performance

CHAPTER 5: DISCUSSION

- 5.1 Statistical performance of models
- 5.2 Economic performance of models
- 5.3 General performance of models
- 5.4 Other considerations

CHAPTER 6: CONCLUSION

REFERENCES

CHAPTER 1: INTRODUCTION AND BACKGROUND

1.1 Introduction

Financial modelling forms an integral part of most modern-day financial decision-making processes. This study examines the effectiveness of different techniques in the estimation of sub-sector correlation structures, which is critical to the sector-allocation decision stage of the portfolio-construction process. As more and more institutions (both financial and other) have come to appreciate the many important applications of modelling, so too has there been an increasing emphasis placed on its importance, especially in recent years.

Financial models often make use of assumptions that allow for the simplification of reality. This simplification enables one to make sense of the relationships existing between various financial variables. The structure and intricacy of these financial models depends upon many factors, not least of which is the purpose to which the model is being put, as well as the economic significance of the model results. Different modelling techniques are adopted, depending on the requirements of the model.

One of the most important requirements of any model is that it be mathematically tractable so that it is not overly complicated and expensive to run. The lack of such tractability has long been one of the major criticisms of use of the full historical correlation matrix approach (described in detail in the next section) in the construction of efficient portfolios.

Portfolio management is a specific area of investment where financial models are of considerable value: specifically, the task of constructing share portfolios that are optimally suited to meeting investment targets. In brief terms, the ‘top-down’ approach involves adopting a structured decision-making process that starts by considering the highest level of asset allocation, i.e. between different asset classes (viz. equities, bonds, property, etc.). Once an appropriate allocation has been decided, the top-down portfolio-construction process continues by considering the split between sectors within each asset class (sector allocation) and then, finally, the split between individual shares within each of those sectors. A geographical or currency selection strategy usually overlays this process and depends on the extent to which the portfolio manager has the ability to invest in foreign assets.

The research examines explicitly the effectiveness of various modelling techniques used to estimate sub-sector correlation structures, which are then applied to the sector-allocation stage of the ‘top-down’ approach to the portfolio selection process. It was decided to focus specifically on the sector selection stage (near the top of this portfolio construction process) rather than down to the point of individual share selection, which lies at the bottom of the ‘top-down’ portfolio selection process. One of the reasons for this focus is its relative importance in overall portfolio performance given its higher position in the ‘top-down’ portfolio selection process. Another is the comparative practicality of collecting data pertaining to sector and sub-sector indices, and hence the calculation of figures necessary to determine their investment performance is simpler. Conducting an analysis based on individual share information, on the other hand, suffers from a wide range of data-induced difficulties. This is not to say, however, that the same investigation cannot be conducted for the individual-share-selection phase of the ‘top-down’ portfolio management process.

1.2 Theoretical Overview

1.2.1 Modern portfolio theory

Markowitz (1959) developed and described a method of constructing efficient portfolios for given levels of risk, provided there are estimates of the relevant parameters of returns on individual assets. The framework he provided is known as modern portfolio theory (MPT) – also referred to as mean—variance portfolio theory.

MPT forms a crucial part of the portfolio construction process. As can be seen from the brief formulation of MPT below, it involves having estimates for the means of, variances of, and pairwise covariances between returns on the universe of available assets as inputs to the optimisation process. It is clear why accurate forecasts of these parameters are important to the construction of optimal portfolios. It is therefore also apparent why the work undertaken in this research report, which focused on the estimation of correlation between sub-sectors, is also significant.

The standard MPT (Markowitz, 1959) portfolio optimisation process was used directly in this research report to assess the economic performance of the various methods used to forecast correlation coefficients between sub-sectors. The standard MPT portfolio optimisation process uses as its appropriate measure of investment risk the variance (or standard deviation) of returns.

However, portfolio optimisation can also be performed using criteria based on alternative measures of investment risk¹. These other measures of investment risk may be useful, depending on the investment objectives one is trying to fulfil. For instance, from an actuarial perspective, of particular interest might be the investment risk relative to one's liabilities – an aspect recently considered by Elton & Gruber (1992) – or another appropriate benchmark. Recent research by Frankfurter et al. (1999) and Kondor et al. (2006) examined the performance of alternative portfolio selection algorithms in the context of portfolio optimisation under various measures of investment risk.

Prior work has also been done on simplifying the portfolio optimisation process through the use of certain assumptions. The capital asset pricing model (CAPM), introduced by Sharpe (1964), is an example of a single-index model that makes some simplifying assumptions about investor behaviour in the market to model the future returns from risky assets in capital markets.

The assumptions made in the CAPM approach are as follows²:

- Capital markets are perfectly efficient.
- No arbitrage opportunities exist.
- Returns on assets are normally distributed.
- All investors have rational expectations.
- All investors have homogeneous expectations about the returns from securities for any given time period.
- Investors are only concerned with level and uncertainty of future wealth.
- Risk-free rates exist with limitless borrowing capacity and universal access. Risk-free rates for borrowing and lending are equal to each other.

The end result of the CAPM approach applied to the sector-allocation process would be the market portfolio³ - i.e. a portfolio consisting of allocation weightings being made to sub-sectors in the same way as they appear in the market, based on their market capitalisation as a percentage of the market as a whole. In other words, the CAPM would result in a passive approach to portfolio selection.

¹ 'Subject 109, Financial Economics,' *Acted Core Reading*

² Wikipedia, (http://en.wikipedia.org/wiki/Capital_asset_pricing_model), 2007

³ 'Subject 109, Financial Economics,' *Acted Core Reading*

Formulation of MPT

The brief formulation of MPT below shows how Markowitz (1959) derived a method to find the optimal security weightings to produce a portfolio with the minimum risk – in this case measured by the standard deviation of portfolio return – for a specified level of expected return.

The return on the portfolio is given by:

$$r_p = \sum_{i=1}^n x_i \cdot r_i$$

where: r_p is the return on the portfolio;

r_i is the return on security i ;

x_i is the proportion of the portfolio invested in security i ; and

n is the number of securities in the market.

The expected return on the portfolio is then:

$$E = E[r_p] = \sum_{i=1}^n x_i \cdot E_i$$

where: E is the expected return on the portfolio; and

E_i is the expected return on security i .

The variance of the portfolio return is:

$$V = Var(r_p) = \sum_{i=1}^n \sum_{j=1}^n x_i x_j \cdot C_{ij}$$

where: C_{ij} is the covariance of returns on securities i and j .

Two constraints are then imposed as part of the portfolio optimisation. Firstly, the proportion of the amounts allocated to each of the securities needs to sum up to one. The other constraint sets the expected return on the portfolio equal to a specific input return, thereby enabling the optimisation process to find the portfolio with the minimum variance for that level of required return.

Lagrangian multipliers can then be used to solve the minimisation problem and produce a system of equations, which, in turn, can be solved through the use of matrix algebra to derive the optimal weightings to each security in the efficient portfolio.⁴

1.2.2 Modelling correlation structures

The major approaches covered in the main literature (Cohen & Pogue, 1967 and Elton & Gruber, 1973) to modelling the correlation structure of future asset (or, in this case, sub-sector) returns are as follows:

- The full historical correlation matrix (HCM) model
- Single-index models – e.g. the capital-asset pricing model (CAPM)
- Multi-factor models:
 - Macroeconomic factor models
 - Fundamental factor models
 - Statistical factor models – including the single-factor version
- Mean Models – involving smoothing, i.e. averaging, of the HCM
 - Overall mean
 - Traditional mean
 - Pseudo-group mean

The research project itself does not investigate the effectiveness of implementing all the aforementioned techniques in the sector-allocation decision. It provides an overview of all models to promote discussion of the merits of those selected for study and to provide suggestions for further research into the topic. The different assumptions underlying each of the models are also described so that further comment can be passed on the relative appropriateness of each. As background to the study, a brief outline of the models follows.

Full HCM approach

This approach uses the historic pairwise covariance between each and every sub-sector to predict future performance of portfolios. As such, it is labour intensive as it requires estimates of the likely correlation between every sub-sector with every other sub-sector. This may not be an important issue when dealing with the sector-allocation stage – since there are unlikely to be many different sectors – but would be more of an issue when dealing with the individual-share-selection stage, when there are likely to be several hundred companies in a sector. With the introduction of

⁴ ‘Subject 109, Financial Economics,’ *Acted Core Reading*

computers, however, this previous criticism (made by Cohen & Pogue, 1967) of the HCM approach has become far less of an issue, yet still something to be borne in mind from the perspective of data quantities required.

Single-index models

Returns on sub-sectors can be modelled as a function of the change in a single index. Any measurable index can be used but the standard is to use the market index, which shows the returns on the entire market of risky assets. For this reason, we refer to the case where the market index is used as the standard single-index (SSI) model.

Multi-factor model approach

This approach involves modelling the returns on each sector as a linear combination of responses to a set of input factors. By introducing factors additional to market variance, the multi-factor approach attempts to improve forecasting by capturing the effects of non-market related influences⁵. These additional factors can be macroeconomic variables (rates of inflation, interest, etc.); sector specific factors seen as fundamental drivers of returns; or statistical factors that are extracted from past data using principal components analysis. The major disadvantage of multi-factor models is that the additional factors may introduce into the forecast more random noise than they do true explanatory power (Elton & Gruber, 1973).

Mean model approach

This approach involves using averaging techniques to smooth the HCM as a means of eradicating random noise, thereby allowing for superior forecasting (Elton & Gruber, 1973). The extent of the smoothing depends on which version of the mean model approach is being used. The primary drawback is that smoothing may lead to the loss of real information, i.e. information that, if discarded, would lead to poorer estimates of future correlation between sectors, thereby affecting the performance of a portfolio constructed on the basis of these poorer estimates.

For any meaningful conclusions to be reached from prevailing market conditions, the models described need to be based on the assumption that the market is in equilibrium. The assumptions necessary under each model to maintain a state of equilibrium in the market requires meticulous description. This is because knowledge

⁵ 'Subject 109, Financial Economics,' *Acted Core Reading*

of these assumptions will help formulate a platform from which to judge the relative merits and applicability of such models to practice.

This research makes use only of historical information up to the dates at which the portfolios are constructed. No attempt is made to formulate future values of individual parameters. Any input values for future variables are based purely on market expectations of future conditions consistent with the lack of arbitrage opportunities. This introduces some of the assumptions associated with strong-form market efficiency, which states that share prices reflect all information, both public and private, in a market. The implication is that investors cannot make profits in excess of normal returns, irrespective of the amount of research or information they have access to⁶.

1.2.3 Characteristics of models

A number of features are desirable in any model⁷:

Representativeness

The model should be representative of reality so that behaviour of assets and other variables under the model reasonably mimics their behaviour in the real world.

Economic interpretation

The behaviour of assets under the model should be consistent with accepted economic theory and principles. For example, asset and variable behaviour under the model should not allow for arbitrage opportunities to arise.

Parsimony

Models should be as simple as possible, while still retaining any key features necessary to predict asset behaviour. A parsimonious model, together with simplicity, makes the model easier to understand. A balance therefore needs to be struck between a model's ability to represent reality and its parsimony.

Simplicity

Closely linked to parsimony, is the need for the workings to be made as simple as possible in order for the results to be more easily understood and communicated to decision-makers, thereby facilitating the optimal course of action to be followed.

⁶ Investopedia, (<http://www.investopedia.com/terms/s/strongform.asp>), 2007

⁷ 'Subject 305, Finance & Investment,' *Acted*

Capacity for update

For models to be viable they need to be able to be updated easily. This necessitates having the capacity for development and refinement without needing to reconstruct the entire model. This would make it too costly and time-consuming to run the model as part of the decision-making process. This is particularly relevant given the trend towards dynamic financial modelling where the latest results are used as an input to building and adapting models for the future, a crucial step in the 'actuarial control cycle', a process widely recognised by the actuarial community for making decisions⁸.

Implementation tools

Good models usually have a range of methods of implementation available to facilitate their parameterisation, their testing and the focus of their results. Implementation tools may include any one, or more, of the following:

- Analytical calculations
- Historical back-testing – the process of assessing a model's effectiveness by using past data to test how it would have performed in the past. Although it has been criticised because results achieved are highly dependent on movements within the past period being tested on⁹, back-testing has been used as part of the assessment of models being compared in this research report.
- Scenario analysis – a deterministic simulation of results under future conditions under different (plausible) scenarios. This also includes sensitivity analysis, which seeks to establish the sensitivity of model results to changes in a single variable.
- Tree-building techniques – these can be used to develop scenarios under which to test the model results.
- Monte-Carlo simulation – a stochastic simulation of future results, where the value of variables in each simulation depends on a random variable from a predefined distribution. It is of particular use where outcomes are path-dependent.

⁸ 'Subject 301, Investment and Asset Management,' *Actuarial Education Company*

⁹ Investopedia, (http://en.wikipedia.org/wiki/Capital_asset_pricing_model & <http://www.investopedia.com/terms/b/backtesting.asp>), 2007

1.2.4 Investment benchmarks

Investors use benchmarks mainly as a standard against which to assess the returns on their actual portfolio, as well as a means by which to explicitly formulate and help understand their investment strategy and targets. The benchmarks that investors use to achieve these objectives are usually set to measure portfolio performance relative to any one of the following¹⁰:

Competitors' or other similar portfolios' performance

Comparing the performance of a particular strategy to that of similar or competitors' portfolios can give an "indication of the costs or benefits of that strategy relative to those of other funds". One has to remember that it is difficult to know exactly which strategies competitors are implementing, although one can usually get some idea from publicly available investment mandates or fund marketing material.

The major point to consider, though, is that these benchmarks are appropriate only if the portfolios on which they are based are similar to the portfolio studied in terms of their constraints and fund objectives, i.e. that they have been matched to any liabilities and have similar risk-tolerance levels, transaction costs (usually a result of total portfolio size), tax positions and any other factors that are likely to affect the returns on a portfolio. Sometimes other portfolios of a sufficient likeness are very difficult to come by.

Published indices

Benchmarks based on published indices are very easy to operate, provided that the data published about the underlying indices are easily available and reliable. A problem arises when there might not be an index in existence that is "consistent with the specific objectives of the investor".

Notional benchmark portfolios

Benchmarks set on this basis are generally more appropriate than those set on either of the above two bases. This is because they can be made to match the liabilities more easily, which is one of the prime objectives from an actuarial perspective. They are normally designed to be "consistent with the specific investment objectives" and constraints of the particular portfolio in question. As a result, comparison of portfolio performance relative to a notional portfolio matched to fund liabilities or constraints should give a fair reflection of relative returns.

¹⁰ 'Subject 301 & 401, Investment & Asset Management,' *Acted*

1.2.5 Passive investment management techniques

Once an appropriate benchmark portfolio has been chosen using one of the aforementioned methods, following a passive investment strategy means trying to match as closely as is (practically) possible the composition of the actual holdings to that of the benchmark portfolio¹¹. The simplest examples of passive investment management are index-tracking funds. In these cases, the benchmark portfolios are the indices they are attempting to track. The success of a particular approach to passive investment is measured by the extent of its tracking error, i.e. the extent to which the returns on the actual portfolio differ from those of the benchmark, rather than outperformance relative to the benchmark. In essence, the aim is to match its performance as closely as possible instead of actually beating it.

Assuming a passive investment management policy has been selected, there are four different basic approaches to doing so, each with its own pros and cons. In the end, however, the most suitable approach will depend on the particular circumstances of the investors as well as those of the market in which they are operating. The four different techniques, as well as some of their advantages and disadvantages, are as follows.

Full replication

As the name suggests, this method involves holding every stock in exactly the same proportions as those in which they appear in the benchmark. Clearly, one would expect this approach to result in “little or no tracking error” – its major advantage. Where a benchmark portfolio consists of a large number of different assets, the full replication approach may lead to a consistent need to readjust actual portfolio proportions back to benchmark values, which may themselves change due to the notional reinvestment of dividends or the entry and exit of various component holdings to the benchmark. This may lead to a very “fragmented portfolio” with numerous small holdings, which is, in turn, likely to result in excessive dealing expenses.

Stratified sampling

This approach involves holding in one’s actual portfolio only a subset of the different assets in the benchmark portfolio. For example, index-tracking funds might seek to

¹¹ ‘Subject 301 & 401, Investment & Asset Management,’ & ‘Subject 305, Finance & Investment,’ *Acted*

match the proportions of their overall portfolio invested in each industry to those of the benchmark index, but hold only a few of the many component benchmark stocks within each industrial sector. It is clear that by taking this approach, an element of tracking error would be introduced. Generally, the benefits of reduced dealing costs are likely to outweigh the drawbacks of any additional tracking error introduced – particularly where there are many stocks within each industry group.

Optimisation

This technique entails holding a portfolio that matches the benchmark in certain specified factors. These factors are usually recognised fundamental drivers of returns. For instance, one might choose to match the overall price-earnings (PE) ratio, degree of systematic risk (beta) or market capitalisations of the actual portfolio with those of the benchmark index whose performance one is trying to replicate. Furthermore, this does not necessarily mean restricting the assets held in the portfolio to a subset of those in the benchmark but potentially means being able to more or less track the performance of the benchmark more efficiently than either of the two methods described before. Clearly, though, the extent to which this is a success depends heavily on the effectiveness of the optimisation process, as well as the continued role of the matched factors in explaining returns.

Synthetic funds

These are “constructed using derivatives on the underlying” benchmark index (or on its component assets or indices if the benchmark is not an index on which derivatives are traded directly). Given the proclivity of derivatives to trade very close to their fair, arbitrage-free values, synthetic funds are a very time-efficient and cost-effective way to track a benchmark in the short term. They may, however, become quite expensive to operate over the long term because of the constant need to roll over derivative contracts at uncertain rates. Further difficulties would be encountered if derivatives traded on the benchmark to be tracked were unavailable.

1.2.6 Active versus passive investment management

The debate about whether active investment management yields better or worse returns in the long term compared with passive investment management is ongoing and one that receives attention frequently, e.g. see Malkiel (2003). The extent to which active management of a portfolio is likely to result in better performance depends on the level of market efficiency, or to be more precise, its inefficiency. Inefficient markets are likely to have more opportunities than other market players

may not have yet identified, therefore meaning that an active investment management strategy can add more value to the investment process and that, as a result, it can provide higher returns than a passive index-tracking strategy.

The research does consider and compare the economic results of following a passive investment strategy with those of implementing the various models but to avoid getting drawn into the debate about active versus passive investment management, on which many a paper has focused in the past, summarised below are the most commonly recognised pros and cons of passive investment management¹².

Firstly, consider the advantages of index-tracking. The costs of running the portfolio are considerably lower because there should be a smaller number and less volume of trades done than for an active investment management strategy. Volatility of passive investment returns should also be lower given the exposure to a larger number of (usually large market capitalisation) shares. This point leads on to another advantage of index-tracking, that of diversification. Depending on the index, one can gain exposure to many different companies, sectors or even geographical regions and the risk of seriously underperforming the index or competitors is reduced. Finally, as already mentioned, in highly efficient markets passive investment management is easier and should lead to better returns given the lower costs involved.

The disadvantages of passive investment include the following. Upside potential can be lost because high-growth shares are not included in the index being tracked and the potential for over-performing the index or competitors is substantially reduced. Exceptionally high returns often come from high-risk, small-capitalisation shares that are not included in most indices. Secondly, it can be difficult to find or, once found, to track an index that is appropriate to the fund objectives or the liabilities it is trying to match. There are inefficiencies in tracking certain indices, particularly those consisting of the top companies by market capitalisation, because these shares generally outperform those shares leaving the index, which therefore weight down performance of the index overall, before entering it and, vice-versa, they underperform those shares entering the index before leaving it. Some indices can also be very difficult to track, e.g. if they consist of restricted shares.

¹² 'Subject 301 & 401, Investment & Asset Management,' & 'Subject 305, Finance & Investment,' *Acted*

1.3 Motivation, aim and objectives

The main aim of this research is to determine whether there consistently emerges – using empirical evidence – an established technique that is superior to others at predicting future correlation structures between sector & sub-sector returns. If such a technique exists, then the theoretical ramifications are that it should enable its users to better select efficient sector allocations in the long run. In other words, the consequences of using such a method ought to be superior returns for those portfolios whose sector allocations are set using it, relative to those using other modelling techniques. The question whether or not the same holds true in the short run also warrants investigation, but is not something that is examined by this research paper, where a buy and hold strategy is assumed for the three-year period considered.

It is also the intention that any modelling technique that yields superior portfolio-selection results be one that is practical to apply. It should not be excessively complex and require overly many input variables. Depending on the circumstances, it may prove to be the case that a trade-off arises between a technique that is sufficiently simple to apply and one that yields satisfactory results when it comes to its capability at selecting efficient portfolios relative to other methods. Fortunately, computers and more user-friendly software make it easier to overcome most data-related, timing-related and cost-related issues of modelling.

The extent to which one is willing to depart from the one objective of the model to achieve the other depends on the level of economic significance of the decision one is faced with. Since sector allocation lies quite near the top of the top-down approach to selecting portfolios, it arguably has a large bearing on financial performance¹³ – although the extent of this effect will need to be investigated from the data analysis. One is therefore likely to favour a model that yields superior returns over one that is very simple and inexpensive to run since the additional returns expected in the long run are likely to justify the extra costs of applying such a technique. All this having been said, it may prove to be the case from the results of the research report that a less complex model provides us with more efficient portfolios – in which instance, such a trade-off will not have to be made after all.

¹³ 'Subject 301 & 401, Investment & Asset Management,' *Acted*

The research question addressed is: Which techniques to modelling correlation between different sub-sectors of the market perform best under the two assessment criteria, statistical and economic performance? The first criterion, statistical performance, evaluates the models' ability to estimate future correlation coefficients between different sub-sectors by analysing the distributions of their absolute forecast errors. The second criterion, economic performance, evaluates the performance of the efficient portfolios (constructed using MPT) resulting from the various models, assuming a three-year buy-and-hold strategy.

CHAPTER 2: LITERATURE REVIEW

This chapter reviews and briefly discusses the literature that was relevant to the topic of this research report. It begins with a broad overview of the literature covered and then focuses on some of the literature that has particular relevance to this report. Specific attention is given to previous research that covered the models investigated in this paper, as well as to their results in the past.

In order to identify the best methods of conducting the research, the literature review is based on publications that are key to that purpose. The selected methods must compare the performance of portfolios constructed via the use of these different modelling methods with those constructed using the standard capital-asset pricing model (CAPM) method (Sharpe, 1964), i.e. a single-factor model with some simplifying assumptions, where overall market variance serves as the proxy to the single source of risk assumed to contribute to sector volatility and the resulting portfolio is the market portfolio. In other words, the CAPM represents a passive investment strategy. The various modelling techniques investigated represent active investment management strategies when it comes to selecting optimal sector allocations within a portfolio – followed in the hope of choosing sector weightings that maximise returns for given levels of portfolio risk. An alternative approach would be to use a passive strategy to select sector allocations. This research covers alternative active investment management strategies, as well as covering briefly a passive investment strategy.

2.1 Overview of literature

Prior literature includes classic research conducted by Elton & Gruber (1971, 1973); Elton, Gruber & Urich (1978); and Cohen & Pogue (1967) on the subject of whether the historical correlation matrix, index models (both multi- and single-) or mean models provide the best prediction of the correlation matrix – thereby allowing for the selection of optimal efficient portfolios, i.e. portfolios that generate the highest expected returns for the levels of risk to which they are subject.

It was generally found by Elton & Gruber (1973) that, despite multi-index models' superiority over single-index models at reproducing the historical correlation matrix, they have not necessarily been better at forecasting the future correlation matrix. In other words, portfolios constructed on the basis of a single-index model generally

outperformed portfolios constructed using a more complicated multi-index approach. In addition, portfolios constructed using the even more data-intensive historical correlation matrix usually underperformed both the multi- and single-index approaches at most risk levels (as shown by Elton & Gruber (1973) and Elton, Gruber & Urich (1978)).

The results of this research are of relevance to the field of investment in many different ways. For instance, as just one example, Farrell (1974) used some of the findings of these papers to help ascertain homogeneous stock groupings for 100 US shares.

2.2 Empirical evaluation of portfolio-selection models

Cohen & Pogue (1967) described an empirical evaluation of alternative portfolio-selection models forty years ago. The main aim of their research paper was to evaluate, on an empirical basis, the ex-ante and ex-post performances of various single-period portfolio-selection models. The idea was that these models would be broadly based upon the Markowitz formulation, but that they would be simplified from the perspective of data preparation and computation.

They tested the results from four different models for both the ex-ante and the ex-post estimates for future periods. These models in decreasing order of computational complexity were the full Markowitz model, two types of multi-index models – the ‘covariance’ form and the ‘diagonal’ form of the multi-index model – and the standard single-index model developed by Sharpe (Sharpe, 1963).

In both Cohen & Pogue (1967) and Elton & Gruber (1973), the performances of the efficient portfolios produced by the various models were compared with those of randomly selected portfolios and, in all cases, were conclusively found to be superior. This result is therefore assumed without testing in this research project.

Cohen & Pogue (1967) observed yearly returns data for a set of 150 and another subset of 75 shares for the period 1947—1964, where the historical observation period ran from 1947—1957 and the period during which the performance of portfolios constructed using various modelling techniques was assessed ran from 1958—1964. As was the case in this research, they used total returns consisting of capital gains and reinvested dividends to measure the yields on shares and ignored the effects of taxation for simplicity. For the purposes of the portfolio selection, they

assumed that the best estimates of future expected returns and variances for the model assessment period were the historic mean returns and the variances of those returns during the historical observation period. They themselves acknowledged in their paper the potential shortcomings of their approach of basing expectations for the future purely on historical values but did so in an effort to remove any subjectivity in the estimates.

They found that while the ex-ante efficient frontier of the more complex Markowitz model dominated those of the simpler methods, the single-index model frontier tended to dominate those of the more complex multi-index models over a wide range of expected returns. Where, for a given level of risk, a particular model produces an efficient portfolio with greater expected returns than that of another model, the former model is said to 'dominate' the latter. They also found that the portfolios constructed using the Markowitz method tended to contain the fewest securities, while the single-index model portfolios tended to contain more securities for a given level of return than those constructed using the multi-index models.

In the case of ex-post performance, they found the picture to be less clear. They found that the more labour-intensive Markowitz approach, i.e. using the full historical correlation matrix, did not necessarily dominate these simpler models. Likewise, they found that for purely 'common stock universes' – i.e. excluding bonds, etc. – multi-index models did not outperform the simpler single-index model. Cohen & Pogue (1967) went on to suggest that, although their research pointed towards the superiority of the single index model over that of the more complex multi-index models, the "richer representation of the variance--covariance matrix permitted by the multi-index models in comparison with the single index model" may become necessary when a more diverse universe of assets is being considered.

2.3 Estimating the dependence structure of share prices

Elton & Gruber (1973) described the implications of estimating the dependence structure of share prices for portfolio selection. Although modern portfolio theory (MPT) has existed since 1952, it has rarely been implemented for individual securities mainly because of the nature of the inputs required. For Markowitz's MPT (Markowitz, 1959) to be used to produce optimum portfolios, accurate estimates are needed of mean returns, variance of returns and covariance of returns – the main challenge.

According to random-walk theory, the best estimates of future means and variances for shares are their historical ones (Osborne (1962); Kendall (1953); Mandelbrot (1966) and Fama (1965)). This assumes that distributions of returns for shares are stable over time. At the same time, analysts believe that they can form expectations for the future of return means and variances. While means and variances have received great attention, scant attention has been paid to estimating correlations. Correlations are, however, difficult to estimate since a large number of estimates are required and there is “no non-overlapping organisational structure that will allow security analysts to produce estimates of correlation coefficients between all pairs of stocks” (Elton & Gruber (1973)). It is also inconclusive whether the use of single- or multi-index models to forecast covariances offers any better forecasts than extrapolation of historical estimates.

Despite these problems, not much attention has been paid to the accuracy of techniques used to estimate share-price correlation structures. Although King (1966) considered the correlation structure of share prices within a particular period, he did not examine the stability or predictive value of correlation structures over time. Cohen & Pogue (1967) examined the predictability of several models but neither examined the accuracy of correlation projections, nor attempted to separate the errors caused by misestimating correlation coefficients from errors in misestimating means and variances of returns. Elton & Gruber (1973) focus on trying to establish which technique best estimates the correlation matrix. Only estimates based on historical data were used because of the difficulty of obtaining subjective estimates. They compared various estimation methods with respect to their ability to forecast correlation matrices and to select efficient portfolios for future periods.

Elton & Gruber (1973) used three basic types of forecasting models, each of which has different underlying assumptions that are pertinent to the selection of the model of choice.

2.3.1 Full historical correlation matrix model

The full historical correlation matrix model assumes that past values are good estimates of future correlation coefficients and thus no assumptions are made as to how or why the relationship between pairs of securities is as it is.

2.3.2 Index models

An alternative to using a direct estimate of correlation coefficients based on historical values is to assume a behavioural model of why securities move together, the parameters of which can be estimated using historical data. Index models are an example of such behavioural models and may be based on single or multiple indices.

Single-index models

The simplest of such behavioural models was developed by Sharpe (1963) following the suggestions of Markowitz (1959). The underlying assumption of the single-index model is that “securities move together only because of a common response to changes in an aggregate index” (Elton & Gruber (1973)).

The estimates of mean returns and the variance of returns produced by the Sharpe model are identical to those produced by direct estimation using historical data but the estimates of correlation coefficients are different. Under the Sharpe model, the covariance between security i and security j , assuming $i \neq j$, is given by:

$$E[(r_i - E[r_i])(r_j - E[r_j])] = \beta_{i1}\beta_{j1}\sigma_M^2 + \beta_{i1}E[e_i e_j] + \beta_{j1}E[e_i e_j] + E[e_i e_j]$$

where: r_i is the return on security i ,

β_{i1} is a measure of the responsiveness of security i to changes in the index;

e_i is a variable with mean of zero and variance σ_i^2 which measures the variability of security i that is not attributable to changes in the index; and

e_M is a variable with mean of zero and variance σ_M^2 , where σ_M^2 measures the variability of the index.

The correlation coefficient between returns on security i and security j is the above expression divided by the product of the standard deviations of the returns on security i and security j , i.e. $\sigma_i \sigma_j$. If the parameters are estimated using least-squares regression then the second and third terms on the right-hand side of the above equation equal zero by construction, “since the expected value of the residuals of a regression are independent of the value assumed by the independent variables” (Elton & Gruber (1973)). There is nothing, however, to guarantee that the fourth term

on the right-hand side of the equation equals zero. The assumption that this term equals zero (i.e. $E[e_i e_j]$ equals zero, for $i \neq j$) under the Sharpe model is the only difference between it and the full HCM model. If the behavioural model is an approximation, but not a perfect representation, of reality, there is a choice, depending on stability:

- To estimate using historical data if that part of the correlation structure not captured by the Sharpe model, $\frac{E[e_i e_j]}{\sigma_i \sigma_j}$, is stable over time;
- If that part is unstable, then assuming that $\frac{E[e_i e_j]}{\sigma_i \sigma_j}$ equals zero might lead to better forecasts of future correlation coefficients; and
- If $E[e_i e_j]$ has a stable component plus random noise, one can use a multi-index model.

Two versions of the single-index model are used by Elton & Gruber (1973). The first of these uses the S&P Industrial Index (adjusted for dividends) as its index (which they refer to as the 'SSI model' – standard single-index model). The second uses the first principal component of the historical correlation matrix as its index (which they call the 'F-1 model'). The first principal component refers to the constructed index that best explains the statistical variance in the past correlation matrix.

Multi-index models

The behavioural model underlying multi-index models assumes securities move together, partly because of economy-wide changes and partly because of an association with some subgroup in the economy, e.g. an industrial sector.

Empirical tests of the CAPM have uniformly found that the single-index model was insufficient in explaining the returns of securities. Douglas (1969) and Miller & Scholes (1972) found that returns on individual shares depended heavily on residual variation, in addition to the systematic (market) risk that the single-index CAPM captures. Work done by Black, Jensen & Scholes (1972) and Fama & MacBeth (1972) found that the explanatory power of models was improved by using beta-related factors. King (1966) also found that multiple indices seem to have real explanatory power. These findings suggest an advantage in using multi-index models.

Estimates of expected returns and variances are identical to estimates from the full historical method of estimation. The covariance between securities i and j , assuming $i \neq j$, is given by:

$$\begin{aligned} E[(r_i - E[r_i])(r_j - E[r_j])] = & \left\{ \sum_{k=1}^m \beta_{ik} \beta_{jk} \sigma_{N+k}^2 \right\} \\ & + \{ \beta_{i1} \beta_{j2} E[\varepsilon_{N+1} \varepsilon_{N+2}] + \beta_{i1} \beta_{j3} E[\varepsilon_{N+1} \varepsilon_{N+3}] \\ & + \dots + \beta_{i, N+k-1} \beta_{j, N+k} E[\varepsilon_{N+k-1} \varepsilon_{N+k}] \} \\ & + \{ \beta_{i1} E[\varepsilon_j \varepsilon_{N+1}] + \dots + \beta_{iN} E[\varepsilon_j \varepsilon_{N+1}] + \beta_{j1} E[\varepsilon_i \varepsilon_{N+1}] \\ & + \dots + \beta_{jN} E[\varepsilon_i \varepsilon_{N+1}] \} \\ & + \{ E[\varepsilon_i \varepsilon_j] \} \end{aligned}$$

where: β_{ik} is a measure of the responsiveness of security i to changes in index k ;

ε_i is a variable with mean of zero and variance σ_i^2 which measures the variability of security i that is not attributable to changes in any index; and

ε_{N+k} is a variable with mean of zero and variance σ_{N+k}^2 , where σ_{N+k}^2 measures the variability of the index k .

Again, the correlation coefficient between returns on security i and security j is the above expression divided by the product of the standard deviations of the returns on security i and security j , i.e. $\sigma_i \sigma_j$. The right-hand side of the above equation has been divided by curly brackets into four sets of terms. If indices are constructed orthogonally, the second set of terms equals zero. For the same reason as before in the case of the single-index model, for least-squares regression, the third set of terms equals zero by construction. The fourth term need not necessarily be equal to zero. The assumption that $E[\varepsilon_i \varepsilon_j]$ for $i \neq j$ equals zero is the only difference between multi-index models and full HCM model estimates.

The single-index model assumes no interaction between shares except that caused by common market movement, i.e. $E[e_i e_j]$ equals zero for $i \neq j$. The multi-index model splits $E[e_i e_j]$ into interaction due to:

- movements of subgroups, $\sum_{k=2}^m \beta_{ik} \beta_{jk} \sigma_{N+k}^2$, (i.e. the first set of terms in brackets minus its first term); and
- a residual $E[\varepsilon_i \varepsilon_j]$, i.e. the final term in curly brackets.

Multi-index models were obtained by Elton & Gruber (1973), who extracted additional indices from principal components that were ordered and orthogonal. These models have been called the F-3 model (3-factor model), the F-8 model (8-factor model) and the F-max model (eigenvalue greater than 1). The best performing model will depend on whether the historical level or zero is a better estimate of future values of the indices and the residuals. The 3-factor and 8-factor index models were chosen by Elton & Gruber (1973) because the percentage of variance explained dropped sharply with the 4th and 9th factors, i.e. the ability of the 4th and 9th factors to explain the historical correlation matrix was somewhat lower than that of the 3rd and 8th factors. The F-max model was determined by keeping in the model all those principal components that had an eigenvalue greater than 1. In the case of Elton & Gruber's (1973) sample period, this meant using 17 or 18 indices to explain variation in returns.

The ability of a model to explain the historical correlation matrix increases with each additional factor. The addition of more factors does not necessarily increase the predictive accuracy of the system though, because additional indices may merely be measuring random noise. This describes the issue of over-parameterisation¹⁴ and this is usually dealt with by means of information criteria (Akaike (1974) and Schwarz (1978)). As a result, if the influences that are attributed to these additional factors are random, they might be better estimated at zero. In other words, a model with fewer factors may result in a better forecast of the future correlation matrix than one with overly many factors.

¹⁴ Personal communication with supervisor, R.J. Thomson.

2.3.3 Averaging (mean) models

Mean models assume that historical data can be useful only to estimate the mean correlation coefficients between groups of shares and that “pairwise differences are random or sufficiently unstable so that zero is a better estimate (than their historical level) of their future value” (Elton & Gruber (1973)).

Overall mean model

The most aggregate averaging possible is to set every correlation coefficient equal to the average of all the correlation coefficients. If the assumption in the single-index model that $E[e_i e_j] = 0$ held, then the overall mean model represents a constrained form of the single-index model where every company’s correlation with the market is assumed to be exactly the same. The response of a share’s return to changes in the market is given by:

$$\beta_{i1} = \frac{\rho_{i,M} \sigma_i \sigma_M}{\sigma_M^2} = \frac{\rho_{i,M}}{\sigma_M} \sigma_i$$

where: β_{i1} is a measure of the responsiveness of security i to changes in the market index;

$\rho_{i,M}$ is the correlation coefficient of the return on security i with that of the market index;

σ_i is the variability of returns on security i ; and

σ_M is the variability of returns on the market index.

If $\rho_{i,M}$ is set equal to a constant for all shares, then β_{i1} is directly proportional to σ_i .

This model “is not inconsistent with the CAPM. Returns and variances are still affected by market moves” and “the concept of efficiency still holds”. Shares “still have different covariances with each other and the market”. In other words, this model represents “an alternative way of estimating β_{i1} where it is assumed to be proportional” to a security’s standard deviation of returns, σ_i , “rather than estimated directly from historic values” (Elton & Gruber (1973)). This may explain some empirical findings by Black et al. (1972), Douglas (1969), Fama & Macbeth (1972) and Miller & Scholes (1972) where the rate of return on a stock was found to be more closely related to its own variance, σ_i^2 , than to β_{i1} .

The primary assumption of the overall mean model that every share's correlation to the market is the same is highly unlikely to be the case, since companies or sub-sectors differ with respect to operational gearing, cyclicity, size etc. These factors are widely accepted to affect the relationship between their particular returns and those of the market as a whole. Any emerging superiority is more likely to stem from sheer coincidence rather than from its actually being a superior representation of reality. Despite this being a very naïve model, however, it can be used as a yardstick against which to measure the performance of more complex models.

Alternative mean models

An alternative approach to the overall mean model used by Elton & Gruber (1973) is a more disaggregated mean model. This involves assuming a common mean correlation coefficient for various subgroups, but that this mean can differ between subgroups. All the shares within the same homogeneous subgroup have a common correlation structure with all other shares in that subgroup. Elton & Gruber (1973) tested three different forms of alternative mean models, averaging within and between these subgroups:

- Traditional mean model - homogeneous groups were formed from the SIC (Standard Industrial Classification) code;
- Pseudo-3 model - a model containing three pseudo-industries was used because the first three principal components accounted for "such a disproportionate amount of the variance in the return data"; and
- Pseudo-7 model - a seven-pseudo-industry model was used because the original data contained seven traditional industries, i.e. the SIC had seven categories.

For the pseudo-3 and -7 models, stocks were divided into pseudo-industries by using multivariate techniques to determine which groups had behaved as homogeneous units. To do this, Elton & Gruber (1973) performed "a varimax rotation of the components" and "firms were assigned to that rotated factor on which they had the largest loading". A varimax rotation refers to a principal component analysis, "a technique used to reduce multidimensional data sets to lower dimensions for analysis"¹⁵, where the varimax criterion is used.

¹⁵ Wikipedia, (http://en.wikipedia.org/wiki/Principal_component_analysis), 2007.

2.4 Evaluating models

Two criteria are used to assess the accuracy of various forecasting techniques: statistical significance and economic significance. Statistical significance means the assessment of a technique's ability to forecast future correlation matrices. Economic significance involves the examination of their "ability to choose portfolios which prove to be efficient in future periods" (Elton & Gruber (1973)), i.e. the assessment of the ex-post performance of portfolios selected using the forecast correlation coefficients from the various models as inputs to the portfolio selection process. Elton & Gruber (1973) used both one-year and five-year estimates of the correlation matrices to test forecasting accuracy over different time spans. To test the consistency of the forecasting results, two separate, non-overlapping five-year-period forecasts and three separate, non-overlapping one-year-period forecasts were studied.

2.4.1 Five-year results – statistical significance

Elton & Gruber (1973) firstly confirmed that the historical correlation matrix contained information of worth about the individual correlations between shares. In every case, the mean absolute error for the HCM was found to be smaller than those of 30 randomly arranged correlation matrices. They then used tests of the statistical significance of the difference in absolute forecast error and comparisons of the cumulative frequency of forecast error between the different techniques.

Elton & Gruber (1973) showed that the SSI model consistently outperformed the full HCM model in the first five-year monitoring period and in the second five-year monitoring period, for both cases of the second five-year monitoring period where a five- and a ten-year observation period were used, at the 5% significance level. The SSI model is not only computationally efficient, but also produces better forecasts of future correlation matrices than the full HCM approach.

Both the SSI and F-1 models consistently outperformed the multi-index models, F-3, F-8 and F-max. The SSI model did so at the 5% significance level for all the periods considered (the first five-year monitoring period and both cases of the second five-year monitoring period). These results indicate that although the multi-index models can explain a greater percentage of the historical correlation, they are worse at predicting future correlations between shares. This means that the adding of further indices to the single-index model just results in further random noise, as opposed to predictive ability. The SSI model also resulted in a statistically significant level of greater accuracy in forecasting than the F-1 model for all the sample five-year

monitoring periods. This means that the future correlation between shares can be better estimated through the use of a broad market influence than on the main influence prevalent in past periods.

The overall mean model (which assumes correlation coefficients between all shares are the same – i.e. the mean correlation coefficient) performed better than the SSI model in both of the five-year monitoring periods. The difference in the first of these monitoring periods was statistically significant. It also resulted in greater forecasting accuracy in the second monitoring period with the longer observation period of ten years (also at a statistically significant level).

Both the pseudo-3 and traditional mean models (which allow for differences between correlation coefficients of shares) resulted in even better forecasts than the overall mean model. The pseudo-3 model performed better than the overall mean model in all monitoring periods sampled, where the difference was statistically significant for both cases of the second five-year monitoring period. The traditional mean model also outperformed the overall mean model at a statistically significant level for the first five-year monitoring period and the second five-year monitoring period with the longer observation period. However, it performed slightly worse than the overall mean model for the second five-year monitoring period when a five-year observation period was used, but not at a statistically significant level. The relative effectiveness demonstrated in Elton & Gruber's (1973) results from using traditional industry groupings in a model further supported the findings a few years before by King (1966).

Although the full historic, traditional mean, overall mean, pseudo-3 and pseudo-7 models all have the same estimate of the average correlation between stocks (the average historic correlation coefficient), this does not apply to other models investigated by Elton & Gruber (1973). They realised that various forecast methods may yield better results because they estimate better the mean level of future correlation coefficients while at the same time estimating differences in correlations between stocks more poorly. To allow for this, they adjusted the correlation matrix resulting from each technique so that all forecast the same average correlation coefficients, by reducing each individual correlation coefficient forecast by the difference between that technique's forecast mean correlation coefficient and that of the full historical correlation matrix. In the case of the five-year forecasts, the relative performance of the various techniques with the adjusted mean correlation

coefficients was very similar to their performance without adjustment, possibly because of the relative stability of five-year correlation coefficients, especially compared with one-year correlation coefficients.

2.4.2 Five-year results – economic significance

Elton & Gruber (1973) investigated the comparative abilities of the different models to aid in the process of selecting efficient portfolios by using actual ex-post values of returns and variances of returns on stocks, as well as the future correlation matrix forecast by each of the various models to select an optimum (efficient) portfolio for given levels of risk. They then monitored the relative performance of the portfolios resulting from the different models to assess whether any order of superiority emerged. Because efficient portfolios may only consist of a subset of the shares in the full correlation matrix, it was possible that the comparative performance of the various techniques at selecting efficient portfolios could vary from their relative ability to predict the future correlation matrix. The ranking of the economic performance of the various techniques' portfolios was found to be quite similar but not identical to their ranking when directly comparing forecast errors in the correlation matrix.

The same three techniques (overall mean, traditional mean and pseudo-3 models) emerged as the best performers. Their performance relative to one another, however, changed. This time, the overall mean model performed best – instead of the traditional mean model. The traditional mean and pseudo-3 models produced portfolios that performed almost identically. Nevertheless, the difference in economic returns produced by all three techniques was relatively insignificant. The Sharpe technique came in fourth position but led to significantly reduced returns (resulting in up to a 25% reduction in annual returns) compared with the first three techniques. The ranking of the four intermediate models (pseudo-7, F-1, F-8 and F-3) varied and differential economic performance was insignificant. The two techniques which clearly emerged as the worst performing were the historic and F-max models.

2.4.3 One-year results – statistical significance

Forecasting techniques have tended to outperform randomly arranged correlation matrices (Cohen & Pogue, 1967 and Elton & Gruber, 1973). In the case of one-year sample periods, however, the statistical dominance of these forecasting techniques was not nearly as great as it was in the case of the five-year periods.

Elton & Gruber (1973) observed that the forecast error was far greater for one-year forecasts than it was for five-year ones. The larger forecast errors for one-year periods stemmed in part from errors in forecasting pairwise deviations from the average correlation coefficient. The main reason for the higher overall forecast errors in one-year predictions though, was due to the larger differences between the average correlation coefficient that was forecast and that which actually materialised. The greater instability of one-year correlation coefficients compared with five-year ones meant that the main reason for greater forecasting errors was a poorer ability to estimate the mean correlation coefficient. The ranking of the various techniques largely depended on their ability to estimate the mean correlation coefficient as opposed to their ability to properly establish the structure of covariances between different shares. To allow for this, Elton & Gruber (1973) considered the relative performances of the various techniques only once they had adjusted each to produce the same average correlation coefficient. The relative performances of these mean-adjusted techniques for one-year forecasts appeared to be much the same as the rankings for five-year forecasts (both those that had and those that had not been adjusted for mean). The pseudo-3, traditional (industry) mean and pseudo-7 modelling techniques all gave better forecasts – in that order – than the overall mean model. After that came the SSI model, which outperformed the full HCM model. The results also showed that multi-index models were outperformed by single-index models.

Comparison of results of one year with five year results

When it came to comparing the results of the various forecasting techniques on the unadjusted data, it was found, as expected, that the relative performances of the models were more erratic, but also that there were far fewer cases of statistical dominance in any of their one-year monitoring periods. Elton & Gruber found no evidence that any of the models they examined could consistently demonstrate a differential ability to forecast the average correlation coefficient. They showed, however, that the relative size of the average correlation coefficient forecast by the various models for a particular period seemed to depend on the model used, regardless of whether the actual mean correlation coefficient for that period lay above or below these estimates. For instance, consistently with the results of Cohen & Pogue (1967), they found that the SSI model produced lower estimates of the average correlation coefficient than any of the other techniques for all of their sample periods.

2.4.4 One-year results – economic significance

Elton & Gruber (1973) found no detectable ordering of the techniques in respect of their ability to select efficient portfolios over a one-year forecast period. This is in stark contrast to their findings for the five-year forecast results but is not surprising given the lack of any statistically significant difference in performance when forecasting one-year correlation coefficients.

2.5 Selection of analytical model

Previous investigations have demonstrated differences in abilities of various techniques to forecast the correlation matrix. Elton & Gruber (1973) found that in the case of forecasting one-year correlation matrices, the relative differences were small enough in comparison with the forecast error that different techniques could not be consistently distinguished on either statistical or economic grounds. However, in forecasting five-year correlation matrices, some of the techniques consistently outperformed others at a level that was both statistically and economically significant. The three averaging techniques appear to be consistently superior to the other techniques investigated, including the two most conventionally employed methods of forecasting correlation matrices, the SSI and the full HCM models.

Elton & Gruber (1973) found the three top-performing models for estimating five-year correlations to be the traditional industry mean, the pseudo-3 and the overall mean models. They suggested that, because of their comparatively simple structures and superior performance, there may be simplified portfolio algorithms which lead to optimal or near-optimal portfolios.

Since then, there has been further research into the effectiveness of these correlation structure estimation techniques, as well as more recently developed ones, in the context of different portfolio selection constraints and different markets. Some research has also been done to find estimation techniques that improve upon those described above.

Ledoit & Wolf (2003) proposed an alternative method of forecasting the correlation structure of returns on shares that has come to be known as shrinkage. It involves estimating the covariance matrix of stock returns with “an optimally weighted average of two existing estimators: the sample covariance matrix (viz. HCM) and single-index covariance matrix”. They saw it as an easy way to account for extra-market

covariance, i.e. correlation between share returns over and above that explained by the market index in the SSI model, without having to resort to using an “arbitrary multi-factor structure”. The shrinkage method was not one of the methods investigated in this research report but further research into its performance could be useful.

Research has also been conducted into the effectiveness of models used to estimate correlation structures under various portfolio selection constraints. Jagannathan & Ma (2002) examined the performance of the models described above under the constraint that weightings to securities in the efficient portfolio were non-negative. They found that under this constraint, the economic performance of the portfolios constructed using the full HCM model was just as good as that of portfolios constructed using multi-index models and more recent shrinkage estimator models.

Prior research has been done into the performance of correlation structure estimation techniques in different markets. Ho & Lee (1995) examined the performance of various methods of forecasting correlation structure in the Japanese market over the period 1977—1991. Unlike Elton & Gruber (1973), they found that the full HCM and industry mean models dominated the overall mean and single-index models. They also found that the full HCM model dominated all other models when the standard MPT portfolio optimisation process was used to construct efficient portfolios. Steiner & Wallmeier (1999) compared the performance of some of the correlation estimation models above with that of multi-factor models that used firm-specific variables. They found that the multi-factor models “did not generally produce better forecasts than ‘naïve’ models”. In fact, they found that the industry mean model significantly outperformed all of the other correlation estimation models in most time periods.

There is much other prior research that also suggests multi-index models are less effective, or at least no more effective, than single-index and other more ‘naïve’ models, in the estimation of correlation structures. As already described, Cohen & Pogue (1967) found that multi-index models did not outperform the simpler single-index model for purely ‘common stock universes’. They suggested that multi-index models may be useful in the estimation of correlation structures between assets when the range of assets is broader than just equities. This research report examines the correlation structure between sub-sectors of the equity market so, based on their suggestion, multi-index models would be unlikely to outperform some of the simpler approaches to forecasting the correlation structure.

Elton & Gruber (1997) expressed the opinion that the versions of multi-index models that were most likely to be used in future were those where the indices were pre-specified observable variables rather than statistically derived factors, i.e. principal components. They did, however, see statistical estimates of factors as useful in helping to confirm that the pre-specified indices “capture all of the major influences”.

In the light of these findings, this research report did not consider the performance of multi-index models based on principal components. Nonetheless, prior literature on multi-index models has been reviewed in order to give context to the estimation models that are investigated in this report. An investigation of the comparative performance of multi-index models based on pre-specified indices or variables could be useful as part of future research into the estimation of sub-sector correlation structures.

Papers have been written over the years with respect to the performance of correlation estimation models applied to various other markets as well, e.g. international markets (Meade & Salkin (2000)) and mutual fund markets (Ahmed (2001)). Meade & Salkin (2000) investigated the performance of various models of asset returns and, more particularly, different input estimation methods. Their results suggested “that the choice of estimation method is more critical than the choice of pricing model”. The focus of this research report, which is specifically on the performance of various techniques of sub-sector correlation estimation, therefore fits in well with this suggestion. It therefore seems that the logical first step to take for research into a new area of application, is the sector-allocation phase of the top-down portfolio construction.

CHAPTER 3: METHODS

3.1 Selection of method

The selection of methods was made to address the research question: Which techniques to modelling correlation between different sub-sectors of the market perform best under the two assessment criteria, statistical and economic performance? The first criterion, statistical performance, evaluates the models' ability to estimate future correlation coefficients between different sub-sectors by analysing the distributions of their absolute forecast errors. The second criterion, economic performance, evaluates the performance of the efficient portfolios (constructed using MPT) resulting from the various models, assuming a three-year buy-and-hold strategy.

The various methods of modelling correlation were the same as some of those covered in the literature. The main difference was that in the literature the models were being applied to the estimation of correlation between shares, whereas in this research they are being applied to the estimation of correlation between sub-sectors.

In this research, the results obtained when firstly the FTSE All-share index (FTSE ALSI) was used as the single market index were compared with those when the FTSE Top 100 index (FTSE100) was used as the single market index.

3.2 Data and sample selection

Data were collected for the relevant analyses by access through the Bloomberg system. Ten years of monthly data were collected for the period 30 September 1994 to 30 September 2004. The monthly data were taken as at the end of the last day of the month or the preceding business day where this fell on a non-business day.

The data include past movements of values on the London Stock Exchange (LSE) pertinent to the investigations conducted. These data included the following information at each point in time (121 observations in all):

- values of the FTSE All-Share Total Return Index (FTSE ALSI TRI);
- values of the FTSE 100 Total Return Index (FTSE100 TRI); and
- values of LSE sector and sub-sector total return indices (further described below).

The total return indices (TRIs) were collected in sterling terms, i.e. the home currency, so as to avoid the effects of any currency movements. Given the volatility of currency markets, these effects could quite easily have overwhelmed any market risk effects. TRIs are calculated under the assumption that any dividend income from shares in the sub-sector are immediately reinvested at no cost and so this forms one of the assumptions of this research.

The sectors and sub-sectors of the LSE for which data were collected are as follows:

1. Resources:
 - Mining; and
 - Oil & Gas.
2. Basic Industrials:
 - Chemicals;
 - Construction & Building Materials;
 - Forestry & Paper; and
 - Steel & Other Metals.
3. General Industrials:
 - Aerospace & Defence;
 - Diversified Industrials (not used);
 - Electro & Electrical Equipment; and
 - Engineering & Machinery.
4. Consumer Goods (Cyclical):
 - Automobile & Parts; and
 - Household Goods & Textiles.
5. Consumer Goods (Non-Cyclical):
 - Beverage;
 - Food Production;
 - Health;
 - Personal Care & Household Products;
 - Pharmaceuticals & Biotechnology; and
 - Tobacco.
6. Services (Cyclical):
 - General Retailers;
 - Leisure & Hotels;
 - Media & Entertainment;

- Support Services; and
 - Transport.
7. Services (Non-Cyclical):
- Food & Drug Retailers; and
 - Telecom Services.
8. Utilities:
- Electricity (not used); and
 - Utilities Other (not used).
9. Information Technology:
- IT Hardware; and
 - IT Software & Technical Services.
10. Financials:
- Banks;
 - Insurance;
 - Life Assurance;
 - Investment Companies;
 - Real Estate; and
 - Speciality & Other Finance.

The 'Diversified Industrials' sub-sector was excluded from the analysis altogether because there was no TRI available after April 2003. There was also no TRI for the 'Utilities Other' sub-sector before January 2003. Due to the fact that there was only one other sub-sector in the 'Utilities' sector, the assumption was made that it was a sector with only one sub-sector and therefore the sector TRI was used as the single sub-sector's TRI.

3.3 Method of analysis

Firstly, monthly total returns were calculated for each of the market indices, the sectors and their sub-sectors from the relevant TRIs. There were 120 observations relating to monthly returns over the 10-year period for each TRI. The following formula was used to calculate the monthly returns for each sub-sector:

$$r_i(t) = \frac{Y_i(t) - Y_i(t-1)}{Y_i(t-1)}$$

where: $r_i(t)$ is the return on sub-sector i in month t ; and

$Y_i(t)$ is the value of the TRI for sub-sector i at time t .

The monthly returns on the market indices were calculated using the same formula except that, instead of using the i^{th} sub-sector TRI, the FTSE ALSI and the FTSE100 were used to calculate the monthly returns for these two indices in month t , denoted by $r_{M_A}(t)$ and $r_{M_{100}}(t)$ respectively,.

The data over the total period of the investigation were divided into three non-overlapping sub-periods of three years. As a result, only the first 108 of the 120 observations of monthly returns were used, 36 for each three-year period. Three years was chosen as a time period because asset manager performance is quite frequently assessed using rolling three-year returns¹⁶ and because it represents a compromise between the one and the five year periods that were used by Elton and Gruber (1973). It also seems a reasonable length of time to carry out a buy-and-hold strategy, which was assumed to be the case for all of the models tested in this research.

For each assessment of the various modelling techniques over a sample period to be carried out, two three-year sub-periods were required. The first sub-period was the phase over which data were analysed, i.e. the observation period, so as to serve as input to the sector allocation decision for the succeeding sub-period, i.e. the monitoring period. The monitoring period was the phase of the investigation over which the statistical and economic performance of the efficient sector allocations, selected using the different modelling techniques, was evaluated. Since there were three three-year sub-periods in all, we were able to perform an assessment of the various models over two monitoring periods.

For ease of reference, the first three-year period over which statistical and economic performance of the models was monitored, i.e. 30 September 1997 to 30 September 2000, has been referred to as 'period 1' (P1). The second three-year period over which performance was monitored, 30 September 2000 to 30 September 2003, has been referred to as 'period 2' (P2).

The observation periods over which data were gathered to construct the various models for these two periods were three years in length. In the case of P1, the observation period, which has been defined as 'period 0' (P0), ran from 30 September 1994 to 30 September 1997. The observation period for P2 was 30

¹⁶ 'Subject 301 & 401, Investment & Asset Management,' *Acted*

September 1997 to 30 September 2000. In other words, P1 was the monitoring period for the P0 observation period and it was also the observation period for the P2 monitoring period.

Similarly to Elton & Gruber (1973), the two three-year observation periods were combined to create a six-year observation period called 'period 0-1' (P0-1), from 30 September 1994 to 30 September 2000. This six-year observation period resulted in different estimates of future correlation between sub-sectors being produced by the various models. Then, exactly as was done for the second of the three-year observation periods above, P1, the statistical and economic performance of the models was monitored over the three-year period that ran from 30 September 2000 to September 2003. The results for P2 using this longer six-year observation period have been referred to as 'combined period 2' (CP2), i.e. P2 using P0 and P1 combined as the observation period.

Past data were analysed to formulate the following for each of the sub-sectors over each three-year sub-period:

- Mean monthly returns;
- Variance of mean monthly returns;
- Beta coefficients with the two market indices; and
- Historical correlation coefficients between that and the various other sub-sectors.

The mean monthly returns for sub-sector i over 'period u ' ($P(u)$), where period u is defined as above, is given by the following formula:

$$\bar{r}_{i,P(u)} = \frac{1}{b_{P(u)} - a_{P(u)}} \sum_{t=a_{P(u)}}^{b_{P(u)}} r_i(t)$$

where: u refers to the sub-period and is equal to 0, 1, 0-1 or 2;

$a_{P(u)}$ is the number of the first month in $P(u)$, so:

$$a_{P(0)} = a_{P(0-1)} = 1$$

$$a_{P(1)} = 37$$

$$a_{P(2)} = 73$$

$b_{P(u)}$ is the number of the last month in $P(u)$, so:

$$\begin{aligned} b_{P(0)} &= 36 \\ b_{P(1)} &= b_{P(0-1)} = 72 \\ b_{P(2)} &= 108 \end{aligned}$$

The sample variance of returns for sub-sector i over $P(u)$ was then calculated using the following formula:

$$\sigma_{i,P(u)}^2 = \frac{1}{(b_{P(u)} - a_{P(u)}) - 1} \sum_{t=a_{P(u)}}^{b_{P(u)}} (r_i(t) - \bar{r}_{i,P(u)})^2$$

Likewise, the mean monthly returns and their variances over the three-year sub-periods were calculated for the two market indices, the FTSE ALSI and the FTSE100.

The covariance of returns for sub-sector i and sub-sector j over $P(u)$ was then calculated using the following formula:

$$Cov_{P(u)}(r_i, r_j) = \frac{1}{(b_{P(u)} - a_{P(u)}) - 1} \sum_{t=a_{P(u)}}^{b_{P(u)}} (r_i(t) - \bar{r}_{i,P(u)})(r_j(t) - \bar{r}_{j,P(u)})$$

The correlation coefficient between returns for sub-sector i and sub-sector j over $P(u)$ can then be calculated from the covariance above by dividing through by the product of the standard deviations of returns on sub-sector i and sub-sector j over that period as follows:

$$\rho_{P(u)}(r_i, r_j) = \frac{Cov_{P(u)}(r_i, r_j)}{\sigma_{i,P(u)} \cdot \sigma_{j,P(u)}}$$

A correlation matrix for the first period (ex-post) was then constructed using the above calculated correlation coefficients derived from the observed correlations between sub-sectors over the first sub-period, P0. This is referred to as the historical correlation matrix (HCM).

This historical correlation matrix was then adjusted under the various models. The models considered were as follows:

- 1) Full historical correlation matrix (HCM) model (see Section 3.1 of Chapter 2);
- 2) Standard single-index (SSI) model – a few different versions (see Section 3.2 of Chapter 2);
- 3) Overall mean model (see Section 3.3 of Chapter 2); and
- 4) Industry mean model (see Section 3.3 of Chapter 2).

In the case of the full HCM model, the correlation coefficients between sub-sectors calculated from the observation periods were used directly as the estimates of the forecast correlation coefficients for the respective monitoring periods, i.e. the HCM derived from the observation period data was used as the forecast correlation matrix for the monitoring period. So, for the full HCM model, the estimated correlation coefficient between sub-sector i and sub-sector j , for $i \neq j$, is derived as follows:

$$\hat{\rho}_{MP}(r_i, r_j) = \rho_{OP}(r_i, r_j)$$

where: $\rho_{OP}(r_i, r_j)$ is the historical correlation coefficient between sub-sector i and sub-sector j for an observation period; and $\hat{\rho}_{MP}(r_i, r_j)$ is the forecast correlation coefficient between sub-sector i and sub-sector j for the corresponding monitoring period.

The overall mean and industry mean models both involved smoothing the HCM, to a different extent in the case of each, before a forecast correlation matrix could be obtained. The overall mean model applied the maximum possible degree of smoothing to the HCM to derive its forecast correlation matrix. It used the average of all correlation coefficients between sub-sectors in the HCM, i.e. excluding variances, as the forecast correlation coefficient between each sub-sector and every other sub-sector. This means that for the overall mean model, the estimated correlation coefficient between sub-sector i and sub-sector j , for $i \neq j$, is derived as follows:

$$\hat{\rho}_{MP}(r_i, r_j) = \frac{2}{n(n-1)} \sum_{i=1}^{n-1} \sum_{j>i}^n \rho_{OP}(r_i, r_j)$$

where: $\rho_{OP}(r_i, r_j)$ is the historical correlation coefficient between sub-sector i and sub-sector j for an observation period;
 $\hat{\rho}_{MP}(r_i, r_j)$ is the forecast correlation coefficient between sub-sector i and sub-sector j for the corresponding monitoring period; and
 n is the total number of sub-sectors in the market, i.e. $n = 33$ in the case of this research.

The industry mean model applied smoothing to the HCM to a lesser extent in that it only smoothed sub-sector correlation coefficients over the industry sector to which that particular sub-sector belongs, rather than having used the average of all observed correlation coefficients between different sub-sectors across all sectors. Algebraically, the estimated correlation coefficient between sub-sector i and sub-sector j , for $i \neq j$, under the industry mean model can be derived using one of the formulae below, dependent on whether sub-sectors i and j are in the same sector, or different sectors.

For sub-sector i and sub-sector j both in the same industry sector:

$$\hat{\rho}_{MP}(r_i, r_j) = \frac{2}{(d_i - c_i + 1)(d_j - c_j)} \sum_{i=c_i}^{d_i-1} \sum_{j>i}^{d_j} \rho_{OP}(r_i, r_j)$$

where: $\rho_{OP}(r_i, r_j)$ is the historical correlation coefficient between sub-sector i and sub-sector j for an observation period;
 $\hat{\rho}_{MP}(r_i, r_j)$ is the forecast correlation coefficient between sub-sector i and sub-sector j for the corresponding monitoring period;
 c_i is the number of the first sub-sector in the industry sector to which sub-sector i belongs; and
 d_i is the number of the last sub-sector in the industry sector to which sub-sector i belongs.

For sub-sector i and sub-sector j in different industry sectors:

$$\hat{\rho}_{MP}(r_i, r_j) = \frac{1}{(d_i - c_i + 1)(d_j - c_j + 1)} \sum_{i=c_i}^{d_i} \sum_{j=c_j}^{d_j} \rho_{OP}(r_i, r_j)$$

where: $\rho_{OP}(r_i, r_j)$ is the historical correlation coefficient between sub-sector i and sub-sector j for an observation period;

$\hat{\rho}_{MP}(r_i, r_j)$ is the forecast correlation coefficient between sub-sector i and sub-sector j for the corresponding monitoring period;

c_i is the number of the first sub-sector in the industry sector to which sub-sector i belongs; and

d_i is the number of the last sub-sector in the industry sector to which sub-sector i belongs.

For the purposes of the SSI models, beta coefficients were also derived for each sub-sector with respect to the two market indices. The beta for sub-sector i with respect to the FTSE ALSI over $P(u)$ was calculated using the following formula:

$$\beta_{i, M_A}(P(u)) = \frac{Cov_{P(u)}(r_i, r_{M_A})}{\sigma_{M_A, P(u)}^2}$$

Similarly, the beta for sub-sector i with respect to the FTSE100 over $P(u)$ is calculated using the following formula:

$$\beta_{i, M_{100}}(P(u)) = \frac{Cov_{P(u)}(r_i, r_{M_{100}})}{\sigma_{M_{100}, P(u)}^2}$$

In the SSI model, correlation coefficients between sub-sectors were forecast for monitoring periods through the use of the historical betas calculated from the corresponding observation periods. The various versions of the SSI model are followed through below only for the case where the FTSE ALSI is used as the single index, but exactly the same was done for the SSI models where the FTSE100 was used. For the unadjusted version of the SSI model, the estimated correlation coefficient between sub-sector i and sub-sector j , for $i \neq j$ was derived as follows:

$$\hat{\rho}_{MP}(r_i, r_j) = \frac{\beta_{i, M_A}(OP) \cdot \beta_{j, M_A}(OP) \cdot \sigma_{M_A, OP}^2}{\sigma_{i, OP} \cdot \sigma_{j, OP}}$$

where: OP is the observation period; and
 MP is the corresponding monitoring period.

Given that the unadjusted SSI model produces a different average forecast correlation coefficient to the other models, alternative versions of the SSI model were also investigated. The unadjusted correlation matrices of the SSI models (using both the FTSE ALSI and the FTSE100 indices as the proxy to the market index) were adjusted so that the average of all correlation coefficients between sub-sectors was equal to the average of the correlation coefficients in the HCM. This was done in two different ways, on a multiplicative and on an additive basis. The additive method is the same as that used by Elton & Gruber (1973), where each correlation coefficient off the main diagonal has added to it the difference between the average of the non-diagonal correlation coefficients in the HCM and the average of those in the unadjusted SSI model's correlation matrix. The multiplicative method is similar, except that instead of adding the difference, each non-diagonal correlation coefficient is scaled up by the factor derived from the average non-diagonal correlation coefficient of the HCM divided by that of the unadjusted SSI model's correlation matrix. In both cases, the result is that the average correlation coefficient of the adjusted matrices is equal to that of the HCM. As pointed out by Elton & Gruber (1973), this eliminates the potential advantage a model may appear to have if its average correlation coefficient ends up being closer to the average correlation between sub-sectors that actually transpires during the model evaluation period, i.e. the period for which a forecast is made.

The statistical performance of the models was gauged by analysing the cumulative distribution of the absolute forecast errors of the correlation coefficients. The absolute forecast error, $FE_{MP}(r_i, r_j)$, for the correlation coefficient between sub-sector i and sub-sector j , where $i \neq j$, was calculated as follows:

$$FE_{MP}(r_i, r_j) = \left| \rho_{MP}(r_i, r_j) - \hat{\rho}_{MP}(r_i, r_j) \right|$$

From these, the average absolute forecast error, AFE_{MP} , was calculated for each of the different models for the different monitoring periods by the following formula:

$$AFE_{MP} = \frac{2}{n(n-1)} \sum_{i=1}^{n-1} \sum_{j>i}^n FE_{MP}(r_i, r_j)$$

The purpose of this research was to investigate the effectiveness of various approaches to modelling sub-sector correlation structures. Means and standard deviations of returns on sub-sectors were therefore assumed to be perfect forecasts in an attempt to isolate the effects of correlation estimates produced by the different models. Once the correlation coefficients had been forecast for the respective monitoring periods, they were used, along with the means and variances of sub-sector and market returns that actually transpired during the relevant monitoring periods, to construct the covariance matrices implied by the correlation structure produced by the various models.

$$\hat{\sigma}_{ij,MP} = \hat{\rho}_{MP}(r_i, r_j) \cdot \sigma_{i,MP} \cdot \sigma_{j,MP}$$

where: $\hat{\sigma}_{ij,MP}$ is the estimated covariance between sub-sector i and sub-sector j for the monitoring period, MP ; and
 $\sigma_{i,MP}$ is the standard deviation of returns on sub-sector i that is actually observed during the monitoring period.

The covariance matrices were needed as inputs to modern portfolio theory (MPT), introduced by Markowitz (1959), to enable the construction of efficient portfolios. After the various projection models mentioned above were developed from the analysis of past information, MPT was used to select the optimal portfolios for given levels of expected return on the portfolio. In other words, the allocations to each sector were determined using MPT. This entailed finding the combination of weightings to each sub-sector that minimises the expected variance of portfolio returns – the risk-exposure proxy – given different levels of expected monthly return on the portfolio (see Section 2.1 of Chapter 1).

In the case of the CAPM, these sector weightings would be the same as those of the market portfolio. From that perspective, it represents a passive investment strategy. The economic performance of the portfolios constructed using the various models

was therefore also compared with the performance of the respective market portfolios, the FTSE ALSI and the FTSE100. The market portfolios were assumed to be the same as those in the two market indices and consequently were established directly from the Bloomberg data.

3.4 Limitations

As already outlined, this research focused on the comparison of these methods of portfolio selection applied only to the sub-sector allocation stage of the ‘top-down’ process. In terms of the overall process of top-down portfolio construction, it could be said that the assumption that was made was that the overall performance of a group of stocks held within a specific sector would follow that of the sector’s index. It is possible that one of the approaches to modelling future dependence structures between sectors could have emerged as being consistently superior over the entire period under investigation, or over sub-periods of it. Were such a modelling technique to be identified as statistically superior relative to others, it would go a long way towards solving the problem of maximising an LSE portfolio’s efficiency. This would, however, be dependent on how well the particular model continued to describe the relationships between different sub-sectors. Therefore, one limitation is that a model that has performed best in the past may not continue to best describe the relationship in the future, in which case it would lose its predictive value.

The quantitative results of the analyses performed are presented in Chapter 4 and include tables and figures to assist in the interpretation of pertinent information that arises from the results.

CHAPTER 4: RESULTS

Microsoft Excel was used to generate the results for the various models. This chapter covers the performance of the various models over the different monitoring periods, specifically with respect to the two assessment criteria identified in the research question. Some brief commentary has been added to assist with further discussion in Chapter 5.

4.1 Statistical performance

4.1.1 Average absolute forecast error of correlation coefficients

Table 1 shows the average absolute error of the forecast correlation coefficients between sub-sectors produced by each of the models tested compared with those that actually materialised in the following three-year periods. Section 3 of the previous chapter described how the average absolute forecast errors were calculated for each of the monitoring periods.

Table 1 – Average absolute forecast error of correlation coefficients

Period 1		Period 2		Combined Period 2	
1 Industry Mean	0.1746	1 Industry Mean	0.2091	1 Industry Mean	0.2097
2 Full HCM	0.1845	2 SSI ALSI (Additive)	0.2193	2 Full HCM	0.2153
3 SSI ALSI (Additive)	0.1864	3 SSI FTSE100 (Additive)	0.2216	3 SSI ALSI (Additive)	0.2221
4 SSI FTSE100 (Additive)	0.1888	4 Full HCM	0.2246	4 SSI FTSE100 (Additive)	0.2246
5 SSI ALSI (Multiplicative)	0.1981	5 SSI ALSI (Multiplicative)	0.2275	5 SSI ALSI (Multiplicative)	0.2273
6 Overall Mean	0.2001	6 Overall Mean	0.2284	6 Overall Mean	0.2312
7 SSI FTSE100 (Multiplicative)	0.2103	7 SSI FTSE100 (Multiplicative)	0.2341	7 SSI FTSE100 (Multiplicative)	0.2363
8 SSI ALSI (Unadjusted)	0.2434	8 SSI ALSI (Unadjusted)	0.3676	8 SSI ALSI (Unadjusted)	0.3660
9 SSI FTSE100 (Unadjusted)	0.2445	9 SSI FTSE100 (Unadjusted)	0.3679	9 SSI FTSE100 (Unadjusted)	0.3661

The results show that the industry mean model clearly emerges as the best performing of the modelling methods in all the periods tested. This result is consistent with Elton & Gruber's (1973) findings when they tested the same model against others. The only difference is that that they were examining correlation structures between shares, whereas we are examining those between sub-sectors.

The full HCM model also performs well. It performed second best in P1, fourth best in P2 and second best in CP2.

The additively adjusted versions of the SSI models performed relatively well. The SSI FTSE ALSI additively adjusted model ranked third in P1, second in P2 and third in

CP2. The SSI FTSE100 additively adjusted model placed fourth, third and fourth in the same respective periods.

The performance of the multiplicatively adjusted SSI FTSE ALSI and FTSE100 models was mediocre. They ranked fifth and seventh respectively in P1, P2 and CP2. The relatively naive, and probably oversimplified, overall mean model ranked between these two models for all sample periods, coming in sixth place.

The unadjusted SSI FTSE ALSI and SSI FTSE100 models consistently performed poorly. The unadjusted SSI FTSE ALSI model ranked eighth and the unadjusted SSI FTSE100 model ranked last in all sample periods. As the table indicates, the average absolute forecast error for these last two models was substantially higher than for the other models tested, particularly in P2 & CP2.

4.1.2 Stochastic dominance of models using absolute forecast errors

For the purposes of comparing the distributions of the absolute forecast errors produced by the alternative estimation models, a non-standard definition of stochastic dominance was used. In this paper, the term 'relative stochastic dominance' is defined as the situation where the cumulative distribution function of the absolute forecast errors of a particular model is generally, but not always, above that of another model. Here, if one says that model A exhibited relative stochastic dominance over model B, this describes the situation where, when the ordered absolute forecast errors produced by the models were compared with each other, the number of times that the ordered absolute errors produced by model A were smaller than those produced by model B exceeded the number of times that the ordered absolute errors of model A were larger than those of model B.

A simple test was developed to test for the significance of the results for relative stochastic dominance. On the simplistic basis that, if model A and model B were exactly equivalent in terms of their ability to estimate correlation coefficients, one would expect model A's cumulative distribution of absolute forecast errors to dominate that of model B's for 50% of the observations of ordered absolute forecast errors.

This meant that, under the null hypothesis that the cumulative distribution functions for the models were equivalent, the number of times model A would have dominated model B was distributed binomially. In fact, the distribution of the number of times the

ordered absolute forecast errors of model A were smaller than their model B equivalents, X , was distributed $Bin(n, p = 0.5)$ where n was the number of observations of absolute forecast error, i.e. the number of pairwise correlation coefficients that were estimated.

Given that there were 33 sub-sectors, i.e. $k = 33$ in the formula below, the number of pairwise correlation coefficients was calculated as follows:

$$n = \frac{1}{2}.k.(k - 1) = \frac{1}{2}.33.32 = 528$$

where: n is the number of estimates made of pairwise correlation coefficients between sub-sectors; and
 k is the number of sub-sectors.

Now that the appropriate null hypothesis distribution for X has been formulated, i.e. $X \sim Bin(n = 528, p = 0.5)$, one can find its expected value and variance under the null hypothesis:

$$E[X] = np = 528.(0.5) = 264$$

$$V(X) = np(1 - p) = 528.(0.5)^2 = 132$$

$$\Rightarrow S.D.(X) = \sqrt{V(X)} = \sqrt{132} = 11.489$$

Under the central limit theorem, the distribution of X tends to a normal distribution with the same parameters. This enabled the critical value of X that corresponded to a 95% level of confidence, i.e. $\alpha = 5\%$ significance level, that model A dominated model B to be found as follows:

$$\Phi^{-1}(1 - \alpha) = \Phi^{-1}(0.95) = 1.6449$$

where: $\Phi^{-1}(z)$ is the inverse cumulative distribution function for the standard normal variable Z , i.e. $Z \sim N(0,1)$;
 α is the level of significance of the one-tailed test.

For a one-tailed test, as this was (since we were considering whether model A dominated model B), the critical value of X was calculated as follows:

$$\begin{aligned} C_X(1-\alpha) &= E[X] + \Phi^{-1}(1-\alpha).(S.D.(X)) \\ C_X(0.95) &= 264 + (1.6449).(11.489) = 282.898 \\ \Rightarrow C_X(0.95) &= 283 \end{aligned}$$

where: $C_X(1-\alpha)$ is the critical value of X corresponding to the $100.(1-\alpha)^{th}$ percentile under the null hypothesis distribution;
 $\Phi^{-1}(z)$ is the inverse cumulative distribution function for the standard normal variable Z , i.e. $Z \sim N(0,1)$;
 α is the level of significance of the one-tailed test.

The corresponding 95% confidence level critical value for the proportion of observations where the ordered absolute forecast errors of a model were smaller than those of another was therefore 53.6% (283 divided by 528). Therefore, to be able to say that model A dominated model B, the proportion of model A's ordered absolute forecast errors that had to be smaller than those of model B had to be 53.6% or above.

It should be pointed out that this analysis of significance is relatively basic in that it took into account for each model those ordered absolute forecast errors that were smaller than those of another model but it did not factor in the actual size of these differences. It might be useful for any further research to also allow for the size of these differences in ordered absolute forecast errors.

Table 2 shows the ranking of the various models in each of the sample periods in terms of their relative stochastic dominance versus the other models tested. Where the relative stochastic dominance of one model over another was not statistically significant at the 95% confidence level, square brackets have been used to show the groupings of models where this was the case. Graphs 1, 2 and 3 of the cumulative distribution functions of absolute forecast errors of correlation coefficient estimates in each of the sample periods also demonstrate the relative stochastic dominance of the various models when compared with each other.

Table 2 – Relative stochastic dominance of correlation coefficients estimates

Period 1		Period 2	Combined Period 2
1 Industry Mean		1 Industry Mean	1 Industry Mean
2 Full HCM		2 Full HCM	2 Full HCM
3 SSI ALSI (Additive)		3 SSI ALSI (Additive)	3 SSI ALSI (Additive)
4 SSI ALSI (Multiplicative)		4 SSI FTSE100 (Additive)	4 SSI FTSE100 (Additive)
5 SSI FTSE100 (Additive)		5 SSI ALSI (Multiplicative)	5 SSI ALSI (Multiplicative)
6 Overall Mean		6 Overall Mean	6 Overall Mean
7 SSI FTSE100 (Multiplicative)		7 SSI FTSE100 (Multiplicative)	7 SSI FTSE100 (Multiplicative)
8 SSI ALSI (Unadjusted)		8 SSI ALSI (Unadjusted)	8 SSI FTSE100 (Unadjusted)
9 SSI FTSE100 (Unadjusted)		9 SSI FTSE100 (Unadjusted)	9 SSI ALSI (Unadjusted)


As appeared to be the case when the average absolute forecast errors were examined, the superiority of the industry mean model in estimating correlation coefficients is confirmed by its relative stochastic dominance over the other modelling methods. It consistently ranks first in P1, P2 and CP2.

Tables 3, 4 and 5 show the percentage of observations where the difference between the ordered absolute forecast errors of correlation coefficients of the model at the top of a column and those of the model at the left of a row were greater than zero. In other words, if the value is close to 100%, the model at the left of that row exhibits greater relative stochastic dominance over the model at the top of that column than if the value is close to 50%. Table 3 shows the relative stochastic dominance of the various models for P1.

Table 3 – Relative stochastic dominance in first three-year period (P1)

	SSI FTSE100 (Unadjusted)	SSI ALSI (Unadjusted)	SSI FTSE100 (Multiplicative)	Overall Mean	SSI FTSE100 (Additive)	SSI ALSI (Multiplicative)	SSI ALSI (Additive)	Full HCM	Industry Mean
1 SSI FTSE100 (Unadjusted)	56%								
2 SSI ALSI (Unadjusted)	90%	88%							
3 SSI FTSE100 (Multiplicative)	82%	82%	51%						
4 Overall Mean	95%	96%	84%	91%					
5 SSI FTSE100 (Additive)	100%	98%	86%	84%	52%				
6 SSI ALSI (Multiplicative)	99%	97%	90%	92%	75%	64%			
7 SSI ALSI (Additive)	100%	100%	99%	95%	74%	93%	70%		
8 Full HCM	100%	98%	95%	94%	95%	94%	89%	73%	
9 Industry Mean	100%	100%	100%	100%	100%	100%	100%	100%	100%

Graph 1 – Cumulative distribution of absolute forecast errors in period 1 (P1)




The full HCM model consistently performs second best of the methods. This is as opposed to the results for the average absolute forecast errors, where it ranked fourth in P2. That having been said, its relative stochastic dominance over the additively adjusted SSI FTSE ALSI and FTSE100 models is very marginal for P2. As can be seen from Table 4, in P2 the ordered absolute forecast errors of the full HCM model are smaller than those of the additively adjusted SSI FTSE ALSI and SSI FTSE100 models in only 51% and 54% of cases respectively. When the observation period, however, was lengthened to six years in CP2, the relative stochastic dominance of the full HCM model over these other two models was more convincing.

Table 4 shows the relative stochastic dominance of the various models for P2.

Table 4 – Relative stochastic dominance in second three-year period (P2)

	SSI FTSE100 (Unadjusted)	SSI ALSI (Unadjusted)	SSI FTSE100 (Multiplicative)	Overall Mean	SSI ALSI (Multiplicative)	SSI FTSE100 (Additive)	SSI ALSI (Additive)	Full HCM	Industry Mean
1 SSI FTSE100 (Unadjusted)									
2 SSI ALSI (Unadjusted)	50%								
3 SSI FTSE100 (Multiplicative)	100%	99%							
4 Overall Mean	100%	100%	54%						
5 SSI ALSI (Multiplicative)	100%	100%	69%	54%					
6 SSI FTSE100 (Additive)	100%	100%	55%	63%	53%				
7 SSI ALSI (Additive)	100%	100%	58%	63%	56%	71%			
8 Full HCM	100%	100%	81%	63%	56%	54%	51%		
9 Industry Mean	100%	100%	91%	85%	98%	78%	79%	74%	
Max	100%	100%	100%	100%	100%	100%	100%	100%	100%

Graph 2 – Cumulative distribution of absolute forecast errors in period 2 (P2)



The additively adjusted SSI FTSE ALSI model consistently performed reasonably well. It ranked third in terms of relative stochastic dominance of absolute forecast errors in P1, P2 and CP2. The statistical performance of the additively adjusted SSI FTSE100 and the multiplicatively adjusted SSI FTSE ALSI models was very similar. The multiplicatively adjusted SSI FTSE ALSI model marginally outperformed the additively adjusted SSI FTSE100 model in P1 and the reverse was true in P2 and CP2.


Table 5 shows the relative stochastic dominance of the various models for CP2.

Table 5 – Relative stochastic dominance in combined period 2 (CP2)

	SSI ALSI (Unadjusted)	SSI FTSE100 (Unadjusted)	SSI FTSE100 (Multiplicative)	Overall Mean	SSI ALSI (Multiplicative)	SSI FTSE100 (Additive)	SSI ALSI (Additive)	Full HCM	Industry Mean
1 SSI ALSI (Unadjusted)									
2 SSI FTSE100 (Unadjusted)	60%								
3 SSI FTSE100 (Multiplicative)	100%	100%							
4 Overall Mean	100%	100%	54%						
5 SSI ALSI (Multiplicative)	100%	100%	79%	56%					
6 SSI FTSE100 (Additive)	100%	100%	54%	63%	53%				
7 SSI ALSI (Additive)	100%	100%	57%	66%	55%	73%			
8 Full HCM	100%	100%	80%	69%	78%	60%	57%		
9 Industry Mean	100%	100%	85%	75%	84%	77%	69%	74%	
Max	100%	100%	100%	100%	100%	100%	100%	100%	100%

The statistical performance of the overall mean model was mediocre. It ranked sixth in all three sample periods and marginally outperformed the multiplicatively adjusted SSI FTSE100 model, which placed seventh in all three cases.

Graph 3 – Cumulative distribution of absolute forecast errors in combined second three-year period (CP2)



As was the case with the average absolute forecast error of the correlation coefficients, the relative stochastic dominance of the unadjusted SSI FTSE ALSI and FTSE100 models clearly demonstrated how poorly they performed in comparison with the other models. The unadjusted SSI FTSE ALSI model ranked second last and the SSI FTSE100 model came last in P1 and P2. Although in P2 exactly 50% of the ordered absolute forecast errors of the unadjusted SSI FTSE ALSI model were smaller than those of the unadjusted SSI FTSE100 model, i.e. the two models appeared equal from the perspective of relative stochastic dominance, the average absolute forecast error of the FTSE ALSI model was smaller and hence judged to have performed marginally better in that period. The unadjusted SSI FTSE100 model marginally outperformed its FTSE ALSI-based counterpart in CP2.

Generally speaking, although some models seem to consistently outperform others, some of the outperformance in estimating future correlation coefficients between sub-sectors is not massive, at least compared to the drastic underperformance of the two unadjusted single-index models. This is particularly apparent from graphs 1, 2 and 3 of the cumulative distribution functions of the absolute errors in forecast correlation coefficients. Some of the models consistently produce smaller errors than others in their forecasts but there definitely appears a distinct difference in all three cases between the grouping of lines for seven of the modelling techniques and those for the two unadjusted single-index models, which clearly lie far out to the right on their own.

4.2 Economic performance

The second criterion by which the forecasting ability of the different models was assessed was their economic performance. Exactly as was done by Elton & Gruber (1973), the correlation coefficients forecast by the various models using an observation period, along with the actual means and variances of sub-sector returns that materialised during the relevant monitoring period, were used as inputs to a portfolio selection model.

Once a portfolio had been selected for each of the models, their economic performance over the three-year monitoring period was compared, assuming no rebalancing of the portfolio over that period. In other words, a three-year buy and hold strategy was used and, through the use of total return indices, all dividend income from a sub-sector was assumed to have been reinvested in that same sub-sector as it was distributed. For the purposes of comparing model performance over the three-year monitoring periods, it was assumed that a total of £1 billion was

available for investment at the time of portfolio selection. At the end of the monitoring period, the accumulated values of the portfolios were compared with each other to assess which models led to better investment returns.

The investment performance of the portfolios selected using each of the models was also compared with that of the portfolio that would have resulted by following a passive investment strategy, i.e. a portfolio that consisted of the sub-sectors in the exactly the same proportions as they contributed to the market index. For this purpose, it was assumed that the portfolio representing a passive investment strategy with respect to a particular market would perform exactly the same as that market's total return index.

This research also assumed that short sales of sub-sectors was possible, i.e. that there were no constraints in terms of shorting the various sub-sector indices. Likewise, it was assumed that there was no limit to the exposure possible in any one sub-sector.

Tables 6 and 7 show the economic performance of the portfolios selected using the various models. They both show the accumulated value at the end of the monitoring periods of a £1 billion notional investment made at the beginning of those periods.

Table 6 shows the performance of the portfolios that were constructed using the actual ex-post return on the FTSE ALSI during the respective monitoring periods as the target return input to the portfolio optimisation. It also shows the economic performance of each model relative to the FTSE ALSI.

Table 6 – Economic performance using FTSE ALSI return as the target return

Period 1 (P1)				Period 2 (P2)				Combined Period 2 (CP2)			
Rank	Model	Accumulated Value	Outperformance of Passive Approach Over Period	Rank	Model	Accumulated Value	Outperformance of Passive Approach Over Period	Rank	Model	Accumulated Value	Outperformance of Passive Approach Over Period
1	Full HCM	1,789,731,876	34.96%	1	Industry Mean	825,486,882	12.71%	1	Full HCM	847,941,502	15.78%
2	SSI FTSE100 (Multiplicative)	1,372,547,017	3.50%	2	SSI FTSE100 (Multiplicative)	810,280,780	10.63%	2	SSI FTSE100 (Multiplicative)	816,473,891	11.48%
3	SSI FTSE100 (Unadjusted)	1,362,295,940	2.72%	3	SSI ALSI (Multiplicative)	803,724,503	9.74%	3	Industry Mean	809,174,405	10.48%
4	Industry Mean	1,352,596,036	1.99%	4	Full HCM	802,879,635	9.62%	4	SSI ALSI (Multiplicative)	808,631,956	10.41%
5	SSI ALSI (Unadjusted)	1,349,758,085	1.78%	5	SSI FTSE100 (Unadjusted)	777,156,598	6.11%	5	SSI FTSE100 (Unadjusted)	779,394,513	6.42%
6	Passive Investment Strategy	1,326,161,255	0.00%	6	SSI ALSI (Unadjusted)	773,674,042	5.64%	6	SSI ALSI (Unadjusted)	774,901,950	5.80%
7	SSI ALSI (Multiplicative)	1,323,479,418	-0.20%	7	SSI FTSE100 (Additive)	760,523,737	3.84%	7	SSI FTSE100 (Additive)	763,852,715	4.29%
8	SSI FTSE100 (Additive)	1,314,151,586	-0.91%	8	SSI ALSI (Additive)	756,357,126	3.27%	8	SSI ALSI (Additive)	758,522,063	3.57%
9	SSI ALSI (Additive)	1,298,528,641	-2.08%	9	Overall Mean	732,755,079	0.05%	9	Passive Investment Strategy	732,401,694	0.00%
10	Overall Mean	1,279,012,259	-3.56%	10	Passive Investment Strategy	732,401,694	0.00%	10	Overall Mean	731,141,913	-0.17%

Table 7 shows the performance of portfolios constructed using the actual ex-post return on the FTSE100 as the target return input to the portfolio optimisation, as well as the economic performance relative to that of the FTSE100.

Table 7 – Economic performance using FTSE100 return as the target return

Period 1 (P1)				Period 2 (P2)				Combined Period 2 (CP2)			
		Outperformance of Passive Approach Over Period				Outperformance of Passive Approach Over Period				Outperformance of Passive Approach Over Period	
Rank	Model	Accumulated Value	Period	Rank	Model	Accumulated Value	Period	Rank	Model	Accumulated Value	Period
1	Full HCM	1,759,875,064	36.69%	1	Industry Mean	794,383,471	11.67%	1	Full HCM	819,505,357	15.20%
2	SSI FTSE100 (Multiplicative)	1,335,099,317	3.69%	2	SSI FTSE100 (Multiplicative)	781,515,379	9.86%	2	SSI FTSE100 (Multiplicative)	787,439,781	10.69%
3	SSI FTSE100 (Unadjusted)	1,322,092,169	2.68%	3	SSI ALSI (Multiplicative)	774,789,726	8.91%	3	SSI ALSI (Multiplicative)	779,479,194	9.57%
4	Industry Mean	1,313,577,172	2.02%	4	Full HCM	771,553,854	8.46%	4	Industry Mean	778,239,226	9.40%
5	SSI ALSI (Unadjusted)	1,309,478,955	1.70%	5	SSI FTSE100 (Unadjusted)	748,633,918	5.24%	5	SSI FTSE100 (Unadjusted)	750,838,945	5.55%
6	Passive Investment Strategy	1,287,538,959	0.00%	6	SSI ALSI (Unadjusted)	745,075,480	4.74%	6	SSI ALSI (Unadjusted)	746,264,937	4.90%
7	SSI ALSI (Multiplicative)	1,284,943,436	-0.20%	7	SSI FTSE100 (Additive)	731,942,098	2.89%	7	SSI FTSE100 (Additive)	735,230,228	3.35%
8	SSI FTSE100 (Additive)	1,274,614,658	-1.00%	8	SSI ALSI (Additive)	727,689,365	2.29%	8	SSI ALSI (Additive)	729,806,695	2.59%
9	SSI ALSI (Additive)	1,258,921,986	-2.22%	9	Passive Investment Strategy	711,376,095	0.00%	9	Passive Investment Strategy	711,376,095	0.00%
10	Overall Mean	1,238,239,746	-3.83%	10	Overall Mean	702,872,447	-1.05%	10	Overall Mean	702,173,387	-1.29%

In most cases, the ranking of the models' economic performance was the same in both the cases where the ex-post FTSE ALSI total returns and the ones where the ex-post FTSE100 total returns were used as the target return inputs to the portfolio selection process. There were a few exceptions to this, which will be commented on, but these had a minimal effect on the general ranking of model performance and so, for that reason, the results for the FTSE ALSI total return comparison in Table 6 were used as the basis for comparison between the various models, with occasional reference made to Table 7 where the ranking differed.

The results show that the full HCM model performed best overall. It ranked first, fourth and first in P1, P2 and CP2 respectively and comfortably outperformed the passive investment strategy in all periods. Its investment performance over P1 comfortably exceeded that of all the other models. Compared with the performance of the FTSE ALSI itself, it resulted in almost 35% greater returns over the three-year monitoring period, equating to an outperformance of roughly 10.5% per year. The next best model in P1 resulted in an outperformance of the FTSE ALSI of only 3.5% over that time, i.e. roughly 1.15% per year. The full HCM performed reasonably well in P2. It only ranked fourth but its level of performance was not much lower than the top three performing models. Over the period, it resulted in 9.6% higher returns than the FTSE ALSI (i.e. 3.1% per year outperformance), whereas the best performing model led to only 12.7%, or 4.1% per year, outperformance of the FTSE ALSI. In CP2, which used the same monitoring period as P2 but with a six-year observation period (P0-1) instead of a three-year one (P1), the full HCM performed best. It

resulted in 15.8% higher returns than the FTSE ALSI, c. 5.0% per year outperformance.

The economic performance of the multiplicatively adjusted SSI models was mixed. The multiplicatively adjusted SSI FTSE100 model performed consistently well. Its economic performance ranked second in P1, P2 and CP2 and it outperformed the FTSE ALSI in all three periods. On the other hand, the multiplicatively adjusted SSI FTSE ALSI model's performance was more erratic and also worse than its FTSE100 equivalent. It ranked seventh in P1, third in P2, fourth in CP2 when the FTSE ALSI ex-post return was used as the target return input in the portfolio optimisation and third in CP2 when the FTSE100 ex-post return was used. It marginally underperformed the passive investment approach in P1 and outperformed it in P2 and CP2.

The industry mean model performed comparatively well in all three periods, although its ranking did fluctuate to a degree. It ranked fourth in P1, first in P2 and third in CP2. Actually, it ranked fourth in CP2 when the FTSE100 ex-post return was used as the target return input but this did little to change the overall impression of its performance relative to the other models. It also outperformed the passive investment approach in all cases. In P2, its outperformance was 12.7%, which was equivalent to just over 4% per year outperformance of the FTSE ALSI. Its outperformance of the FTSE ALSI over CP2 was slightly lower at 10.5%, or 3.4% per year, and over P1 was decidedly lower at 2%, i.e. 0.7% per year.

The economic performance of the unadjusted SSI models was mediocre. The unadjusted SSI FTSE100 model ranked third in P1 and fifth in P2 and CP2. It also outperformed the passive investment strategy in all three periods, but not to the same extent as some of the other models. Likewise, the unadjusted SSI FTSE ALSI model also outperformed the passive investment approach in all three periods, although not quite to the same degree as the unadjusted SSI FTSE100 model. It ranked fifth in P1 and sixth in P2 and CP2.

The economic performance of the additively adjusted SSI FTSE ALSI and FTSE100 models was below average compared with the other active investment strategy models. However, their performance appeared quite similar to, if anything slightly better than, that of the passive investment strategy. The SSI FTSE100 additive model just beat its FTSE ALSI equivalent in all three periods. It placed eighth in P1

and seventh in P2 and CP2, whereas the SSI FTSE ALSI additive model placed ninth in P1 and eighth in P2 and CP2. Compared with the passive investment strategy, they both underperformed very slightly in P1 and both outperformed slightly in P2 and CP2.

The worst economic performance can categorically be said to have come from the overall mean model. It performed worst in P1 and CP2 and second worst in P2 when the FTSE ALSI ex-post return was used as the target return input. Only the passive investment strategy performed worse than the overall mean model in P2, and only by the smallest of margins. In P2, it only outperformed the passive investment strategy by 0.05% over the entire three-year period. In P1, it led to 3.6% lower returns than the passive investment strategy. Its underperformance of the index in CP2 was less severe, where it resulted in only 0.2% lower returns over those three years. When the FTSE100 ex-post return was used as the input to target return, the overall mean model ranked last in all periods and consistently underperformed the FTSE100 index by at least 1% over the three-year periods investigated.

Generally speaking, most of the models investigated performed better than the passive investment strategy. The only model that appeared to underperform it over the sample periods was the overall mean model. As mentioned previously, despite the naïveté of the overall mean model, it serves as a useful benchmark against which to compare the performance of the other more complex models. From the results of the economic performance of the various models, it appears that the relative complexity of the other models was justifiable since they led to significantly better economic returns.

When the historical observation period for the second period was lengthened to six years in CP2, the economic performance of most of the models investigated improved compared with their performance over the same three-year monitoring period when only a three-year observation period was used in P2. The only exceptions to this were the two smoothing models. Both the industry mean and the overall mean models performed worse in CP2 than they did in P2.

For the models that performed better in CP2 than in P2, i.e. those that produced better economic performance when a six-year versus a three-year observation period was used, the difference in their economic performance in each case over the same three-year monitoring period was, for the most part, noticeable but, at the same time,

not massive. Except in the case of the full HCM model, the difference in outperformance between using a six-year and a three-year observation period was under 1% over the entire three-year monitoring period. The difference in the full HCM model's outperformance of the FTSE ALSI when a six-year compared with a three-year observation period was used was more than 6% over three years.

For the two smoothing models, which both performed worse when a longer observation period was used, the difference in performance for the various lengths of observation period was varied. On the one hand, the difference in performance of the industry mean model over P2 and CP2 was more than 2% over the three-year monitoring period. On the other, the difference in performance of the overall mean model over P2 and CP2 was less than 0.25% over the three-year monitoring period.

It would therefore seem that, although the length of observation period made only a minor difference to the economic performance of most of the models tested, it had a somewhat more substantial effect on the full HCM and the industry mean models, which were also the two models with the top statistical performance. In the case of the full HCM, a longer observation period led to better economic performance. In the case of the industry mean model, it led to poorer economic performance.

In the context of the focus of this research having been on the estimation of correlation coefficients between sub-sectors rather than between individual shares, it seems reasonable that the results might differ slightly from those of some prior research.

CHAPTER 5: DISCUSSION

Elton & Gruber (1973) expressed their reticence to come to any conclusion regarding the relative performance of correlation estimation methods when it came to one-year forecasts based on the size of the errors (which were far larger) resulting from each of the estimation techniques, and there not being any statistically significant differences in their relative performance.

Their results showed that for five-year periods some of the methods tested were statistically superior. They found that, in some cases, choosing the incorrect technique could cost an investor as much as 50% of returns produced by the best of these techniques. What follows below is a discussion of the results of the research into the performance of the various models when applied to sub-sector selection, as opposed to individual share selection as carried out by Elton & Gruber (1973) and Cohen & Pogue (1967).

5.1 Statistical performance of models

The industry-mean model consistently performed best at estimating future correlation coefficients between sub-sectors of the different methods that were tested. This result agrees with Elton & Gruber's (1973) findings which were that the smoothing of the correlation matrix so that correlation coefficients between firms were the same for all firms within the same traditional industries gave the best estimates of future correlation coefficients. This result is also in keeping with what one may intuitively expect, as the correlation between one sub-sector and another from a different sector is likely to be very similar to that between two other sub-sectors, each from the same sectors as the original two sub-sectors.

The full HCM model also performed fairly well. This is contrary to the results of Elton & Gruber's (1973) research. They found the full HCM model to be the worst performing of the models they tested. The difference in results could potentially be due to the fact that their research was done in the context of individual share selection, whereas this paper focuses on sub-sector allocation. Over time one would expect the relationships between different sub-sectors to be more stable than the relationships between different individual shares. In this sense, it is perhaps not surprising that the full HCM model performs well and indeed better, when it comes to estimating sub-sector correlation, than many of the other models that brought Elton &

Gruber (1973) greater success in estimating the correlation structure between shares.

Single-index models based on the FTSE All-share index (FTSE ALSI) generally outperformed the equivalent single-index models that used the FTSE Top 100 index (FTSE100) as the market proxy. This is probably as one expects, in that some of the sub-sectors, whose correlation with other sub-sectors we need to forecast as accurately as possible to construct efficient portfolios, only contain smaller companies that might not be included in the hundred largest companies by market capitalisation. As a result, one may expect the performance of companies from certain sectors, for example the financials and resources sectors, to dominate the performance of the FTSE100 and, for that reason, it may not be as appropriate as the FTSE ALSI as the index on which to forecast correlation between certain smaller-cap sub-sectors. In any case, although the FTSE100 is updated more frequently than the FTSE ALSI and is often used as a basis for derivative products, both indices are easily available and, given that we are considering model performance from the perspective of a three-year buy-and-hold strategy instead of a short-term trading one, the less frequent daily FTSE ALSI quote is of very little consequence.

That having been said, although the relative outperformance of the FTSE ALSI models over their FTSE100 equivalents in terms of forecast error in the correlation coefficients is consistent, it is overshadowed by a more significant effect. This is the effect of the adjustment method used to bring the average correlation coefficient of the SSI model's correlation matrix in line with that of the historical correlation matrix, i.e. whether correlation coefficients are left unadjusted, adjusted additively or adjusted multiplicatively. This is clear from the fact that the difference in rank and average absolute forecast error between SSI models that used the same index and different adjustment methods, tended to be larger than the difference in rank and average absolute forecast error between SSI models that used the same adjustment method and a different index.

As far as the statistical performance of the various SSI models in forecasting correlation coefficients goes, those that were adjusted to have the same average correlation coefficient as the historical correlation matrix far outperformed the ones that were left unadjusted in terms of estimating future correlation coefficients. In fact, during the sample periods, both versions of the unadjusted SSI models drastically

underperformed all the other models tested. This is very clear from graphs 1, 2 and 3 showing the cumulative frequency of absolute forecast errors.

The additively adjusted SSI models consistently outperformed those that were multiplicatively adjusted in the sample periods. The additively adjusted SSI model is the same one used by Elton & Gruber (1973) and one that they found consistently outperformed the full HCM model in the estimation of correlation coefficients between individual shares. This was certainly not the case in the sample periods observed in this research paper, which examined the effectiveness of various models in estimating correlation coefficients between sub-sectors. This is probably for the reason mentioned above that suggests why the full HCM model performed well in selecting sub-sectors relative to the other models tested.

5.2 Economic performance of models

Naturally, one might expect there to be some degree of tracking error in the performance of a portfolio held to follow the market exactly, caused by transaction costs, timing differences, tax, etc.. For the purposes of this research though, these differences were ignored and a perfect, albeit relatively naive, passive investment strategy was used for comparison. Further investigation into the effects of these various costs and a comparison of the effectiveness of the different passive investment methods used in practice (that were covered in Section 2.5 of Chapter 1) could form the basis of future research.

The assumption was made in this research that short sales of sub-sectors were possible and that exposure to individual sub-sectors was unlimited. To achieve negative or very high exposures to particular sub-sectors, one would probably have to use derivatives. In reality, there may be a limited extent to which this is possible. There are often restrictions on asset managers to prevent them from holding unhedged short positions, perhaps as a result of many high profile risk management failures caused by derivative misuse. Such restrictions may be imposed either by fund mandates or even by regulation. In this case it would be necessary to build constraints into the portfolio optimisation. From a practical perspective, however, this research ignored these constraints and it was assumed that negative exposure to sub-sectors could be achieved without limit. Similar to the research done by Jagannathan & Ma (2002), further research into portfolio constraints such as non-negative sub-sector weightings may be useful. Further research into the practical

issues associated with using derivatives to obtain the targeted sub-sector exposures could also prove to be valuable.

Market liquidity, or rather lack thereof, may also result in self-imposed constraints by asset managers on the levels of investment in particular sub-sectors, particularly for the smaller sub-sectors. For example, if a £1billion investment were made in the UK banking sub-sector, it would probably have very little effect on the banking sub-sector index. If the same sized investment were made in the household goods and textiles sub-sector though, it could move the price of that sub-sector to the extent that it no longer gave the same returns as it would have done. This is driven by the extent to which there is supply in the market to meet the level of demand. This was not something that was taken into account in this research but it may be an aspect that may also warrant further consideration.

5.3 General performance of models

Taking into account both the statistical and economic performance of the various models, it is immediately obvious that their ranking was not the same. As was the case for Elton & Gruber (1973), the ranking of the economic performance of the portfolios produced by the various models was found, to some extent, to be similar, but certainly not identical, to their ranking when directly comparing forecast errors in the correlation matrix, i.e. their statistical performance.

The reason for this is that they reflect different assessment criteria. The statistical performance reflects the ability of models to produce accurate estimates of the future correlation matrix between sub-sectors., whereas the economic performance reflects the ability of models to aid in sub-sector selection as part of the efficient portfolio construction process.

The statistical performance of the various models from one sample period to the next was more consistent than their economic performance. The less consistent economic performance of models could partially have been due to the cyclical nature of markets. Some models may inherently result in portfolios that perform better in rising markets than those from other models, whereas others may produce portfolios that fare better in falling markets.

Overall, the full HCM model with a six-year observation period performed best of the models investigated. This is in stark contrast to the relatively poor performance of the

full HCM model in Elton & Gruber's (1973) results. The especially strong performance by the full HCM in this research compared with that in Elton & Gruber's is probably because the link between historical and future correlation of sub-sectors is much stronger than it is between that of individual shares. One expects that there is a far more stable relationship between sub-sectors over time than there is between individual shares.

The results of this research resemble closely those of Ho & Lee (1995) who examined the performance of various methods of forecasting correlation structures in the Japanese market over the period 1977—1991. They too found that the full HCM and industry mean models dominated the overall mean and single-index models. Furthermore, they also found that the full HCM model dominated all other models when the standard MPT portfolio optimisation process was used to construct efficient portfolios.

The models investigated represent active investment strategies and generally outperformed the passive investment strategy. The general pros and cons of active versus passive investment were outlined in the theoretical overview (in Section 2.6 of Chapter 1). From an investor's perspective, the additional return resulting from some of the models versus that of the passive investment strategy may be insufficient to cover the higher charges that active investment managers would likely charge.

5.4 Other considerations

This research has focused on the ability of the different models to estimate future correlation between sub-sectors in the context of perfect estimates for expected returns and their variance. The relative performance of models may have differed if, for instance, past returns and their variances were used as best estimates of expectations for their future values. Given that estimates of future expected returns and variances of returns are entirely subjective, the assumption of perfect knowledge in this respect is clearly naive but nonetheless necessary for the sake of simplicity. It is also needed in order to isolate the effect of the various models' ability to estimate future correlation structures between sub-sectors.

The 'Diversified Industrials' & 'Utilities Other' sub-sectors did not have all the necessary data points and so were ignored altogether. This is not ideal as they may have had some effect, albeit quite minor, on returns in the market.

This research, like that of Elton & Gruber's (1973), used rates of return to formulate the various models. There is a very strong argument to be made for using forces of return rather than rates of return, as was made by Thomson (1996). Forces of return are additive, whereas rates of return are not. Ideally speaking, averaging of an additive variable would be far more meaningful than averaging of a non-additive variable. An investigation into the performance of the various models using forces of return would seem to be the next logical step in terms of further research.

The effect of the length of observation period on the performance of each of the models was varied and was assessed by comparing the performance of the models in P2, which used a three-year observation period, versus that in CP2, which used a six-year observation period. With respect to the statistical performance of the models, the longer observation period resulted in a much lower average absolute forecast error only in the case of the full HCM model. It also led to marginally lower forecast errors for the unadjusted SSI models and the multiplicatively adjusted SSI FTSE ALSI model. The majority of the models, however, had higher average absolute forecast errors when the longer observation period was used. It is therefore not clear as to whether a longer observation period aids or hinders the estimation of future correlation coefficients. From the perspective of economic performance, a lengthened observation period generally resulted in superior performance from all models, except for the two smoothing models. This suggests that using a six-year observation period might be better in most cases. Further research to help ascertain the optimal length of observation period for the various models would be useful, but would require a far larger dataset than the one used in this paper.

This research assumed a buy-and-hold strategy, where the holding period was three years. In reality, this is unlikely to be the practice of asset managers, who generally rebalance their portfolios far more frequently. If one were to try and allow for more frequent rebalancing, the modelling required would need to be dynamic and therefore far more sophisticated. Elton & Gruber (1973) analysed the performance of models using one-year and five-year holding periods. Likewise, the effectiveness of the various models using holding periods of different lengths could be investigated further.

In terms of building up a portfolio, it could potentially be difficult to gain exposure to specific sub-sectors by taking positions in individual shares. This is because quite often companies may be involved in more than one sub-sector. Some, such as multi-

national conglomerates, might even have operations in more than one sector and can have very different business mixes which makes them difficult to classify by industry. This research ignores this issue but in reality it could make portfolio construction more difficult if exposure were not to be achieved synthetically, i.e. using derivatives.

The different models should also be compared across investment portfolios with different levels of risk tolerance. This is because, although a particular method may prove to be superior at selecting portfolios with risk equivalent to that of the market as a whole (i.e. $\beta = 1$), the same may not hold for portfolios with either higher ($\beta > 1$) or lower ($\beta < 1$) risk exposure than the market. Different modelling techniques may perform better at different risk levels. As a result, further research into the same analysis conducted at various risk levels would be useful.

Occasionally there may be fundamental shifts in the market paradigm and the way it perceives and prices risk and correlations between sub-sectors. For example, these shifts can be the result of invention and innovation or regulatory changes within industries. This can lead to a change in expected returns and variance of returns for sub-sectors, as well as the underlying relationships between sub-sectors, i.e. correlation coefficients between sub-sectors. For the purposes of this research, this possibility has been ignored but one should make allowance for the likely effects of any potential shifts in the prevailing paradigm when sub-sector performance is being forecast.

One should also bear in mind that the results of this research merely represent the performance of the various models over the sample periods. For there to be more robust conclusions drawn, a greater number of periods would need to be examined as part of any further research.

CHAPTER 6: CONCLUSION

The results of this research suggest there is some merit in using certain of the models tested to forecast future correlation between sub-sectors, whether it be from the perspective of the statistical or the economic performance criterion. The ranking of the models differed to some degree between these two assessment criteria. That having been said, two models in particular, the full HCM model and the industry mean model, stood out from the rest in terms of their performance.

From the perspective of the statistical performance criterion, which aimed to measure the models' ability to estimate future correlation coefficients between different sub-sectors by analysing the distributions of their absolute forecast errors, the industry mean model consistently performed the best of the models investigated. The full HCM model also performed well in comparison with the other models and, generally speaking, could be said to have performed second best.

With respect to economic performance, all the models, with the exception of the overall mean model, outperformed the passive investment strategy of holding the market portfolio. The economic performance of the full HCM model was best overall and that of the industry mean model was also strong relative to the other models tested.

The strong performance from the industry mean model reiterates the findings of Elton & Gruber (1973) when they applied various models to the estimation of future correlation between individual shares rather than sub-sectors. All things considered though, the strongest overall performance in this research came from the full HCM model. This is in contrast to Elton & Gruber's (1973) findings, where the full HCM model performed poorly. One expects the link between historical and future correlation of sub-sectors to be much stronger than it is between that of individual shares. This explains why a far more stable relationship between sub-sectors than between individual shares over time results in the exceptionally strong performance by the full HCM model in this research.

Single-index models are widely used because of the simplicity of their inputs and practical application. Although the performance in this research of the various SSI models tested was not as strong as that of the full HCM and the industry mean models, they can still be useful for those reasons. The effects of varying two key

aspects of the SSI model on its performance were investigated. These two aspects were, firstly, the method used to adjust the average correlation coefficient to be equal to that of the HCM and, secondly, which index was used as the market proxy.

Of these two aspects, the one that had more of an impact on the SSI model's performance was the adjustment method applied to correlation coefficients between sub-sectors so that the average correlation coefficient was equal to that of the HCM. The two adjustment methods tested were where the correlation coefficients were adjusted additively and multiplicatively. The statistical performance of the additively adjusted SSI model was better than the multiplicatively adjusted SSI model and the ranking of the models was reversed in respect of their economic performance.

The performance of the unadjusted SSI model relative to the two adjusted versions was mixed but, in most cases, the adjusted SSI models performed better than the unadjusted SSI model. The statistical performance of the unadjusted SSI model was by far the worst of all the models tested. The economic performance of the unadjusted SSI model was worse than the multiplicatively adjusted but better than the additively adjusted SSI model. The unadjusted SSI model's performance, however, is not directly comparable with any of the other models as its relative performance will depend to a large degree on whether the average correlation between sub-sectors that materialises during the forecast period is closer to the average correlation it produces or the average correlation that prevailed during the observation period. One cannot predict with any certainty which of the two is a better estimate of future correlation between sub-sectors and so it makes more sense to use one of the adjusted SSI models.

Which index was used as the market proxy in the SSI models had less of an effect than the adjustment method. Aside from the practical considerations of whether to use the FTSE ALSI or the FTSE100 (which were discussed briefly in Section 1 of Chapter 5), it was not entirely clear from the results which one led to better performance. The statistical performance of the SSI model was better when the FTSE ALSI was used as the market proxy, whereas its economic performance improved when the FTSE100 was used.

The length of observation period used had an effect on the performance of the various models. This was assessed by comparing the results of the same models using three-year and six-year observation periods. The longer observation period had

a mixed effect on performance. Purely in the context of the top two performing models, it improved both the statistical and economic performance of the full HCM model whereas it reduced the performance of the industry mean model under both assessment criteria.

There are several aspects that have been identified as potentially requiring further investigation. Some of those identified were the effects of using forces of return rather than rates of return in the construction of models, the determination of an optimal length of observation period and the effectiveness of the models when a more dynamic portfolio construction approach is used.

In conclusion, while prior research has established that the industry-mean model is useful in forecasting future correlation between individual shares, this research found that the industry-mean model also has value in forecasting future correlation between sub-sectors. The most important conclusion of this research, however, is that, despite demonstration in some prior research of the full HCM model's poor ability to estimate future correlation between individual shares, it certainly cannot be ignored and is indeed one of the most effective models when it comes to forecasting correlation between sub-sectors. Both of these models are therefore likely to lead to more efficient portfolios for given levels of risk and, as a result, benefit investors in the sub-sector allocation stage of the top-down approach to financial portfolio construction.

REFERENCES

- Ahmed, P. (Jun. 2001), 'Forecasting Correlation among Equity Mutual Funds,' *Journal of Banking and Finance*; 25; 6: 1187-1208.
- Akaike, H. (Dec. 1974), 'A New Look at the Statistical Model Identification,' *IEEE Transactions on Automatic Control*; 19; 6: 716-723.
- Black, F., Jensen, M. & Scholes, M. (1972), 'The Capital Asset Pricing Model: Some Empirical Tests,' *Studies in the Theory of Capital Markets*, Praeger Press, New York.
- Cohen, K.J. & Pogue, J.A. (Apr. 1967), 'An Empirical Evaluation of Alternative Portfolio-Selection Models,' *The Journal of Business*; 40; 2: 166-193.
- Douglas, G.W. (1969), 'Risk in the Equity Markets: An Empirical Appraisal of Market Efficiency,' *Yale Economic Essays*; 9: 3-45.
- Elton, E.J. & Gruber, M.J. (Oct. 1971), 'Improved Forecasting Through the Design of Homogeneous Groups,' *The Journal of Business*; 44; 4: 432-450.
- Elton, E.J. & Gruber, M.J. (Dec. 1973), 'Estimating the Dependence Structure of Share Prices – Implications for Portfolio Selection,' *The Journal of Finance*; 28; 5: 1203-1232.
- Elton, E.J., Gruber, M.J. & Ulrich, T.J. (Dec. 1978), 'Are Betas Best?' *The Journal of Finance*; 33; 5: 1375-1384.
- Elton, E.J. & Gruber, M.J. (1992), 'Optimal Investment Strategies with Investor Liabilities,' *Journal of Banking & Finance*; 16: 98-112.
- Elton, E.J. & Gruber, M.J. (1997), 'Modern Portfolio Theory, 1950 to Date,' *Journal of Banking & Finance*; 21: 1743-1759.
- Fama, E. (Jan. 1965), 'The Behaviour of Stock Market Prices,' *The Journal of Business*; 38; 1: 34-105.

Fama, E. & Macbeth, J. (Jul. 1972), 'Risk, Return and Equilibrium: Empirical Tests,' University of Chicago.

Farrell, J.L. (Apr. 1974), 'Analysing Covariation of Returns to Determine Homogeneous Stock Groupings,' *The Journal of Business*; 47; 2: 186-207.

Frankfurter, G.M., Phillips, H.E. & Faulk, G. (Dec. 1999), 'The Ex Post Performance of Four Portfolio Selection Algorithms,' *Review of Quantitative Finance and Accounting*; 13; 4: 347-366.

Ho, R.Y.K. & Lee, R.S.K. (1995), 'Correlation Structure Forecasting & Ex Ante Portfolio Selection Strategies in the Japan Market,' *Asia-Pacific Financial Markets*; 2; 1: 1-14.

Jagannathan, R. & Ma, T. (Jan. 2002), 'Risk Reduction in Large Portfolios: A Role for Portfolio Weight Constraints,' AFA 2002 Atlanta Meetings.

Kendall, M.G. (1953), 'The Analysis of Economic Time Series,' *Journal of the Royal Statistical Society*; 116; 1: 11-25.

King, B. (Jan. 1966), 'Market and Industry Factors in Stock Price Behaviour,' *The Journal of Business, Supplement*; 139-190.

Kondor, I., Pafka, S., Karadi, R. & Nagy, G. (2006), 'Portfolio Selection in a Noisy Environment Using Absolute Deviation as a Risk Measure,' *Practical Fruits of Econophysics*, Springer, Tokyo; 220-225.

Ledoit, O. & Wolf, M. (Dec. 2003), 'Improved Estimation of the Covariance Matrix of Stock Returns with an Application to Portfolio Selection,' *Journal of Empirical Finance*; 10; 5: 603-621.

Malkiel, B.G. (2003), 'Passive Investment Strategies and Efficient Markets,' *European Financial Management*; 9; 1: 1-10.

Mandelbrot, B. (Jan. 1966), 'Forecasts of Future Prices, Unbiased Markets, and "Martingale" Models,' *The Journal of Business*; 39; 1: 242-255.

Markowitz, H. (1959), *Portfolio Selection: Efficient Diversification of Investment*, John Wiley and Sons, Inc., New York.

Meade, N. & Salkin, G.R. (Sep. 2000), 'The Selection of Multinational Equity Portfolios: Forecasting Models and Estimation Risk,' *The European Journal of Finance*; 6; 3: 259-279.

Miller, M.H. & Scholes, M. (1972), 'Rates of Return in Relation to Risk: A Re-examination of Some Recent Findings,' *Studies in the Theory of Capital Markets*, Praeger Press, New York.

Osborne, M.F.M. (May-Jun. 1962), 'Periodic Structure in the Brownian Motion of Stock Prices,' *Operations Research*; 10; 3: 345-379.

Schwarz, G. (Mar. 1978), 'Estimating the Dimension of a Model,' *The Annals of Statistics*; 6; 2: 461-464.

Sharpe, W.F. (Jan. 1963), 'A Simplified Model for Portfolio Analysis,' *Management Science*; 9; 2: 277-293.

Sharpe, W.F. (1964), 'Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk,' *Journal of Finance*; 19; 3: 425-442.

Steiner, M. & Wallmeier, M. (Mar. 1999), 'Forecasting the Correlation Structure of German Stock Returns: A Test of Firm-Specific Factor Models,' *European Financial Management*; 5; 1: 85-102.

Thomson, R.J. (1996), 'Stochastic Investment Modelling: The Case of South Africa,' *British Actuarial Journal*; 2; 3: 765-801.