Nottingham University Business School

MSc Finance and Investment

Masters Dissertation

Financial Algorithms for Dynamic Investment Reallocation

Alfranso K. Lindsey

Nottingham University Business School
University of Nottingham, UK
Email: lixal2@nottingham.ac.uk

Word Count: 19968

Abstract

This dissertation provides the fundamentals of dynamic investment reallocation and their applications to trading investment practitioners' decisions in the financial services sector. Our objective is to explain the concepts and techniques that can be applied to real-world dynamic investment decision making. While empirical findings on dynamic reallocation in investment theory are currently limited, we take a bold step towards a new approach to allocation in respect to trading investments. We develop a dynamic reallocation model to explore its implications for allocators. We present our findings in a way which poses questions as to the importance of further research in this field. We explore the findings of both simulated and published data to derive important conclusions.

Acknowledgements

I would like to thank:

Dr. Robert Young, my dissertation supervisor and mentor, for his guidance, encouragement and support throughout my academic career at Nottingham University. I would not have made it through without his patience and belief in me.

Dr. Peter Blanchfield, my undergraduate supervisor, for his encouragement, enthusiasm and for motivating me towards my chosen academic goals.

Teresa Bee, my student support officer and academic counsellor, for supporting me unconditionally in the difficult times of my academic career at Nottingham University.

Rebecca Maguire, for her uncontested belief in my abilities (from day one), and her support to contribute towards my goal of reaching my full potential.

The following Charitable Trusts: Snowdon Award Scheme, Houston Charitable Trust, Nottingham Gordon Memorial Trust and Professional Classes Aid Council for their financial support.

All my friends, colleagues and non-academic staff members who help brighten up my life over the years at Nottingham University.

The Catering Staff, especially, for their open-hearted kindness and friendship throughout my course of study at Nottingham.

Most importantly, my family, for their everlasting belief in me and their patience during my long absence from home.

Contents

1 Introduction 9

1.1 Motivation 10

1.2 Research Methodology 11

1.3 Outline of the Dissertation 13

2 Literature Review 14

2.1 Portfolio Allocation 14

2.1.1 Mental Accounting 15

2.1.1.1 Individual Stock Accounting 16

2.1.1.2 Portfolio Accounting 20

2.1.2 Home Bias 22

2.1.3 Diversification – Demographic Variables 22

2.2 Dynamic Decision Making 23

2.2.1 Characteristics of Dynamic Decision Making Environments 24

2.2.2 Dynamic Decision Making in Micro-worlds 25

2.2.3 Dynamic Decision Making in the Real-world 26

2.2.4 Learning theories in Dynamic Decision Making 27

2.2.4.1 Strategy-Based Learning Theory 27

2.2.4.2 Connectionist Learning Theory 27

2.2.4.3 Instance-based Learning Theory 28

2.2.5 Individual Differences in Dynamic Decision Making 29

2.3 Literature Conclusions 29

3 Data and Methodology 31
3.1 Research Questions 31
3.2 Synopsis of the Approach 32
3.3 Data Collection 33
3.3.1 Duration of Collecting Data 33
3.3.2 Sample of CTAs 33
3.4 Methodology Research Design 34
3.4.1 Theoretical Basis Model 35
3.4.1.1 CTA Response – Four Cases 37
3.4.2 Methodological Framework Implications 41
3.4.3 Rule Construction 42
3.4.3.1 Rule Definition 43
3.4.4 Simulation Development 50
3.4.4.1 Monte Carlo Methods 50
3.4.4.2 Simulation Model Implementation 51
4 Rules of Allocation 53
4.1 Constant Allocation Rule 53
4.2 Allocation Rule 1 53
4.3 Allocation Rule 2 54
4.4 Allocation Rule 3 54
4.5 Allocation Rule 4 54
4.6 Allocation Rule 5 55
4.7 Allocation Rule 6 55

4.8 Allocation Rule 7 ... 56

4.9 Allocation Rule 8 ... 56

4.10 Allocation Rule 9 ... 56

4.11 Allocation Rule 10 ... 57

4.12 Allocation Rule 11 ... 57

4.13 Allocation Rule 12 ... 58

5 Rule Categorisation ... 59

5.1 Augmented Allocation Rule 3 ... 60

5.2 Augmented Allocation Rule 7 ... 60

5.3 Augmented Allocation Rule 9 ... 60

5.4 Augmented Allocation Rule 12 ... 60

6 Weight Allocations ... 62

6.1 Weight Generation ... 62

6.2 Optimal Weighting using Composite Allocation Algorithm ... 63

6.3 Calculating Optimal Weights using Alternative Population Means ... 64

7 Incorporating Serial Correlation ... 67

7.1 Calculating Optimal Weights with Serial Correlation ... 68

8 Empirical Analysis of CTA Performance ... 71

8.1 Calculating Descriptive Statistics for CTA Data Set ... 72

8.2 Rule Categorisation – Incorporating Serial Correlation ... 73

9 Conclusions ... 77

9.1 Future Investigation Outline ... 80

Appendicies ... 82

References..82

List of Tables

Table 1: List of all CTAs in the Data Set 34

Table 2: Summary of Four Cases 41

Table 3: Comparison of Results of Allocations Rules with Augmented Parameters 59

Table 4: Weight Computation for Composite Allocation 62

Table 5: Weighting for the Composite Allocation – 10 Simulations 63

Table 6: Weighting for the Composite Allocation for Different variations of *n* 64

Table 7: Comparisons of Optimal Weight Allocation with Different Population Means 65

Table 8: Optimal Weight Allocation with Full Range of Population Means 65

Table 9: Comparisons of Optimal Weight Allocation with Serial Correlation 68

Table 10: Comparisons of Optimal Weight Allocation with Zero Serial Correlation 68

Table 11: Further Comparisons of Optimal Weight Allocation with Serial Correlation 69

Table 12: Descriptive Statistics for CTA Data Set 71

Table 13: Optimal Weight Allocation with Serial Correlation 73

Table 14: Categorising CTAs using Population Means and Serial Correlation 74

Table 15: Comparison of Static Allocation and Dynamic Allocation using Rule 3 75

Chapter 1

Introduction

One of the fastest growing sectors of the financial services industry is the hedge-fund (Commodity Pool Operator) or the alternative investments sector. These hedge funds are now attracting major institutional investors such as large state and corporate pension funds and university endowments, and efforts are underway to make hedge-fund investments available to smaller investors through more traditional mutual-fund investment vehicles (Getmansky *et al.* 2003). Many hedge funds accomplish high returns by maintaining both long and short positions in securities and other instruments, hence the term “hedge” fund which, hypothetically, gives investors an opportunity to profit from both positive and negative information while, at the same time, providing some degree of ‘market neutrality’ because of the simultaneous long and short positions.

The assessment of portfolio performance is essential to both investors and funds managers. This also applies to Commodity Pool Operators (CPOs) and Commodity Trading Advisors (CTAs). Traditional portfolio measures present some limitations when applied to both parties. For example, the Sharpe ratio uses the excess reward per unit of risk as measure of performance, with risk represented by the variance (standard deviation). The mean-variance approach to the portfolio selection problem developed by Markowitz (1952) has frequently been the subject of undue criticism due to its utilization of variance as a measure of risk exposure when examining the non-normal returns of funds of hedge funds. These empirical properties may have potentially significant implications for assessing the risks and expected returns of hedge-fund investments, and can be traced to a single common source: significant serial correlation in their returns.

Such implications may come to some surprise because serial correlation is often associated with market inefficiencies, implying a violation of the Random Walk Hypothesis and the presence of predictability in returns. This seems inconsistent with the popular belief that the sector attracts the best and the brightest fund managers in the financial services sector. In particular, if a CPO's returns are predictable, the implication is that the manager's investment policy is not optimal; if their returns next month can be reliably forecasted to be positive, he or she should increase their positions to CTAs this month to take advantage of this forecast, and vice versa for the opposite forecast. By taking advantage of such predictability the CPO will eventually eliminate it, along the lines of Samuelson's (1965) original "proof that properly anticipated prices fluctuate randomly". Thus, with the hefty financial incentives of hedge-fund managers to produce profitable investment strategies, the existence of significant unexploited sources of predictability seems unlikely.

However, assuming that prices follow random walks, there is still to be considered from two perspectives the response of a CTA to profits and losses. First there is the human factor. Most CTAs are dominated by one or a few individuals. It seems plausible that individuals will adapt their trading decisions in the light of pass profits and losses. Moreover, the matter goes deeper than this. A hypothetical rational CTA will find it optimal to reflect past profits and losses in their trading decisions as long as they are risk averse. Even a purely rational hypothetical CTA would respond to profits and losses unless he was risk neutral. In the light of this, serial correlation in CTA performance is not merely an aberration[1].

1.1 Motivation

In this dissertation we develop an investment theory that integrates old portfolio theory with post-1952 technical offerings of academia to reshape the way in which resources are distributed and re-distrusted by CPOs. Such a theory would recognise that client/beneficiary value creation of optimal reallocation for trading investments. The theory of optimising a portfolio of assets in terms of expected yield and risk is well established. This Modern Portfolio Theory combines the work of Harry Markowitz, Merton Miller and William Sharpe to arrive at mathematical conclusions that are intended to help you invest efficiently across all asset classes. The theory combines several assumptions with historical performance figures for groups of investments to build an optimal portfolio.

[1] This proposition finds empirical verification in the results of Chapter 8.

Tobin (1958) added to the portfolio theory by introducing the "Efficient Frontier". According to the theory, every possible combination of securities can be plotted on a graph comprising of the standard deviation of the securities and their expected returns on its two axes. The collection of all such portfolios on the risk-return space defines an area, which is bordered by an upward sloping line. This line is termed as the efficient frontier. The collection of portfolios which fall on the efficient frontier are the efficient or optimum portfolios that have the lowest amount of risk for a given amount of return or alternately the highest level of return for a given level of risk.

One of the problems faced during the application of the original portfolio theory was that when a portfolio is created based on only statistical measures of risk and returns, the results obtained are overly simplistic. This problem was overcome in the Black-Litterman Model by simply postulating that the initial expected returns are the basically required returns so as to maintain equilibrium of the portfolio with that of the market (Black and Litterman 1991; 1992).

However, in our context there is the added complexity that if CTA decisions are affected by past profits and losses then the efficient frontier is continually shifting. Therefore, without conflicting with established theory, our focus in this dissertation is on reallocation rather than allocation.

1.2 Research Methodology

What lies behind portfolio theory is the factoring of yield as follows. Consider the yield on an asset during a time period T, which may be a month. We assume that the portfolio is reconsidered monthly, but not during any month. Therefore we have:

$$Y_T = H_T \, \Delta P_t \tag{1}$$

Y_T is the yield from an asset in time period T, H_T is the holding of the asset in period T and ΔP is the net change in the price of the asset during the period. However, when the investment is a trading investment this explanation of yield needs to be reinterpreted and extended. The essence of a trading investment is that the holding of the asset changes frequently. Therefore, we start by reinterpreting H_T, the mean holding of the asset during the time period. We can now replace equation (1) with the following:

$$Y_T = H_T \, \Delta P_t + \int_{T-1}^{T} h(t) \, \dot{p}(t) \, dt \tag{2}$$

In equation (2), *h(t)* is the deviation of the holding from its mean at time point *t* which is during time period $T. \dot{p}(t)$ is the rate of change of price at time point *t.* The second term on the right hand side of equation (2) is essentially the covariance (during period *T*) of the size of the holding and the rate of change of price. It is convenient to refer to this as the coordination of holding and price change. By comparison with equation (1), equation (2) shows a fundamental distinction between asset investments and trading investments. However, one observation that applies to both sorts of investment is that yield depends symmetrically on price movements and holdings. In the case of trading investments, this amounts to trading decisions and allocation decisions being equally important.

Much of the literature on dynamic allocation for investments is limited. Brown *et al.* (2001), study CPO and CTA managers' variance strategy based on past performance and survival, however, they fail to incorporate a framework which involves reallocation. Authors such as Rachev *et al.* (2004) focus on issues with optimal asset allocation whilst Jangmin *et al.* (2006) developed a stock trading method that incorporates dynamic asset allocation. The latter provides some foundation however still not directly related to trading investments and more specifically reallocation to trading investments.

The purpose of the dissertation is to generalise the established literature on portfolio optimisation so that it applies to trading investments. We are interested in CPO reallocation to CTAs. This study differs from Ding and Ma (2010); Jangmin *et al.* (2006) as we focus exclusively on trading investment reallocation instead of asset reallocation. Similar to Elton *et al.* (1987), who apply a portfolio approach of Markowitz to examine portfolio allocation in commodity markets, we adopt a portfolio approach to develop our dynamic reallocation framework for trading investments.

1.3 Outline of the Dissertation

Before describing our dynamic reallocation model, we provide a review of the pertinent literature in Chapter 2. In Chapter 3, there is a theoretical model of the ways in which past profits and losses can affect the CTA so as to have consequences for the expected value of future profits. If CTAs do not respond to their own profits and losses, but keep trading as if nothing had happened, then we would not have a theoretical basis for dynamic reallocation. However, the theoretical model examines several ways in which a CTA might plausibly respond to past performance such that future expected performance is altered. This constitutes the theoretical basis for the proposition that dynamic reallocation can enhance the performance of a trading investment. We show that with such a framework, we are able to apply the methodology to analyse both simulated and published data. The model is simulated in Chapter 4 and we derive its implications for further development. Several methods of categorising our findings are proposed in Chapter 5. We devise a composite weight allocation system to investigate how a dynamic reallocation system could be adapted to the performance characteristics of a CTA in Chapter 6. Chapter 7 analyses the effects of incorporating serial correlation in CPO returns for the dynamic reallocation model and we apply these methods to a dataset of 10 CTAs spanning over a 10 year period. These findings are summarised in Chapter 8, and concluded in Chapter 9.

Chapter 2

Literature Review

In this chapter, we review the foundations of portfolio allocation from a normative point of view. We then investigate the research literature on dynamic decision making. It is extremely difficult to find useful normative theories for these kinds of decisions which directly focus on dynamic investment allocation or even reallocation as this field is still in its infancy. Therefore our research focuses on descriptive issues of dynamic decision making as a whole which can bolster the framework of this study.

2.1 Portfolio Allocation

Modern Portfolio Theory (MPT) (Markowitz 1959) began laying the foundations of portfolio allocation from the normative point of view. In MPT one of the most influential concepts emphasized by Markowitz is diversification. According to Markowitz, this is a risk-management technique where various investments are combined in order to reduce the risk of the portfolio. It is argued that many investors do not adequately diversify their portfolios. This may be due to beliefs that the risk is defined at the level of an individual asset rather than the portfolio level, and that it can be avoided by hedging techniques, decision delay, or delegation of authority (De Bondt 1998). Reallocation decisions are comparable to variable hedging, except that the technique can be used to increase exposure as well as reduce it.

Since the findings of Markowitz, many subsequent authors have sought to investigate diversification in order to uncover underlying strategies which could improve MPT. Benartzi and Thaler (2001) studied naive diversification strategies in the context of defined contribution saving plans. They found evidence of 1/n heuristic, as a special case of diversification heuristic, in which an investor spreads their contributions evenly across available investment possibilities. Benartzi and

Thaler further elicit that such a strategy can be problematic both in terms of ex ante welfare costs and ex post regret (in case the returns differ from historical norms). It is important to remember that naive diversification does not imply coherent decision making. Although it may be a reasonable strategy for some investors, it is unlikely that the same strategy would be suitable for all investors, who obviously differ on their risk preferences and other risk factors, such as age (Lovrić 2011). 1/n heuristic can produce a portfolio that is close to some point on the efficient frontier. Nonetheless, a naive diversification strategy stands as a very strong benchmark, as shown by DeMiguel *et al.* (2007). By comparing out-of-sample performance of various optimising mean-variance models, they illustrated that no single model consistently beats the 1/n strategy in terms of the Sharpe ratio or the certainty-equivalent return. Poor performance of these optimal models is due to errors in estimating means and covariances (DeMiguel *et al.* 2007).

Although dynamic reallocation to a group of CTAs is, in a sense, beyond the scope of this dissertation, the reallocation rules we consider do amount to an active approach to diversification. Even if the same dynamic reallocation rule is applied to two or more CTAs, differences in CTA performance lead to a changing distribution of the portfolio across CTAs. In other words, a dynamic approach to diversification is implicit in the work which follows.

2.1.1 Mental Accounting

Mental Accounting is an economic concept established by Thaler (1980) which argues that individuals divide their current and future assets into separate, non-transferable portions. The theory purports individuals assign different levels of utility to each asset group, which affects their consumption decisions and other behaviours. Rather than rationally viewing every dollar as identical, mental accounting helps explain why many investors designate some of their dollars as ‘safety’ capital which they invest in low-risk investments, while at the same time treating their ‘risk capital’ quite differently. Benartzi and Thaler (2001) also found a support for mental accounting on the company stock: when company stock is in the array of available investment options, the total exposure to equities is higher than when it is not available. It seems that company stock is given a separate mental account different from the rest of equity classes.

2.1.1.1 Individual Stock Accounting

Extensive experimental work suggests that loss aversion and narrow framing are important features of the way people evaluate risky gambles. Barberis and Huang (2001) incorporate these two ideas to help to divide mental accounting in two forms; Individual Stock accounting and Portfolio accounting. To do this, they state that when the investor is loss averse over individual stock fluctuations and chose consumption C_t and an allocation $S_{i,t}$ to stock i to maximize:

$$E\sum_{t=0}^{\infty}\left[\rho^t \frac{C_t^{1-\gamma}}{1-\gamma}+ b_0\, \bar{C}_t^{-\gamma}\, \rho^{t+1} \sum_{i=1}^{n} v\left(X_{i,t+1}, S_{i,t}\,, z_{i,t}\right)\right] \tag{3}$$

The first term in this preference specification, utility over consumption C_t, is standard in asset-pricing models. The parameter ρ is the time discount factor, and $\gamma > 0$ controls the curvature of utility over consumption. The second term models the idea that the investor is loss averse over changes in the value of individual stocks. "The variable $X_{i,t+1}$ measures the gain or loss on stock i between time t and time $t + 1$, a positive value indicating a gain and a negative value, a loss. The utility the investor receives from this gain or loss is given by the function v, and it is added up across all stocks owned by the investor. It is a function not only of the gain or loss itself, but also of $S_{i,t}$, the value of the investor's holdings of stock i at time t, and of a state variable $z_{i,t}$, which measures the investor's gains or losses on the stock prior to time t as a fraction of $S_{i,t}$. By including $S_{i,t}$ and $z_{i,t}$ as arguments of v, we allow the investor's prior investment performance to affect the way subsequent losses are experienced" (pp. 1256).

Another pertinent literature which supports the findings and analysis of Barberis and Huang is the work of Barberis *et al.* (2001). We shall use both sources to draw a clearer understanding to relevant concepts and bolster the model description which they present. Barberis *et al.* further formalised the notion of loss aversion in a model of the aggregate stock market. The essential structure their specification expresses the gain or loss on stock *i* between time *t* and *t* + 1 is measured as:

$$X_{i,t} = \quad S_{i,t}\, R_{i,t+1} - S_{i,t}\, R_{f,t}\,. \tag{4}$$

We can derive from the above equation that the gain is the value of stock *i* at time *t* +1 minus its value at time *t* multiplied by the risk-free rate. By multiplying by the risk-free rate, this design models the idea that investors may only view the return on a stock as a gain if it exceeds the risk-free rate. The *t* in this instance is a year, therefore gains and losses are measured annually. Barberis *et al.* insist that while the investor may check his holdings much more often than that, even several times a day, they explicitly make the assumption that it is only once a year, perhaps at tax time, that he confronts his past performance in a serious way.

In a way analogous to Barberis *et al.*, we show assuming that dynamic reallocation decisions are made at the end of each month. It is extremely likely that CPOs will be monitoring CTA performance day by day or even more often, but we assume that a serious reappraisal of allocations is made only monthly. At the end of each month the CPO has available a sufficient amount of new performance data to warrant changing the allocation from that made the previous month.

The term $z_{i,t}$ tracks prior gains and losses on stock *i*. It is the ratio of another variable, $Z_{i,t}$, to $S_{i,t}$, so that $z_{i,t} = Z_{i,t} / S_{i,t}$. Barberis *et al.* refer to $Z_{i,t}$ as the “historical benchmark level” for stock *i*, to be thought of as the investor’s memory of an earlier price level at which the stock used to trade. When $S_{i,t} > Z_{i,t}$ or $z_{i,t} < 1$, the stock price today is higher than what the investor remembers it to be, making him feel as though he has accumulated prior gains on the stock, to the tune of $S_{i,t} - Z_{i,t}$. When $S_{i,t} > Z_{i,t}$, or $z_{i,t} < 1$, the current stock price is lower than it used to be, so that the investor feels that he has had past losses, again of $S_{i,t} - Z_{i,t}$ (Barberis and Huang 2001).

The variable $z_{i,\ t}$ is introduced to allow v to capture experimental evidence suggesting that the pain of a loss depends on prior outcomes. Barberis *et al.* illustrate this by defining v in the following way. When $z_{i,\ t} = 1$,

$$v(X_{i,t+1}\ S_{i,t}\ , 1) = \{X_{i,t+1}\ for\ X_{i,t+1}\ \geq 0 \quad \lambda X_{i,t+1}\ for\ X_{i,t+1} < 0'$$

(5)

$with\ \lambda > 1. For\ z_{i,t} < 1,$

$$v(X_{i,t+1}, S_{i,t}\ , z_{i,t}) = \{S_{i,t}\ R_{i,t+1} - S_{i,t}\ R_{f,t}\ for\ R_{i,t+1}$$
$$\geq z_{i,t}\ R_{f,t} \qquad S_{i,t}\ (z_{i,t}\ R_{f,t} -\ R_{f,t}) + \lambda S_{i,t}\ (R_{i,t+1} - z_{i,t}\ R_{f,t})\ for\ R_{i,t+1} < z_{i,t}\ R_{f,t}{}'$$

(6)

$and\ for\ z_{i,t} > 1,$

$$v(X_{i,t+1}, S_{i,t}\ , z_{i,t}) = \{X_{i,t+1}\ for\ X_{i,t+1}\ \geq 0 \qquad \lambda(z_{i,t})X_{i,t+1}\ for\ X_{i,t+1} < 0'$$

(7)

$with$

$$\lambda(z_{i,t}) = \ \lambda + k(\ z_{i,t} - 1),$$

(8)

$and\ k > 0.$

"When $z_{i,\ t} < 1$, the investor has accumulated prior gains on stock *i*. The form of $v\ (X_{i,\ t+1}, S_{i,\ t}\ , z_{i,\ t})$ is the same as for $v\ (X_{i,\ t+1}, S_{i,\ t}\ , 1)$ except that the kink is no longer at the origin but a little to the left; how far to the left depends on the size of the prior gain" (Barberis and Huang 2001: pp 1258). What Barberis and Huang imply here is that prior gains may cushion subsequent losses. Since a loss is cushioned by the prior gain, they presume it is less painful. However if substantial losses are incurred which subsequently deplete the investor's entire reserve of prior gains, it is once again penalized at the more severe rate of $\lambda > 1$.

Barberis and Huang show that in the case where $z_{i,\,t}$ >1, where stock *i* has been losing value, the form of v ($X_{i,\,t+1}$, $S_{i,\,t}$, $z_{i,\,t}$) has a kink at the origin just like v ($X_{i,\,t+1}$, $S_{i,\,t}$, 1) but differs from v ($X_{i,\,t+1}$, $S_{i,\,t}$, 1) in that losses are penalized at a rate more severe than λ, capturing the idea that losses that come after other losses are more painful than usual. How much higher than λ the penalty is, is determined by equation (8) and in particular by the constant *k*.

In this dissertation, our search for reallocation rules can be regarded as corresponding to Barberis and Huang's view of past performance as being important. On the one hand, Barberis and Huang are saying in effect that a CTAs trading may be influenced by past performance and, in particular, accumulation of profits and losses. On the other hand, we ourselves in seeking allocation rules take accumulated performance into account, and we consider doing so in several alternative ways.

To complete the model description, an equation for the dynamics of $z_{i,\,t}$ is needed. Again we refer to Barberis *et al.* who use:

$$z_{i,t+1} = \eta\left(z_{i,t}\,\frac{\bar{R}_i}{R_{i,t+1}}\right) + (1-\eta)(1), \tag{9}$$

Where $\bar{R}_i$ is a fixed parameter and $\eta \approx 1$. Note that if the return on stock *i* is particularly good, so that $R_{i,\,t} > \bar{R}_i$, the state variable $z_{i,\,t} = Z_{i,\,t} / S_{i,\,t}$ falls in value. This means that the benchmark level $Z_{i,\,t}$ rises less than the stock price $S_{i,\,t}$, increasing the investor's reserve of prior gains. What the authors are implying here is that equation (9) captures the idea that a particularly good return should increase the amount of prior gains the investor feels he has accumulated on the stock. They also insist that a particularly poor return depletes the investor's prior gains: If $R_{i,\,t+1} < \bar{R}_i$, then $z_{i,\,t}$ goes up, showing that $Z_{i,\,t}$ falls less than $S_{i,\,t}$, decreasing $S_{i,\,t} - Z_{i,\,t}$. The parameter η controls the persistence of the state variable and hence how long prior gains and losses affect the investor. If $\eta \approx 1$, a prior loss, say, will increase the investor's sensitivity to further losses for many subsequent periods (Barberis and Huang 2001: pp. 1259).

Embedded in equation (9) is the assumption that the evolution of $z_{i,t}$ is unaffected by any actions the investor might take, such as buying or selling shares of the stock. In several cases, this may perhaps be a reasonable assumption. What Barberis *et al.* imply at this point is that if the investor sells some shares for consumption purposes, it is plausible that any prior gains on the stock are reduced in proportion to the amount sold—in other words, that $z_{i,t}$ remains constant. However, more excessive transactions (such as selling one's entire holdings of the stock) might plausibly affect the way $z_{i,t}$ evolves. Thus, in order to keep their analysis tractable Barberis and Huang make a strong assumption that they do not.

The parameter $\bar{R}_i$ *is* not a free parameter, but is determined endogenously by imposing the requirement that in equilibrium, the median value of $z_{i,t}$ be equal to one. The idea behind this is that half the time, the investor should feel as though he has prior gains, and the rest of the time as though he has prior losses (Barberis and Huang 2001: pp. 1259). Thus $\bar{R}_i$, is typically of similar magnitude to the average stock return.

2.1.1.2 Portfolio Accounting

The second form of mental accounting which is considered in this study is portfolio accounting. This form of narrowing framing implies that investors are loss averse only over portfolio fluctuations, in particular, they choose consumption C_t and an allocation $S_{i,t}$ to stock i to maximize:

$$E\sum_{t=0}^{\infty}\left[\rho^t\,\frac{C_t^{1-\gamma}}{1-\gamma}+\,b_0\,\bar{C}_t^{-\gamma}\,\rho^{t+1}\sum_{i=1}^{n}v(X_{t+1},S_{\,t}\,,z_t)\right]. \tag{10}$$

The variable X_{t+1} is the gain or loss on the investor's overall portfolio of risky assets between time t and time t + 1, $S_t = \sum_{i=1}^{n} S_{i,t}$ is the value of those holdings at time t, and the z_t term measures prior gains and losses on the portfolio as a fraction of S_t . Once again, Barberis and Huang interpret v as a non-consumption source of utility, which in this case is experienced over changes in overall portfolio value and not over changes in individual stock value. Portfolio gains and losses are measured as:

$$X_{t+1} = \quad S_t\, R_{t+1} - S_t\, R_{f,t},$$

(11)

where R_{t+1} is the gross return on the portfolio. When z_{t+1}, v is defined as:

$$v\left(X_{i,t+1}\, S_t\,, 1\right) = \{X_{t+1}\, for\, X_{t+1}\ \geq 0 \quad \lambda X_{t+1}\, for\, X_{t+1} < 0'$$

(12)

$with\, \lambda > 1. For\, z_t < 1,$

$$v(X_{t+1}, S_t\,, z_t) = \{S_t\ R_{t+1} - S_t\, R_{f,t}\, for\, R_{t+1}$$

$$\geq z_t\, R_{f,t} \qquad S_t\left(z_t\, R_{f,t} -\ R_{f,t}\right) + \lambda S_t\left(R_{t+1} - z_t\, R_{f,t}\right) for\, R_{t+1} < z_t\, R_{f,t}{}'$$

(13)

$and\, for\, z_t > 1,$

$$v(X_{t+1}, S_t\,, z_t) = \{X_{t+1}\, for\, X_{t+1}\ \geq 0 \qquad \lambda(z_t) X_{t+1}\, for\, X_{t+1} < 0'$$

(14)

$with$

$$\lambda(z_t) = \quad \lambda + k(\, z_t - 1),$$

(15)

$and\, k > 0. Finally, the\, dynamics\, of\, z_t\, are\, given\, by$:

$$z_{t+1} = \ \eta\left(z_t\ \frac{\bar{R}}{R_{t+1}}\right) + (1 - \eta)(1).$$

(16)

To conclude, the functional forms are identical to what they were in the case of individual stock accounting. The only difference between the two formulations is that in equation (3), the investor experiences loss aversion over changes in the value of each stock that he owns, while in equation (10), he is loss averse only over overall portfolio fluctuations.

In our study, we develop an approach to dynamic reallocation that is comparable with individual stock accounting. A generalisation of this corresponding to portfolio accounting is beyond the scope of the dissertation, but remains a potential interesting line of development for the future.

2.1.2 Home Bias

Within the concept of portfolio allocation another robust finding which has been present in the literature is home bias. Regardless of the advantages of international portfolio diversification, French and Poterba (1991) found that the actual portfolio allocation of many investors is often too concentrated in their domestic market. There appears to be a limitation in the evidence on this concept because thus far, the literature has not provided a generally accepted explanation for the observed home bias. Huberman and Jiang (2006) argue that “familiarity breeds investment,” and that a person is more likely to invest in the company that he or she thinks they know. Instances of this familiarity bias are investing in domestic market, in company stocks, in stocks that are visible in investors’ lives, and stocks that are discussed favourably in the media. This effect would correspond to a CPO favouring a CTA with whom he has an established business relationship and familiarity. To some extent, we test for this in Chapter 8 where we allow for the possibility of a CPO using different reallocation rules for different CTAs.

2.1.3 Diversification – Demographic Variables

Goetzmann and Kumar (2001) examined the diversification of investors with respect to demographic variables of age, income, and employment. Their conclusions found that low income and non-professional categories hold the least diversified portfolios. They also found that young active investors tended to be more over-focused and inclined towards concentrated, undiversified portfolios, which might be a manifestation of overconfidence. It is questionable whether such overconfidence is justified. This underlines one of the purposes of this dissertation which is to introduce a greater degree of scientific justification in place of emotional justification.

2.2 Dynamic Decision Making

Whilst much of the literature on portfolio allocation is limited in context to dynamic reallocation, providing a thorough understanding of how decisions are made dynamically can perhaps bring us one step closer to bridging the gap in the literature. Dynamic decision-making (DDM) can be formerly defined as interdependent decision-making that takes place in an environment that changes over time either due to the previous actions of the decision maker or due to events that are outside of the control of the decision maker (Brehmer 1992; Edwards 1962). In this sense, dynamic decisions, unlike (simple) conventional one-time decisions, are typically more complex and often occur in real-time. Such decision making usually involves observing the extent to which people are able to use their experience (past history) to control a particular complex system, which could further include the types of experience that lead to better decisions over time.

Much of the research literature on DDM uses computer simulations which are laboratory analogues for real-life situations. These computer simulations are defined by Turkle (1984) as “micro-worlds” and are used to observe individuals’ behaviours in a simulated real world settings where individuals typically try to control a complex system where later decisions are affected by earlier decisions (Gonzalez *et al.* 2005). Examples of real world DDM scenarios include managing factory production and inventory, air traffic control, climate change, driving a car and a military command and control in a battle field. Research in DDM has focused on investigating the extent to which decision makers use their experience to control a particular system; the factors that underlie the acquisition and use of experience in making decisions; and the type of experiences that lead to better decisions in dynamic tasks. Applying this focus to trading investments, experience would in fact be substituted with the past investment time period ($t - n$) in order to aid CPOs in their decisions for reallocation. The decisions could usually only be understood as part of an ongoing process, and thus the decision problems of CPOs conform to Edwards’ (1962) classic description of dynamic decision making in that:

(1) A series of decisions is required to reach the goal. That is, to achieve and maintain control is a continuous activity requiring many decisions, each of which can only be understood in the context of the other decisions.

(2) The decisions are not independent. That is, later decisions are constrained by earlier decisions, and, in turn, constrain those that come after them.

(3) The state of the decision problem changes, both autonomously and as a consequence of the decision maker's actions.

Rapoport (1975) later provided a more formal definition, but this did not change the general meaning of dynamic decision making compared to Edwards' original definition. However, the three characteristics mentioned by Edwards did fully capture the essence of the decision problems we studied.

2.2.1 Characteristics of Dynamic Decision Making Environments

The primary characteristics of dynamic decision environments are dynamics, opaqueness, complexity, and dynamic complexity (Brehmer 1992). The dynamics of the environments refers to the dependence of the system's state on its state at an earlier time. Dynamics in the system could be driven by positive feedback (self-amplifying loops) or negative feedback (self-correcting loops). Simple examples of these could be the accumulation of interest in a fixed income investment or assuage of hunger due to eating respectively. In the case of a trading investment, the state of the system changes continually as profits and losses accumulate.

In a DDM context, the nature of opaqueness refers to the physical invisibility of some aspects of a dynamic system. Such a characteristic might also be dependent upon a decision maker's ability to acquire knowledge of the components of the system. In the situations which we are concerned, opaqueness is particularly important because it is inherent and cannot be removed. In some systems opaqueness arises because some aspect of the system is not being observed, but could be observed. In the case of a trading investment, the CTAs state of mind cannot in principle be observed – it is inherently opaque.

Complexity in principal refers to a collection of interconnected elements within a system that can make it difficult to predict the behaviour of the system. However, the definition of complexity can still itself have problems as system components can vary in terms of how many components there are in the system, number of relationships between them, and the nature of those relationships. Complexity may also be a function of the decision maker's ability (Brehmer and Allard 1991).

In contrast, dynamic complexity refers to the decision maker’s ability to control the system using the feedback the decision maker receives from the system. Diehl and Sterman (1995) separated dynamic complexity into three components. The opaqueness present in the system might cause unintended side-effects. There might be non-linear relationships between components of a system and feedback delays between actions taken and their outcomes. The dynamic complexity of a system might eventually make it hard for the decision makers to understand and control the system.

2.2.2 Dynamic Decision Making in Micro-worlds

Early studies on DDM used the term “micro-worlds[2]” to describe the complex simulations used in controlled experiments designed to study dynamic decisions (Dörner *et al.* 1983; Dörner *et al.* 1986). Research in dynamic decision-making is mostly laboratory-based and uses computer simulation micro-world tools – i.e., Decision Making Games. Such environments become the laboratory analogues for real-life situations and aid investigators in the study of decision-making by compressing space and time whilst simultaneously maintaining experimental control.

Many studies on DDM reveal that people have been shown to perform below the optimal levels of performance, if an optimal could be ascertained or known. For example, in a forest fire-fighting simulation game, participants frequently allowed their headquarters to be burned down (Brehmer and Allard 1991). Gonzalez and Vrbin (2007) studied DDM in a medical context and conveyed that

[2] Synthetic Task Environments, High Fidelity Simulations, Interactive Learning Environments, Virtual Environments and Scaled Worlds

participants acting as doctors in an emergency room allowed their patients to die while they kept waiting for results of test that were actually non-diagnostic. An interesting insight into decisions from experience in DDM is that mostly the learning is implicit, and despite people's improvement of performance with repeated trials they are unable to verbalize the strategy they followed to do so (Berry and Broadbent 1984). This is increasingly becoming the general consensus in the case of technical analysis and fund managers' performance.

2.2.3 Dynamic Decision Making in the Real-world

Much of the emphasis in the past has been on DDM using laboratory micro-world environments to investigate dynamic decisions however, there has been recent emphasis placed on DDM research which focuses on decision making in the real world. This does not demerit research in the laboratory micro-worlds, however instead it reveals the broader conception of the research underlying DDM. Under the DDM in the real world, individuals tend to be more interested in processes such as planning, perceptual and attention processes, forecasting, goal setting and many more (Gibson *et al.* 1997). The study of these processes brings DDM research closer to situation awareness and expertise.

In a real world research study, McKenna and Crick (1991) found that motorists who have more than 10 years of experience or expertise are faster to respond to hazards than drivers with less than three years of experience. Moreover, owing to their greater experience, such motorists tend to perform a more effective and efficient search for hazards cues than their not so experienced counterparts (Horswill and McKenna 2004). A potential explanation for such behaviour can be based upon the premise that situation awareness in DDM tasks makes certain behaviours automatic for people with expertise. Endsley (2006) also documented this behaviour for pilots and platoon commanders reporting considerations of novice and experienced platoon commanders in a virtual reality battle simulator show that more experience was associated with higher perceptual skills, higher comprehension skills. Thus, experience on different DDM tasks makes a decision maker more situational aware with higher levels of perceptual and comprehension skills.

2.2.4 Learning theories in Dynamic Decision Making

One of the main research activities in DDM has been to investigate the extent to which people are able to learn to control a particular simulated system and investigating the factors that might explain the learning in DDM tasks. Such a study of learning forms an integral part of DDM research and stands at the core of how a dynamic allocation system could eventually be integrated in a CPO's decision making structure.

2.2.4.1 Strategy-Based Learning Theory

The theory of strategy-based learning uses rules (or strategies) of action that communicate to a particular task. Such rules will often specify the conditions under which specifically designed rules or strategies will apply. These rules hold the form if you recognize situation 'A', then carry out action 'B'. In a study carried out by Anzai (1984) rules were implemented which performed the DDM task of steering a ship through a certain set of gates. The results of these rules were successful as they were able to mimic the performance on the task by human participants reasonably well. In a similar study, Lovett and Anderson (1996) illustrate how people use production rules of the "if – then" type in the building-sticks task which is an isomorph of "Lurchins' waterjug problem" (Lurchins 1942; Lurchins and Lurchins 1959).

2.2.4.2 Connectionist Learning Theory

Connectionist theory (or connectionism) is another explanatory learning theory which is believed by some to be a means of explaining learning in DDM tasks. It attempts to explain the connections between units, whose strength or weighing depend upon previous experience. Thus, the output of a given unit depends upon the output of the previous unit weighted by the strength of the connection. Gibson *et al.* (1997) studied a connectionist neural network machine learning model and concluded that such a model does a good job to explain human behaviour in the Berry and Broadbent's Sugar Production Factory task.

2.2.4.3 Instance-based Learning Theory

The Instance-Based Learning Theory (IBLT) is a theory of how humans make decisions in dynamic tasks. According to IBLT, individuals rely on their accumulated experience to make decisions by retrieving past solutions to similar situations stored in memory (Gonzalez *et al.* 2003). Gonzalez and Dutt (2011) extended this theory into two different paradigms of dynamic tasks, called sampling and repeated-choice. They show that in these dynamic tasks, IBLT provides the best explanation of human behaviour and performs better than many other competing models and approaches. Thus, decision accuracy can only improve gradually and through interaction with similar situations.

IBLT assumes that specific instances or experiences or exemplars are stored in the memory (Dienes and Fahey 1995). These specified instances have very concrete structures defined by three distinct parts which include the situation, decision, and utility (or SDU[3]).

IBLT also relies on the global, high-level decision making process, consisting of five stages: recognition, judgment, choice, execution, and feedback (Gonzalez and Dutt 2011). Gonzalez and Dutt's study found that when people are faced with a particular environment's situation, people are likely to retrieve similar instances from memory to make a decision. In typical situations (those that are not similar to anything encountered in the past), retrieval from memory is not possible and people would need to use a heuristic (which does not rely on memory) to make a decision. In situations that are typical and where instances can be retrieved, evaluation of the utility of the similar instances takes place until a necessity level is crossed (Gonzalez and Dutt 2011).

The necessity level is typically determined by the decision maker's "aspiration level," similar to Simon and March's satisficing strategy. However such a necessity level may also be determined by external environmental factors such as time constraints. As soon as the necessity level is reached, the decision involving the instance with the highest utility is made. The resulting outcome is then used to update the utility of the instance that was used to make the decision in the first place (from expected

[3] Situation refers to the environment's cues. Decision refers to decision maker's actions applicable to a particular situation. Utility refers to the correctness of a particular decision in that situation, either the expected utility (before making a decision) or the experienced utility (after feedback on the outcome of the decision has been received)

to experienced). This generic decision making process is assumed to apply to any dynamic decision making situation, when decisions are made from past experience.

The computational representation of IBLT relies on several learning mechanisms proposed by a generic theory of cognition. At present, some authors have implemented decision tasks into IBLT that have reproduced and explained human behaviour accurately (Martin *et al.* 2004; Gonzalez and Lebiere 2005).

2.2.5 Individual Differences in Dynamic Decision Making

Individual performance on DDM tends to be accompanied by significant degree of variability, which might be a result of the varying amount of skill and cognitive abilities of individuals who interact with the DDM tasks. While many studies have shown that individual differences are present in DDM tasks, there has been some debate on whether these differences arise as a result of differences in cognitive abilities. Some studies have failed to find evidence of a link between cognitive abilities as measured by intelligence tests and performance on DDM tasks. However, subsequent studies argued that this lack is due to absence of reliable performance measures on DDM tasks (Rigas *et al.* 2002; Gonzalez *et al.* 2005).

Gonzalez (2005) suggested a relationship between workload and cognitive abilities. His study found that low ability participants are generally outperformed by high ability participants. Furthermore, Gonzalez found that under demanding conditions of workload, low ability participants do not show improvement in performance in either training or test trials. An early study found that low ability participants use more heuristics particularly when the task demands faster trials or time pressure and this happens both during training and test conditions (Gonzalez 2004).

2.3 Literature Conclusions

DDM is a well established part of both decision theory (French 1988) and the field of decision analysis (Clemen and Reilly 2001). Throughout the remainder of the dissertation it will serve as one of the two key perspectives. Our concern is going to be with both CTA decisions and results, and the allocation decisions of CPOs which combine with them. In the following chapter we consider some specific theory relating to CTA decision making, and then we proceed to develop CPO decision rules.

The other, complementary, perspective is that of the area of the literature concerned with portfolio allocations. The way in which these two perspectives combine is that we shall be using the ideas of DDM to make a contribution to the literature on portfolio allocations. This of course will be in our specific area of dynamic allocations to trading investments.

Chapter 3

Data and Methodology

Collis and Hussey (2003) define 'methodology' as the overall approach to the research process, from the theoretical underpinning to the collection and analysis of the data. The purpose of this chapter is to discuss in detail the methodology which will be developed in order to present our findings for analysis. The study will follow a progressive pattern beginning with an overview of the questions being researched: for which the answers may be hard to find, along with possible explanations of the problems faced during the investigation of the answers. A brief synopsis of the approach used and the methodology applied for the purpose of research in relation to the topic will be given. The discussion will then be focused on the research methods rather than the research questions. A detailed discussion of the methods of acquisition of the data and the manner in which it would be arranged and/or classified will be presented. We focus on the detailed methodology applied to analyse the data acquired and how this will assist in drawing conclusions for the analysis.

3.1 Research Questions

The research questions proposed for this study centralise on the methodology of simulations. In constructing a simulation, McLeish (2011) outline distinct number of steps useful for this study;

1. Formulate the problem at hand. Why do we need to use simulation?

2. Set the objectives as specifically as possible. This should include what measures on the process are of most interest.

3. Suggest candidate models. Which of these are closest to the real-world? Which are fairly easy to write computer code for? What parameter values are of interest?

4. If possible, collect real data and identify which of the above models is most appropriate. Which does the best job of generating the general characteristics of the real data?

5. Implement the model. Write computer code to run simulations.

6. Verify (debug) the model. Using simple special cases insure that the code is doing what you think it is doing.

7. Validate the model. Ensure that it generates data with the characteristics of the real data.

8. Determine simulation design parameters. How many simulations are to be run and what alternatives are to be simulated?

9. Run the simulation. Collect and analyse the output.

10. Are there surprises? Do we need to change the model or the parameters?

11. Finally we document the results and conclusions in the light of the simulation results.

The above questions provide a sequential methodology which is pertinent to our approach in this study. Thus we will use McLeish's framework as a guideline as we develop the model in this chapter. Any problems which may arise at any specific step will be highlighted when necessary.

3.2 Synopsis of the Approach

The exploration for a successful dynamic allocation system will be conducted using a three stage approach. The first stage will devise and compare a set of rules in each of which the allocation to a CTA changes from month to month in response to monthly profit or loss. A series of monthly CTA percentage performance figures will be simulated with specified mean and variance. Using each of the twelve rules, this series of simulated returns will translate into monthly profits or losses. This series will further be used for comparing the dynamic allocation rules with one another whilst the CTA monthly percentage performance is held the same in all twelve cases.

The second stage will be to identify a tractable number of candidate rules and combine them into a dynamic allocation system and investigate how such a system could be adapted to the performance characteristics of a CTA.

The third stage and final stage will repeat the simulations of the second stage except with serial correlation introduced into the simulated monthly percentage performance figures. At this point, actual CTA performance data will be introduced.

3.3 Data Collection

The data used in this study can be categorised into two parts. Firstly, in regards to the sample needed to develop our model we will use a series of monthly CTA percentage performance figures which have been simulated as a series of identically and independently distributed observations on a Gaussian distribution. Next, in the secondary research, we will use actual CTA performance data for 10 CTAs from the Altegris Clearing Solutions, LLC database.

3.3.1 Duration of Collecting Data

As for duration of the data gathering period, this runs from January 2002 to December 2011. All figures obtained are classed as percentages of net unit profit or losses.

3.3.2 Sample of CTAs

The following table lists each CTA used in this study by name with the addition of the ‘default’ sample CTA data set. The default data set, in turn, refers to the simulated identically and independently distributed observations on a Gaussian distribution.

Table 1: List of all CTAs in the Data Set

Quest Partners
Dreiss Research Corp.
Kelly Angle Inc.
Saxon Investment Corp.
Hamer Trading, Inc.
Tactical Investment Management Corp.
Clarke Capital Management, Inc.
Mulvaney Capital Management Ltd
James H. Jones
GIC, LLC
Default

3.4 Methodology Research Design

Our starting point begins with the default CTA data, more specifically, the generation of a random series from the standard normal distribution which will then be scaled to give a standard deviation and mean. This process is vital to ensuring that our default series is appropriately scaled to emulate CTA returns. We use a random number function which provides us with a series which can then be altered. To generate the returns, we begin by specifying a variance (σ) and a population mean (μ). We set σ to 0.5 and the μ to 1. To calculate the returns, we use the following formula:

$$r_t = (N(x)\sigma) + \mu$$

(17)

From returns r_t we also compute the lagged returns giving a series of returns from r_{t-1}. The lagged returns will become useful later in this chapter when we develop the algorithmic framework for the chosen rules. Before devising the algorithmic rules, we shall first outline the theoretical basis for changing the allocation in response to profits or losses.

3.4.1 Theoretical Basis Model

Consider a CTA whose trading is based on a sequence of standard trading propositions based on a unit position size:

$$\eta = p\pi - (1 - p)\lambda$$

(18)

π = amount of contingent profit

λ = amount of contingent loss

p = probability of profit

η = expected value of profit or loss

The CTAs of interest are generally those who make a profit in the long run, so assume that:

$\pi > \lambda;\ p > 0.5.$

The question we need to address is whether there is any theoretical basis for changing the allocation to this CTA in response to profits or losses. The following analysis suggests that there is, and the argument consists of two stages. To connect the question of allocation to future performance, we assume that a CPO will want to have a larger allocation to the CTA when the expected value of future profit is greater. This leaves the issue of connecting future expected profit to past profit or loss.

This issue can be approached by considering the possible effect of past performance on the CTA's future trading. To do this, we can think of a pair of trades by the CTA, one following the other. The expected outcome after two trades keeping position size constant is:

$$p(\pi + \eta) + (1-p)[-\lambda + \eta] \tag{19}$$

This turns out to be twice the expected profit from a single trade, i.e. 2η. However, suppose now that in making his second trade the CTA is affected by whether the first trade is a profit or loss. As a general model of the effect of the CTA's response, the expected value of the profit or loss on the second trade is:

$$\varepsilon = p\alpha\pi - (1-p)\beta\lambda \tag{20}$$

Recalling that,

$$\eta = p\pi - (1-p)\lambda \tag{21}$$

Equation (20) can be rearranged to:

$$\varepsilon = \beta\eta + (\alpha - \beta)p\pi \tag{22}$$

According to equation (22), the CTA's response to a loss on the first trade has two effects. The first effect is that it scales the expected profit on a standard trade (η) by the factor β. If β<1, η is scaled down. If this were the only effect of the CTA's response then it would amount to reducing the expected profit on the second trade. However, there is also a second effect of the response, which is to add to the expected profit on the second trade the amount (α-β)pπ. If α>β the effect of this is to add to the expected profit on the second trade. If α<β this second effect reduces the expected profit on the second trade. In the following analysis, it will be convenient to refer to the first of the two effects as the scaling effect and the second effect as the addition or subtraction effect.

Comparing the expected value of profit on the second trade (ε) with the expected profit on the first trade (η), ε>η if and only if

$$\beta\eta + (\alpha - \beta)p\pi > \eta$$

$$\Leftrightarrow \qquad (\alpha - \beta)p\pi > (1 - \beta)\eta \tag{23}$$

There are many ways in which the CTA might respond to making a loss on the first trade. Four of these are considered in the following special cases. In each of the four cases we assume that if the first trade is profitable the second trade is a repeat of the standard profile.

3.4.1.1 CTA Response – Four Cases

Case A: cutting down position size

If the loss on the first trade makes the CTA more risk averse, one appropriate response would be to reduce the position size for the second trade. This would have the effect of scaling down the profit and loss on the second trade. In this case:

$$\alpha = \beta < 1 \tag{24a}$$

The expected profit on the second trade is found by substituting (24a) into equation (22). This gives:

$$\varepsilon = \beta\eta + (\beta - \beta)p\pi$$

$$\varepsilon = \beta\eta \tag{25a}$$

Because, in this case, $\alpha=\beta$ the second term in the equation for ε is zero. The only effect of scaling down the position size is to scale down the expected value of the profit. In other words, there is only a scaling effect, no addition or subtraction effect.

Case B: Snatching Profit

Following a loss on the first trade, the CTA might be inclined to accept a smaller profit rather than hold out for the full potential profit π. This would make:

$$\alpha < 1;\ \beta = 1 \tag{24b}$$

In this case equation (22) becomes

$$\varepsilon = 1\eta + (\alpha - 1)p\pi$$

$$\varepsilon = \eta - (1 - \alpha)p\pi \tag{25b}$$

Equation (25b) shows two things. First, there is no scaling effect. Secondly, there is a subtraction effect. Going back to condition (23) which is what is required for ε to be greater than η, we have the condition:

$$(\alpha - 1)p\pi > \eta \tag{26a}$$

Since $\eta>0$ and $\alpha<1$, condition (26a) cannot be satisfied. Snatching profit always reduces the expected profit from the second trade.

Case C: Cutting Loss

The loss on the first trade may make the CTA more loss averse when it comes to the second trade so that he cuts a loss sooner – before loss reaches the limit λ. This would make:

$$\alpha = 1;\ \beta < 1$$

(24c)

Substituting (24c) in equation (22) gives

$$\varepsilon = \beta\eta + (1-\beta)p\pi$$

(25c)

In the case of loss cutting there is a downward scaling effect but there is also an addition effect. These two effects act on the standard expected profit (η) in opposite directions and it is not clear which effect is larger. However, if the loss cutting parameters are substituted in condition (23) (the condition for ε being greater than η) it becomes:

$$(1-\beta)p\pi > (1-\beta)\eta$$

$$\Leftrightarrow \qquad p\pi > \eta$$

(26b)

Now recall that,

$$\eta = p\pi - (1-p)\lambda$$

The second term on the right hand side of this equation for η ($-(1-p)\lambda$) is always negative. It follows that $p\pi$ is always greater than η. So ε is always greater than η. In other words, loss cutting always creates an addition effect that outweighs the downward scaling effect.

Case D: Running Loss

An alternative response to loss cutting is that, having just taken one loss, the CTA may be more reluctant to take a second loss. If the second position has moved against him, the only way for the CTA to avoid taking a second loss is to continue to run the position. If the market continues to run against his position, he will end with a loss greater than the standard loss limit λ. This would mean

$$\alpha = 1;\ \beta > 1$$

(24d)

Substituting in equation (22), the expected profit on the second trade is now:

$$\varepsilon = \beta\eta + (1-\beta)p\pi$$

$$\varepsilon = \beta\eta - (\beta-1)p\pi$$

(25d)

The equation for ε is the same as in loss cutting (Case C) but now β>1 and it is more convenient to rewrite the equation as in version (25d). There is now an upward scaling effect – the standard expected profit η is scaled up by the factor β which is greater than 1. However, there is now a subtraction effect and it is not clear which of these two effects outweighs the other.

If the loss running parameters are substituted in condition (23) (the condition for ε being greater than η) it becomes,

$$(1-\beta)p\pi > (1-\beta)\eta$$

However, since β>1 the factor 1-β is negative so that cancelling 1-β on both sides reverses the direction of the inequality and gives

$$p\pi < \eta$$

(26c)

as the condition for ε being greater than η. Just as condition (25d) is always true, so condition (26c) can never be satisfied. It follows that loss running always has a net negative effect on expected profit. In other words, the subtraction effect always outweighs the scaling effect.

3.4.2 Methodological Framework Implications

The results derived in each of the four cases considered above can be summarised as in the following table.

Table 2: Summary of Four Cases

Response of CTA to loss	*Effect on second trade expected profit*
Cutting down position size	*Reduced*
Snatching profit	*Reduced*
Cutting loss	*Increased*
Running loss	*Reduced*

The first observation to be drawn from these results is that in each of the four cases there is a definite affect of past performance (the first trade) on future performance (the second trade). If a CPO allocating to this CTA wishes to relate the size of the allocation to the expected future performance, so as to improve profit, then the CPO should adjust the allocation in response to past profits and losses.

However, the second observation is that the appropriate response depends on the way in which the CTA himself responds to profits and losses. If the CTA responds to a loss by cutting position size, snatching profit or running loss, then this implies a reduction in expected future profit following a past loss. The appropriate reallocation by the CPO would be to reduce the allocation to the CTA. On the other hand, if the CTA's response is to cut losses following a loss, this will improve future expected profit and the CPO should respond by increasing the allocation.

The point about this second observation is that the CPO may well not know how the CTA responds to his own profit or loss, and he may not be able to discover it. The foregoing theoretical analysis has demonstrated that there are potential gains from dynamic reallocation whenever CTAs are affected by their profits or losses. However, it leaves the question of how such dynamic reallocation ought to be done to be discovered empirically and by the simulation analysis that follow.

3.4.3 Rule Construction

We now begin the development of the algorithmic rules which will primarily be responsible for the allocation and reallocation to our trading investments.

Each of the rules defined will consist of a position (P) and the corresponding result (π) for the specified month (t). The position will be initially set to 100 units. The result will be computed by multiplying the position by the corresponding monthly returns:

$$\pi_t = P_t \, r_t \tag{27}$$

The base case of this framework holds a constant or static allocation. Therefore, whilst all the rules developed in this framework will begin with an allocation position of 100, the base case's *P* will not change irrespective of profits or losses. Such a base case provides an adequate benchmark to make comparisons to all other rules developed on the basis of performance.

Similar to Lovett and Anderson (1996), we will use the "if – then" type from the Strategy-Based Learning Theory. We introduce patterns for the specification of properties within each rule in Structured English (Flake *et al.* 2000). These patterns are based on frequently used specification patterns identified by Dwyer *et al.* (1998; 1999). Each rule will be specified and represented using the Structured English framework which embodies a pseudo code algorithm loop. These self-correcting loops (Goetz 2011) embrace the concept of 'negative feedback' (Mees 1981) as the chain of cause-and-effect creates a circuit designed to stabilise the system and improve monthly performance. Ramaprasad (1983) defines this feedback generally as "information about the gap between the actual level and the reference level of a system parameter which is used to alter the gap in some way", emphasising that the information by itself is not feedback unless translated into action. In this instance our dynamic allocation algorithm presents a complete causal path that leads from the initial detection of the profit or loss to the subsequent modification of the position.

The algorithms will produce 3 data types: constants, deterministic variables and random variables[4]. McLeish (2005) insists that constants are relatively straightforward to define. One merely needs to select some scalar values that seem reasonable based on subjective theory or previous research and implement them in the simulation. In this instance our constant will be the starting position which is the allocation to the first month. Deterministic variables are vectors of values that take a range of values in a pre-specified non-random manner (McLeish 2005). In this context our deterministic variables will be a column vector of monthly positions which are computed using the prescribed rule. Finally, the generation of random variables tends to be formidably difficult for the following two reasons outlined by McLeish. Firstly, he insists it is often difficult to determine how a variable is distributed and which of the many standard distributions best represents it. Secondly, the computer algorithms to generate random variables are a great deal more complex than those needed to generate deterministic or constant variables. With respect to random variables in this study, we will use random values on the normal (or Gaussian) distribution function to generate random returns.

3.4.3.1 Rule Definition

Constant Allocation Rule states that if a profit or loss is made then the position will remain the same.

Algorithm:

```
{
 Public static final int STARTING_UNITS = 100;
      Public static int p;
      Public static int r;
      p = STARTING_UNITS;
      IF (r > 0) THEN {p = p} ELSE {p = p};
}
```

[4] The random component of most social processes is what makes statistical estimation problematic

The foregoing algorithm clarifies the dimensions to be explored by the further 12 rules set out below. In essence, what each rule is looking to do is increase or decrease position size depending on past performance. Thus, there are 3 dimensions. The first of these is scale; should the adjustments to position size be small or large and should it be related to the size of accumulated performance. The second is asymmetry; should increases in position size and decreases be equal or different. Thirdly, there is the question of time scale; should we look at performance last month or performance extending further back.

Dynamic Allocation Rule 1 states that if a profit is made then the position will increase by 5 however if a loss is made, the position will decrease by 5.

Algorithm:

```
{
 Public static final int STARTING_UNITS = 100;
      Public static int p;
      Public static int r;
      p = STARTING_UNITS;
      IF (r > 0) THEN {p = p + 5} ELSE {p = p - 5};
}
```

Dynamic Allocation Rule 2 states that if a profit is made then the position will remain the same however if a loss is made, the position will increase by 5.

Algorithm:

```
{
 Public static final int STARTING_UNITS = 100;
      Public static int p;
      Public static int r;
      p = STARTING_UNITS;
      IF (r > 0) THEN {p = p} ELSE {p = p + 5};
}
```

Dynamic Allocation Rule 3 states that if a profit is made then the position will increase by 5 however if a loss is made, the position will decrease by 10.

Algorithm:

```
{
 Public static final int STARTING_UNITS = 100;
      Public static int p;
      Public static int r;
      p = STARTING_UNITS;
      IF (r > 0) THEN {p = p + 5} ELSE {p = p - 10};
}
```

At this point, the subsequent pattern of the algorithm changes and now incorporates a series of monthly returns in the profit or loss query function. This rule structure incorporates connectionism (Marcus 2001; Elman *et al.* 1996) which is a theory which attempts to explain the connections between units, whose strength or weighing depend upon previous experience. Thus, the output position depends upon the result output of the previous unit(s) weighted by the strength of the connection.

Dynamic Allocation Rule 4 states that if a profit is made over a specified 3 month period then the position will increase by 5 however if a loss is made over this specified 3 month period, the position will decrease by 5.

Algorithm:

```
{
 Public static final int STARTING_UNITS = 100;
      Public static int p;
      Public static int r1, r2, r3, r3Month;
      p = STARTING_UNITS;
      r3Month = r1 + r2 + r3;
      IF (r3Month > 0) THEN {p = p + 5} ELSE {p = p - 5};
}
```

Dynamic Allocation Rule 5 states that if a profit is made over a specified 3 month period then the position will increase by 5 however if a loss is made over this specified 3 month period, the position will remain the same.

Algorithm:

```
{
 Public static final int STARTING_UNITS = 100;
      Public static int p;
      Public static int r1, r2, r3, r3Month;
      p = STARTING_UNITS;
      r3Month = r1 + r2 + r3;
      IF (r3Month > 0) THEN {p = p + 5} ELSE {p = p};
}
```

Using this same pattern we will now implement the following 2 rules using a longer period of 6 months.

Dynamic Allocation Rule 6 states that if a profit is made over a specified 6 month period then the position will increase by 5 however if a loss is made over this specified 6 month period, the position will decrease by 5.

Algorithm:

```
{
 Public static final int STARTING_UNITS = 100;
      Public static int p;
      Public static int r1, r2, r3, r4, r5, r6, r6Month;
      p = STARTING_UNITS;
      r6Month = r1 + r2 + r3 + r4 + r5 + r6;
      IF (r6Month > 0) THEN {p = p + 5} ELSE {p = p - 5};
}
```

Dynamic Allocation Rule 7 states that if a profit is made over a specified 6 month period then the position will increase by 10 however if a loss is made over this specified 6 month period, the position will decrease by 5.

Algorithm:

```
{
 Public static final int STARTING_UNITS = 100;
      Public static int p;
      Public static int r1, r2, r3, r4, r5, r6, r6Month;
      p = STARTING_UNITS;
      r6Month = r1 + r2 + r3 + r4 + r5 + r6;
      IF (r6Month > 0) THEN {p = p + 10} ELSE {p = p - 5};
}
```

The pattern of the algorithms will change even further and in this case the lagged returns computed in section 3.4 will be used to construct the remaining rules.

Dynamic Allocation Rule 8 states that if a profit is made then the position will increase by the lagged returns multiplied by 5 however if a loss is made, the position will decrease by the lagged returns multiplied by 2.

Algorithm:

```
{
 Public static final int STARTING_UNITS = 100;
      Public static int p;
      Public static int r;
      Public static int rLag;

      p = STARTING_UNITS;
      IF (r > 0) THEN {p = p + (rLag * 5)} ELSE {p = p - (rLag * 2)};
}
```

Dynamic Allocation Rule 9 states that if a profit is made then the position will increase by the lagged returns multiplied by 5 however if a loss is made, the position will decrease by the lagged returns multiplied by 10.

Algorithm:

```
{
 Public static final int STARTING_UNITS = 100;
      Public static int p;
      Public static int r;
      Public static int rLag;

      p = STARTING_UNITS;
      IF (r > 0) THEN {p = p + (rLag * 5)} ELSE {p = p - (rLag * 10)}
}
```

Dynamic Allocation Rule 10 states that if a profit is made then the position will increase by the lagged returns multiplied by 2.5% of the previous position however if a loss is made, the position will decrease by the lagged returns multiplied by 5% of the previous position.

Algorithm:

```
{
 Public static final int STARTING_UNITS = 100;
      Public static int p;
      Public static int _p;
      Public static int r;
      Public static int rLag;

      p = STARTING_UNITS;
      IF (r > 0) THEN {p = p + (rLag * (_p * 0.025))} ELSE {p = p -
         (rLag * (_p * 0.05))};
      _p = p;
}
```

Dynamic Allocation Rule 11 states that if a profit is made then the position will remain the same however if a loss is made, the position will decrease by the lagged returns multiplied by 5% of the previous position.

Algorithm:

```
{
 Public static final int STARTING_UNITS = 100;
      Public static int p;
      Public static int _p;
      Public static int r;
      Public static int rLag;

      p = STARTING_UNITS;
      IF (r > 0) THEN {p = p} ELSE {p = p - (rLag * (_p * 0.05))};
      _p = p;
}
```

Dynamic Allocation Rule 12 states that if a profit is made then the position will increase by the lagged 'squared' returns multiplied by 5% of the previous position however if a loss is made, the position will decrease by the lagged squared returns multiplied by 10% of the previous position.

Algorithm:

```
{
 Public static final int STARTING_UNITS = 100;
      Public static int p;
      Public static int _p;
      Public static int r;
      Public static int rLag;


      p = STARTING_UNITS;
      rLag = rLag^2;
      IF (r > 0) THEN {p = p + (rLag * (_p * 0.05))} ELSE {p = p -
         (rLag * (_p * 0.1))};
      _p = p;
}
```

3.4.4 Simulation Development

The central part of our dynamic reallocation system is the rules which we have created. Whilst time constraints restrict the development of a full 'micro-world' (Turkle 1984), the established rules provide a fundamental starting point in DDM for CPOs in respect to trading investments. The conceptual framework incorporates important variables which may perhaps provide some new insights when we analyse each of their performances individually in Chapter 4. With this said, this process emphasises the importance of developing a competent results simulator.

3.4.4.1 Monte Carlo Methods

A Monte Carlo simulation may perhaps be a useful approach to adopt at this point for a number of reasons. Logically it allows the population of interest to be simulated. From this pseudo-population, repeated random samples can be drawn (Mooney 1997). Although the logic may not be hard to grasp, the execution on the other hand, can be. Monte Carlo work is highly computer intensive and therefore complicated models can consume large amounts of time.

In this study, Monte Carlo simulation offers an alternative to analytical mathematics for understanding a statistic's (CTA's) sample distribution and evaluating its behaviour in random samples. It does this empirically by using random samples from known populations of simulated data to track a CTAs performance. Although the simulation approach involves a computational burden, it is at least tractable whereas the various non-linearities in the rules make an analytical approach infeasible.

3.4.4.2 Simulation Model Implementation

The simulation development process begins with creating an input-form which takes input values for the CTA data set and the number of simulations (n). When prompted the user will input such values and run the simulation. The 'Run' function will load the corresponding data from a CTA and input n into the counter for the computational loop. Each rule will consist of a 'Total Result' and an 'Accumulative Result' which stores the profit/loss totals for the iteration and a sum of all completed iterations thus far, respectively. The simulation repeats until $n = 0$ then the Run simulator will end and display an output listing each rule and its corresponding accumulative total.

An additional feature of the simulator consists of an option to switch the parameters around. For example, if a rule consists of an increase of 5 and a decrease of 10, the parameter switch will swap over giving an increase of 10 and a decrease of 5. The purpose of this approach will become evident in the subsequent Chapters 4 & 5.

To choose candidate rules and combine them into a dynamic allocation system we will generate a list of random weights which will be allocated to each of the top 4 rules. The first complication which arises is the random number generator produces numbers containing 10 decimal places. In order to avoid further complications a separate sub function will be written in order to perfectly generate random numbers between 0 and 1 which only contains one decimal point. This sub function is then called in the weight generation algorithm which stores the overall performance of all the rules which have been allocated weights. While such a process is accurate, it may not perhaps be optimal. Therefore the weight generation algorithm will loop 1000 times to find the most optimal weight allocation which provides superior performance. After this process is complete the algorithm generates an output listing each rule, its corresponding weight and the superior performance result total.

The final aspect of the simulator incorporates serial correlation (φ) into the monthly returns. This is inputted using the following formula:

$$r_t = \left[\left(\phi\, r_{t-1} + N(x)\right)\sigma\right] + \mu \tag{28}$$

The φ variable will be later manipulated to uncover results which will be subject to analysis in the subsequent Chapters 7 & 8.

The methodology created in this chapter provides a robust framework which will be used to investigate each different aspect of the 3 stage approach in greater detail. Subsequent chapters will attempt to undercover original findings and present corresponding analysis where necessary. With the simulator now complete, our next objective is to use the developed rules to generate performance results which can be categorised, analysed and evaluated.

Chapter 4

Rules of Allocation

In this chapter, we briefly outline each allocation rule, giving a description of their structure, starting position (*P*) at a specific point in time (*t*) and the results which are provided using *n* simulations. We do not intend to give an in depth analysis of the resulting simulations but rather outline the performance of the allocation rule against constant allocation and any other allocation rule which provides similar performance. We propose each rule will provide results which fall into the following categories: High, Medium and Low performers. Such categorisation will be extended in Chapter 5, when discussing more alternative underlying dynamics.

4.1 Constant Allocation Rule

The purpose of this rule is to illustrate what happens to returns if the allocation position remains the same throughout the investment period. Therefore position is not dependent on performance in any way.

> **Outcome:** This rule provides returns which we can use to make comparisons with the dynamic allocation approaches.
>
> **Performance Category:** Low

4.2 Allocation Rule 1

This rule is structured so if profit is made our position increases by 5 and if a loss is made it decreases by 5.

Outcome: This rule provides the closest results to the constant allocation rule signifying that an equal change upward or downward in a profit or loss state respectively, doesn't provide any significant difference in returns.

Performance Category: Low

4.3 Allocation Rule 2

This rule is structured so if profit is made our position remains the same however if a loss is made it increases by 5.

Outcome: This rule provides improved results more than 3 times greater than the constant allocation rule.

Performance Category: Medium

4.4 Allocation Rule 3

This rule is structured so if profit is made our position increases by 5 and if a loss is made it decreases by 10.

Outcome: This rule provides results which are approximately twice as large as Rule 2 due to the relationship between of the position increases. This rule would follow the same behavioural pattern as Rule 2 just with more allocation to the position. By making the loss component greater this tends to provide an improvement in results.

Performance Category: High

4.5 Allocation Rule 4

This rule is constructed using a 3 month parameter. This parameter shifts down one position each time it verifies positive or negative returns. Therefore if profit is made in a selected 3 month period our position increases by 5 and if a loss is made it decreases by 5.

Outcome: This rule provides results which are twice as great as Rule 2 which holds the same allocation amount to the position. The 3 month parameter however, provides a significant improvement for this allocation rule over Rule 2 although results are still below other rules.

Performance Category: Low

4.6 Allocation Rule 5

This rule is constructed again using a 3 month parameter. The parameter shifts down one position each time it verifies positive or negative returns. Consequently if profit is made in a selected 3 month period our position increases by 5 and if a loss our position remains the same.

Outcome: This rule provides improved results on Rule 4 which holds the same 3 month parameter however the allocation adjustment to the position seems to be a contributory factor to this improvement. Moreover it provides more than 3 times the returns of constant allocation.

Performance Category: Medium

4.7 Allocation Rule 6

This rule is constructed using a 6 month parameter. This parameter shifts down one position each time it verifies positive or negative returns. Therefore if profit is made in a selected 6 month period our position increases by 5 and if a loss is made it decreases by 5.

Outcome: This rule provides results which are twice as great as Rule 2 which holds the same allocation amount to the position. The 6 month parameter however, provides results closely similar to Rule 4 which hold the same allocation state to the position. Yet it still doesn’t provide significantly improved results in comparison to other rules.

Performance Category: Low

4.8 Allocation Rule 7

This rule is constructed using a 6 month parameter. Therefore if profit is made in a selected 6 month period our position increases by 10 and if a loss is made it decreases by 5.

Outcome: This rule provides significantly positive results. Whilst being 6 times greater than constant allocation, its results are also close to those of Rule 3. Surprisingly, these two rules have opposite allocation for profit and loss. This leads us to believe that the 6 month parameter has had a significantly positive effect on returns for this case.

Performance Category: High

4.9 Allocation Rule 8

This rule is constructed using the lagged returns as a function which determines the allocation. Therefore if profit is made our position increases by lagged returns multiplied by 5 and if a loss is made it decreases by lagged returns multiplied by 2.

Outcome: This rule provides results which are significantly better than most rules. The lagged returns function allows the use of returns in $t - 1$ as a weight when increasing or decreasing allocation.

Performance Category: Medium

4.10 Allocation Rule 9

This rule is constructed using the lagged returns as a function which determines the allocation. Therefore if profit is made our position increases by lagged returns multiplied by 5 and if a loss is made it decreases by lagged returns multiplied by 10.

Outcome: This rule again provides results which are significantly better than most rules. It also provides an improvement on the results of Rule 8. The lagged returns function allows the use of returns in $t - 1$ as a weight when increasing or decreasing allocation. By increasing the allocation component if a loss is made this provides improved results overall.

Performance Category: High

4.11 Allocation Rule 10

This rule is constructed using the lagged returns as a function which determines the allocation. In addition, it also relies on the previous position (P_{t-1}) as a function of the allocation component. Therefore if profit is made our position is increased by lagged returns multiplied by 2.5% of the previous position and if a loss is made it decreases by lagged returns multiplied by 5% of the previous position.

Outcome: This rule again provides results which are significantly better than most rules. It provides results which are similar to those of Rule 8 and justifiably better that constant allocation.

Performance Category: Medium

4.12 Allocation Rule 11

This rule is constructed using the lagged returns as a function which determines the allocation. In addition it also relies on the P_{t-1} as a function of the allocation component. Therefore if profit is made our position remains the same, however if a loss is made it decreases by lagged returns multiplied by 5% of the previous position.

Outcome: This rule provides results which are slightly below similar rules which use lagged returns as a function. By holding our position constant this noticeably reduces the performance of this rule.

Performance Category: Medium

4.13 Allocation Rule 12

This rule is constructed using the lagged 'squared' returns as a function which determines the allocation. In addition it also relies on the P_{t-1} as a function of the allocation component. Therefore if profit is made our position is increased by lagged squared returns multiplied by 5% of the previous position and if a loss is made it decreases by lagged squared returns multiplied by 10% of the previous position.

> **Outcome:** This rule provides results which are significantly larger than all rules which is expected are the returns are squared. Looking past this factor results are still an improvement on previous results which use the lagged returns as a function.
>
> **Performance Category:** High

This algorithmic framework constitutes an excellent test case, since the established rules are able to provide a good indicator of which rule can be also expected to work well in further experimental simulations.

Chapter 5

Rule Categorisation

In exploring the results of the twelve rules in Chapters 3 & 4, one issue was asymmetry, i.e. unequal responses to profit and losses. The results in Chapter 4 show that all four high performing rules are asymmetrical. However, the amounts by which position sizes are adjusted were largely arbitrary and for the purpose of exploring each type of rule. We therefore now consider the potential benefits of changing the adjustments amounts.

In this chapter, we extract the allocation rules which provide only high performance compared to the residual rules. The purpose of this approach is to investigate further into the discovery of any fundamental ingredients which can be extracted and manipulated in order to improve performance. We categorise each rule as a top performer and computationally alter the parameters to create augmented rules (*a*) in order to make comparisons with the original results generated by the allocation rule. Setting σ to 0.5 and μ to 1, we use n simulations to generate the following results for analysis:

Table 3: Comparison of Results of Allocations Rules with Augmented Parameters

	Original Allocation Rule			*Augmented Allocation Rule*		
	n = 50	n = 100	n = 200	n = 50	n = 100	n = 200
Rule 3	14243.6	28607.93	57009.23	14516.96	28927.04	57110.08
Rule 7	14687.86*	28300.14	56800.04	13853.97*	28397.85	56937.23
Rule 9	12916.86	25842.85	51485.57	16050.77	30812.56	60610.42
Rule 12	47590.59	94814.48	188150.95	209193.53	344823.37	724176.9

5.1 Augmented Allocation Rule 3

This rule is structured so if profit is made our position increases by 10 and if a loss is made it decreases by 5. Compared to the original allocation Rule 3, this rule provides results which are closely similar, however overall the parameter switch has improved the performance.

5.2 Augmented Allocation Rule 7

This rule is again constructed using a 6 month parameter. However on this occasion if profit is made in a selected 6 month period our position increases by 5 and if a loss is made it decreases by 10. As can be seen from the table this rule performs quite similar to the original rule with only one exception[5]. Whilst this exception occurred using 50 iterations, subsequent experiments consistently showed that this rule provides improved results.

5.3 Augmented Allocation Rule 9

This rule is constructed using the lagged returns as a function which determines the allocation. However on this occasion if profit is made our position increases by lagged returns multiplied by 10 and if a loss is made it decreases by lagged returns multiplied by 5. The table shows that this rule had a significant improvement on performance compared to the original. This was consistently illustrated overall using all values for n.

5.4 Augmented Allocation Rule 12

This rule is constructed using the lagged 'squared' returns as a function which determines the allocation. In addition, it also relies on P_{t-1} as a function of the allocation component. Using augmented parameters if profit is made our position is increased by lagged squared returns multiplied by 10% of the previous position and if a loss is made it decreases by lagged squared returns multiplied by 5% of the previous position. While it is clear that this rule has a very significant improvement compared to the original allocation rule, this was expected due to the exponential growth from the

[5] Marked with a: *

squared factor. As a result at this stage this rule will be considered as unstable, for this reason, and will be withdrawn from our top performers' categorisation model.

Overall the augmented rules have collectively illustrated improved performance. Thus, from this point forward, we will use these rules as the base algorithmic computations of further experiments and analysis.

Chapter 6

Weight Allocations

The simulations thus far have produced four high performing rules. Among these, Rule 12 tends to produce unstable returns. On this ground, we have dropped Rule 12 and carry forward Rules 3, 7 and 9. However all three of these remaining rules are candidates to be the most preferred dynamic allocation rule. Since these are three candidates, we now explore the potential for combining them.

In this chapter, we utilise the extraction technique used in Chapter 5 to determine the most superior performing rules to attempt to develop a combined rule for further analysis. The purpose of this combined rule is to, month by month, take an allocation defined by each of the refined rules in order to make up a composite allocation which is a weighted average of the 3 allocations – e. g. (0.3, 0.5 and 0.2). We formulate the weighting function to generate a weight which sums to 1. These weights are generated using the following algorithms for random number (Rand) generation:

Table 4: Weight Computation for Composite Allocation

	Weight Definition	*Weight Computation*	*Weight Result*
Rule 3	w_1	= Rand for w_1	= w_1_result
Rule 7	w_2	= Rand * (1 – w_1)	= w_2_result
Rule 9	w_3	= (1 – (w_1 + w_2))	= w_3_result
Total			= $w_1 + w_2 + w_3$

6.1 Weight Generation

The weight generation algorithm first generates a random number between 0 and 1 and rounds this number to 1 decimal point. This weight is then initialised to a type ‘Double’ in order to handle decimals and the result is assigned to Rule 3. For Rule 7, the weight generation algorithm generates another random number between 0 and 1 and rounds this number to 1 decimal point, and multiplies this by 1 minus w_1. This weight is then initialised to a type 'Double' and the result is assigned to Rule 7. Rule 9 is generated by summing the previous w_1 and w_2 and subtracting this total from 1. This result is assigned to Rule 9 as the final w_3. Because of the rounding function computed when the weight is generated, this ensures that total of all 3 weights always sums to 1 precisely.

6.2 Optimal Weighting using Composite Allocation Algorithm

As the weight generation algorithm now generates 3 weights, the subsequent task must now simulate this process whilst incorporating the performances of each rule to devise an optimal composite allocation. After computing the first 3 weights the result is stored and compared to the next result generated from the $n + 1$ simulation. If this result is greater, then that result becomes the new optimal weighting for the composite allocation. Using the previous performance figures, we generate weights for an initial 10 simulations below:

Table 5: Weighting for the Composite Allocation – 10 Simulations

	Performance (P)	*Weight (w)*	*(P * w)*
Rule 3	28027.04	0.2	5605.41
Rule 7	28397.85	0.2	5679.57
Rule 9	30812.56	0.6	18487.54
Total			29772.51

As seen in the above table the results are computed by multiplying the performance of each rule by the optimal weight generated when $n = 10$. The best overall performance of the chosen composite rule gives a total of 29772.51. We will now simulate more variations of n using the corresponding performance for each rule to determine whether our results can be improved.

Table 6: Weighting for the Composite Allocation for Different variations of *n*

	n = 50		*n = 100*		*n = 200*		*n = 1000*	
	(P)	*(w)*	*(P)*	*(w)*	*(P)*	*(w)*	*(P)*	*(w)*
Rule 3	14516.96	0.2	28027.04	0.1	57110.08	0.1	256629	0.1
Rule 7	13853.97	0.1	28397.85	0.1	56737.23	0.1	253204.5	0.1
Rule 9	16050.77	0.7	30812.56	0.8	60610.42	0.8	261235.1	0.8
Total	15524.33		30292.54		59873.07		259971.44	

The above table shows different variations of *n* and the performance results for each rule using that variation. The composite allocation total provides the total performance for the optimal allocation which has been computed by the algorithm using *w* for each rule. At low values of *n*, distributed weights are suggested. For example, with 50 iterations Rule 9 dominated, but Rule 3 nevertheless gets twice the weight of Rule 7. However, our results show that as *n* increases, the weight generation algorithm chooses one rule as the 'dominant' rule and allocates it a maximal 0.8 weighting. Consequently, the remaining two rules are given two weight allocations of 0.1 each. Thus, we have discovered that our weight generation algorithm now selects the best rule of allocation (dominant rule) based on the performance of the composite allocation.

6.3 Calculating Optimal Weights using Alternative Population Means

The above success of disentangling performance ingredients to optimally allocate weights based on high performing rules now brings us to question whether the population mean (μ) plays a significant role in the composite allocation. To further investigate the effect which this may have we shall begin by using a set of performance results generated from n = 100[6] and using an additional two values for μ.

Table 7: Comparisons of Optimal Weight Allocation with Different Population Means

	μ = 0.1		μ = 0.5		μ = 1	
	(P)	*(w)*	*(P)*	*(w)*	*(P)*	*(w)*
Rule 3	3004.88	0.8	14345.53	0.8	28430.34	0.1
Rule 7	2581	0.1	14179.08	0.1	28669.46	0.1
Rule 9	1437.11	0.1	10722.01	0.1	30279.07	0.8

As our performance results shows that the population mean can in fact affect our P which in turn affects our optimal choice of weights. We can see from the above table that when μ tends to be lower, the rule which is allocated the dominant weight (0.8) is Rule 3. Conversely when μ is higher, Rule 9 is allocated the dominate weight. This gives further justification to investigate results with more variations for μ for the full range of –1 to +1 in order to have a clearer picture of the optimal weight allocation. The table below illustrates the corresponding results:

Table 8: Optimal Weight Allocation with Full Range of Population Means

	μ from -1 to -0.1	μ = 0	μ from 0.1 to 0.8	μ from 0.9 to 1
	(w)	*(w)*	*(w)*	*(w)*
Rule 3	0.1	0.8 or 0.1 or 0.1	0.8	0.1
Rule 7	0.1	0.1 or 0.8 or 0.1	0.1	0.1
Rule 9	0.8	0.1 or 0.1 or 0.8	0.1	0.8

As can be seen from the above table when μ is negative our dominant weight is allocated to Rule 9. However, a μ of zero tends to make no difference which rule is chosen and the results show that

[6] In table 6, further increases in n to 200 and 1000 left the weights unchanged

either Rule 3, 7 or 9 could be allocated the dominant weight. On the other hand, whenever μ ranges from 0.1 to 0.8 Rule 3 holds the dominate weight allocation. Finally, when μ reaches 0.9 to 1, the dominant rule returns back to being Rule 9. It might appear that taking Rule 9 to be the dominant rule would be a fair summary of the results, but the alternative choice of Rule 3 is of great practical importance as we shall see when we progress to analyse actual data from CTAs in Chapter 8.

The weight generation algorithm has shed some light on the need for examining dynamic variables which are incorporated in the calculation of returns. With this said, the fundamental problem which arises is attempting to measure something which cannot be observed. We will subsequently have to evaluate the possibly of even answering the question itself. In the next chapter another such variable will be incorporated within this study to aid in the discovery of how other factors may affect the rule which holds the dominant weighting in the composite allocation.

Chapter 7

Incorporating Serial Correlation

In this chapter, we adopt the similar methodology established in section 6.3 to compute optimally allocated weights based on high performing rules whilst varying the μ from -1 to +1. However in this section we incorporate a new variable which tends to be found in CPO returns called Serial Correlation (φ). Serial correlation is the degree to which each month's returns for a CTA mirror the results of the month before. If in such an instance a CTA generates the exact same returns amount every month it is said to be perfectly serially correlated. With this kind of performance—a nice, smooth line going up no matter what the market does— this is a good indicator that you should take a closer look. Lo (2008) shows that from the pattern of historical returns in hedge-fund databases, when funds' returns grow too consistently, this is a significant sign that the investments are either very hard to value accurately and the returns are just guesses, or, worse, that they've been manipulated in a way that smoothes them artificially. In this study, there is an additional reason to consider serial correlation. Our starting point was the proposition that a CTA's trading might be affected by past results. If this proposition carries over to a monthly time scale, e.g. if trading this month is affected by last month's profit or loss, then we should expect to see serial correlation.

In order to apply serial correlation to our study the weights are generated using a set of performance results generated from n = 50, 100 and 200, using two values for φ (0.5 and 0.9). The different variations of n are used to establish robustness to the weight allocation and thus the actual performance results generated from these are not included as our focus here is on the optimal weight allocation.

Table 9: Comparisons of Optimal Weight Allocation with Serial Correlation

$\varphi = 0.5$

	μ from -1 to -0.1	$\mu = 0$	μ from 0.1 to 0.7	$\mu = 0.8$	μ from 0.9 to 1
	(w)	*(w)*	*(w)*	*(w)*	*(w)*
Rule 3	0.1	0.8 or 0.1 or 0.1	0.8 or 0.1	0.8 or 0.1 or 0.1	0.1
Rule 7	0.1	0.1 or 0.8 or 0.1	0.1 or 0.8	0.1 or 0.8 or 0.1	0.1
Rule 9	0.8	0.1 or 0.1 or 0.8	0.1	0.1 or 0.1 or 0.8	0.8

For ease of comparison the results from $\varphi = 0$ were as follows:

Table 10: Comparisons of Optimal Weight Allocation with Zero Serial Correlation

$\varphi = 0$

	μ from -1 to -0.1	$\mu = 0$	μ from 0.1 to 0.7	$\mu = 0.8$	μ from 0.9 to 1
	(w)	*(w)*	*(w)*	*(w)*	*(w)*
Rule 3	0.1	0.8 or 0.1 or 0.1	0.8	0.8	0.1
Rule 7	0.1	0.1 or 0.8 or 0.1	0.1	0.1	0.1
Rule 9	0.8	0.1 or 0.1 or 0.8	0.1	0.1	0.8

7.1 Calculating Optimal Weights with Serial Correlation

When we incorporate φ at 0.5 into our simulations the above table shows performance results which are seemingly identical to table 8 in section 6.3. While this may have been anticipated, the 3rd column containing μ values is significantly different in two ways. Firstly, the range for the 3rd column is no longer 0.1 to 0.8, it now omits 0.8 and categorises this separately[7]. Secondly, both Rule 3 and 9 can be anticipated to have the superior performance at any stage from 0.1 to 0.7. Moreover, from 0.1 to 0.6 these two rules perform seemingly identical with only a small fraction separating them at any given time. By 0.7, the pattern begins to change again and although Rules 3 and 7 still provide superior

[7] Columns 3 & 4 of the above table have been modified in order to allow easier comparisons. See table 8 in section 6.3 for its original structure.

performance, Rule 9 seems to evidently be very close behind them in performance. At μ = 0.8, Rule 9 becomes increasing stronger and by now all three rules seems to produce similar results with either rule outperforming the others at any given time. Finally, as we've seen before when μ is relatively high (0.9 to 1), Rule 9 becomes the dominant rule again.

As this experiment provides some insightful findings, this validates its usefulness and gives further justification to incorporate φ at 0.9. The following table lists the weight allocation findings for such simulations:

Table 11: Further Comparisons of Optimal Weight Allocation with Serial Correlation

φ = 0.9

	μ from -1 to -0.1	μ = 0	μ = 0.1	μ = 0.2 to 0.3	μ = 0.4 to 0.5	μ = *0.6 to 1*
	(w)	*(w)*	*(w)*	*(w)*	*(w)*	*(w)*
Rule 3	0.1	0.1	0.8 or 0.1	0.1	0.8 or 0.1	0.1
Rule 7	0.1	0.1	0.1	0.1	0.1	0.1
Rule 9	0.8	0.8	0.1 or 0.8	0.8	0.1 or 0.8	0.8

While it is evident that serial correlation can have an impact on the optimal weight allocation, the size of this serial correlation factor still requires further investigation. This may perhaps be beyond the scope of this project. Nonetheless, by incorporating a larger φ, we can begin to see some startling insights from the above results. The above table illustrates that similar to the previous simulations when μ from -1 to -0.1, the weight allocation algorithm still selects Rule 9 as the dominant rule. However on this occasion, a μ which equals 0 also selects Rule 9 as the dominant rule, while previous results showed all 3 rules could have been dominant at this μ level. Furthermore the pattern changes again at 0.1 when Rule 3 and Rule 9 can be allocated the dominant rule component. Subsequently, when μ is 0.2 to 0.3 the dominant rule returns back to being Rule 9. Again, the pattern reverses and we see both Rule 3 and Rule 9 being allocation 0.8 at any given time for a μ between 0.4 and 0.5. Finally, the pattern switches again back to Rule 9 being the most dominant rule. In addition to the above interpretation, performance results generated between a μ of 0.8 to 1 revealed Rule 9 outperformed the other rules by at least a factor of 2 on most simulations.

Our composite allocation algorithm seems to have provided results which are strongly centred on Rule 9 when serial correlation tends to be high. With the interpretation of the above results now complete, this allows us to build a framework to analyse actual CTA data performance in the next chapter.

Chapter 8

Empirical Analysis of CTA Performance

In this chapter, we substitute our Rand results for actual CTA performance results obtained from Altegris Clearing Solutions, LLC. The purpose of this approach warrants the need to incorporate CTA data to determine the usefulness of our developed framework. In addition this will allow for the categorisation of each CTA based on the μ and φ which are calculated using the standard deviation (σ) and the average of the data set ($\bar{x}$). The data used for this chapter is obtained from the sample period from January 2002 to December 2011.

Table 12: Descriptive Statistics for CTA Data Set

	$\bar{x}$	σ	μ	φ
Quest Partners	1.055167	6.388291	0.165172	-0.19768
Dreiss Research Corp.	1.999917	8.878613	0.225251	0.105815
Kelly Angle Inc.	1.81575	10.71336	0.169485	0.086683
Saxon Investment Corp.	1.692333	6.989112	0.242139	0.035563
Hamer Trading, Inc.	0.983833	5.923923	0.166078	0.090315
Tactical Investment Management Corp.	1.782	8.036235	0.221746	-0.00589
Clarke Capital Management, Inc.	1.10575	6.737651	0.164115	-0.10383
Mulvaney Capital Management Ltd	1.757833	9.710328	0.181027	0.150105
James H. Jones	1.1925	9.502597	0.125492	0.03702
GIC, LLC	1.125417	5.782725	0.194617	-0.02519

8.1 Calculating Descriptive Statistics for CTA Data Set

We begin by calculating the $\bar{x}$ for each CTA using the following formula:

$$\bar{x} = \frac{\sum x}{N} \tag{29}$$

Where the x values are published monthly rates of return and N is the sample size which is 120.This gives the mean total across the sample period for that CTA. The σ is computed using the following formula, which provides the standard deviation of returns of the sample period.

$$\sigma = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N}(x_i - \bar{x})^2} \tag{30}$$

The μ is computed by dividing the $\bar{x}$ by the σ of the CTA, giving the population mean using the following formula:

$$\mu = \frac{\bar{x}}{\sigma} \tag{31}$$

The φ is more complex to compute as it requires preliminary computation to be done before it can be calculated using the following correlation formula:

$$Correlation = \frac{(\sigma_{AB})}{sd(A)\ sd(B)} \tag{32}$$

Firstly, we compute the deviation from $\bar{x}$ for each month of the CTA data set by subtracting the $\bar{x}$ from each month's performance figure. This gives a monthly series of deviations from $\bar{x}$ which can then be lagged by one month to create the lagged deviation from $\bar{x}$. Using these two components, we then compute the cross-product of the deviation from $\bar{x}$ and the lagged deviation from $\bar{x}$. This computation now provides a series of monthly cross-products for each month beginning from the first lagged month onwards. The average of these cross-products is then computed and divided by the σ^2 giving a final column for φ.

8.2 Rule Categorisation – Incorporating Serial Correlation

The results generated in Chapter 7 will now allow the possible integration of serial correlation into our performance results when μ ranges from -1 to +1. Such a framework is illustrated in the below table and indicates which rule was allocated the dominant weight of 0.8.

Table 13: Optimal Weight Allocation with Serial Correlation

	$\varphi = 0$	$\varphi = 0.5$	$\varphi = 1$
μ from -1 to -0.1	*Rule 9*	*Rule 9*	*Rule 9*
μ = 0	*Rule 3, 7 or 9*	*Rule 3, 7 or 9*	*Rule 9*
μ = 0.1	*Rule 3*	*Rule 3 or 7*	*Rule 3 or 9*
μ from 0.2 to 0.3	*Rule 3*	*Rule 3 or 7*	*Rule 9*
μ from 0.4 to 0.5	*Rule 3*	*Rule 3 or 7*	*Rule 3 or 9*
μ from 0.6 to 0.7	*Rule 3*	*Rule 3 or 7*	*Rule 9*
μ = 0.8	*Rule 3*	*Rule 3, 7 or 9*	*Rule 9*
μ from 0.9 to 1	*Rule 9*	*Rule 9*	*Rule 9*

The above results illustrate that when serial correlation is high, Rule 9 is more desirable for the CPO as it provides the superior performance results. Recall that Rule 9 makes adjustments to the allocation in proportion to the size of profits and losses. It may be that such a rule is more effective when there is greater persistence in runs of profit. Similarly when serial correlation tends to be closer to zero, Rule 3 is allocated the dominance factor from the weight allocation algorithm. Between these two parameters either Rule 3, 7 or 9 is proven to provide superior results which could lead to a dominant weight allocation of 0.8 to either rule. Using this framework we are now able to categorise each CTA using the previously calculated φ and μ. The following table illustrates the results found after we categorise each CTA using the developed framework.

Table 14: Categorising CTAs using Population Means and Serial Correlation

	μ	φ	*Dominant Rule*
Quest Partners	0.165172	-0.19768	*Rule 3*
Dreiss Research Corp.	0.225251	0.105815	*Rule 3*
Kelly Angle Inc.	0.169485	0.086683	*Rule 3*
Saxon Investment Corp.	0.242139	0.035563	*Rule 3*
Hamer Trading, Inc.	0.166078	0.090315	*Rule 3*
Tactical Investment Management Corp.	0.221746	-0.00589	*Rule 3*
Clarke Capital Management, Inc.	0.164115	-0.10383	*Rule 3*
Mulvaney Capital Management Ltd	0.181027	0.150105	*Rule 3*
James H. Jones	0.125492	0.03702	*Rule 3*
GIC, LLC	0.194617	-0.02519	*Rule 3*

Our results thus far has deduced that based on the calculated μ and φ, all CTAs which have been examined in this study have been categorised under a single rule. This rule is Rule 3 unanimously. This does not by any means imply that every CTA would be allocated the dominant weight using this weight allocation algorithm, as further research is necessary to examine and larger pool of CTAs. Nonetheless this conclusion does bring to light the effectiveness of the composite weight allocation algorithm and the need for further development. In considering the foregoing outcome, it may be worth bearing in mind that all of the 10 CTAs chosen had track records of at least 10 years and they were all among the top performers. It may well be that less successful CTAs would benefit more from other rules.

We will now apply Rule 3 to the historical data collected for each CTA and make a comparison of the results generated from the dynamic allocation.

Table 15: Comparison of Static Allocation and Dynamic Allocation using Rule 3

	Static Allocation	*Dynamic Allocation*	*Percentage Change (%)*
Quest Partners	127	364	+186.61%
Dreiss Research Corp.	240	713	+197.08%
Kelly Angle Inc.	218	540	+147.71%
Saxon Investment Corp.	203	598	+194.58%
Hamer Trading, Inc.	118	430	+264.41%
Tactical Investment Management Corp.	214	625	+192.06%
Clarke Capital Management, Inc.	133	317	+138.35%
Mulvaney Capital Management Ltd	211	673	+218.96%
James H. Jones	143	286	+100.00%
GIC, LLC	135	335	+148.15%

On the basis of the historical results we have applied Rule 3 and compared the bottom line performance with a static allocation – Appendix A. The above table shows that in every case the performance with the reallocation rule exceeded that with a static allocation.

The effectiveness of the dynamic allocation system varies between CTAs. The smallest increase in performance is 100% while the largest is 264%. The average improvement across all 10 CTAs is 179%. The fact that the percentage improvement differs from CTA to CTA raises the question of how much the performance enhanced by dynamic allocation depends on the CTA in question. As a preliminary remark, the CTA whose performance was improved by the smallest factor (*James H. Jones*) produced a better static allocation performance than the CTA whose performance was most increased (*Hamer Trading, Inc.*). This raises the suggestion that perhaps relatively weak CTAs gain most from dynamic allocation. However, across the 10 CTAs there is correlation of +0.9 between the dynamic allocation performance and the static allocation performance. The appearance is therefore that the enhanced results depend on both the dynamic allocation system and the underlying CTA performance, and that these combine in a complex way specific to each CTA.

Chapter 9

Conclusions

The profit or loss on a trading investment depends on the performance of the CTAs and the allocations made by the CPO. The two factors combine multiplicatively so that they are both of equal importance. It is usual for trading decisions to be at least partly systematic, in other words, rule based. In contrast, allocations and reallocations are made on an ad hoc basis. In this dissertation, we have taken a step into the development of a scientifically validated system of rules and algorithms for dynamic allocations.

The fundamental axiom of dynamic allocation is that past performance affects future performance through the state of mind of a CTA. This creates the possibility of a rule based dynamic allocation system improving on the profitability of a static allocation. In Chapter 3 there is presented a theoretical model of the response of a CTA to past profit or loss such as to alter the expected value off future performance. The model implies that there are such effects, but that the direction and size of appropriate reallocations must be determined empirically. This, starting in Chapter 4, is what the dissertation proceeds to do. The objective of the dissertation was to determine whether such a dynamic allocation system could be found in respect of a varying allocation to a CTA.

The search for a successful dynamic allocation system was conducted in three stages. In keeping with the fundamental axiom, the first step was to devise and compare a set of rules in each of which the allocation to a CTA changes from month to month in response to monthly profit or loss. Twelve such rules were devised. In the first instance, a series of monthly CTA percentage performance figures was simulated as a series of identically and independently distributed observations on a Gaussian distribution with specified mean and variance. This one series of percentage returns was translated into monthly profits or losses using each of the twelve rules in turn. Thus, when we compared the twelve profit and loss series, we were comparing the dynamic allocation rules with one another, the CTA monthly percentage performance being the same in all twelve rules.

This simulation was repeated 50, 100 and 200 times and the results aggregated. The result of this was that three of the twelve rules emerged as being superior to the other rules in terms of accumulated profit and stability. The three rules are:

Rule 3: If profit is made our position increases by 5 and if a loss is made it decreases by 10.

Rule 7: If profit is made in a selected 6 month period our position increases by 10 and if a loss is made it decreases by 5.

Rule 9: If profit is made our position increases by lagged returns multiplied by 5 and if a loss is made it decreases by lagged returns multiplied by 10.

The performances of the 3 high performing rules were further altered with a parameter switch component. This provided improved performance for a new set of augmented rules. These three rules are:

Rule 3(a): If profit is made our position increases by 10 and if a loss is made it decreases by 5.

Rule 7(a): If profit is made in a selected 6 month period our position increases by 5 and if a loss is made it decreases by 10.

Rule 9(a): If profit is made our position increases by lagged returns multiplied by 10 and if a loss is made it decreases by lagged returns multiplied by 5.

Having identified a tractable number of candidate rules, our second step was to combine them into a dynamic allocation system and investigate how such a system could be adapted to the performance characteristics of a CTA. The rules were combined by sub-allocating a fraction of the CTA's allocation to be governed by each of the three rules. The profit-maximising weights on the three rules were found by way of repeated simulations. Weights ranging from 0.1 to 0.8 subject to the constraint of summing to 1 were considered. The first phenomenon to emerge from this step was that in each case one and only one rule or another was selected for a maximum weight of 0.8.

We then ran more simulations, varying the mean monthly percentage profit of the simulated CTA relative to the variance (which was held constant). This yielded a relationship between the dominant rule and mean profit. The essential results were that with a negative mean monthly rate of return, Rule 9 dominates, whereas with a positive mean monthly rate of return, Rule 3 usually dominates. At this stage, we first had a system for selecting a dynamic reallocation system based on the performance characteristics of any given CTA.

Our third and final stage was to repeat the simulations of step 2 but with serial correlation introduced into the simulated monthly percentage performance figures. Our findings in this respect were as follows. With a negative monthly rate of return, Rule 9 dominates whatever the strength of the serial correlation (if any). However, with positive rates of return, the presence of serial correlation introduces Rules 7 and 9 as close competitors to Rule 3. At this stage, a credible system for prescribing a dynamic allocation system for any given CTA was complete.

It was now time to confront the simulation-based prescription system with real CTA monthly percentage performance data. Ten years of monthly figures were downloaded from the Altegris Clearing Solutions, LLC database for each of ten CTAs. Importantly, the data which were used to evaluate the prescription and dynamic allocation systems had not played any part in the development of the systems. For each CTA, a dynamic allocation system was prescribed and applied. In each case, the result was an increase in profitability ranging from 100% to 264%.

Reverting to the original question of whether it is possible to devise a profitable dynamic allocation system, we are driven to an affirmative response, because that is what the dissertation has done.

9.1 Future Investigation Outline

There are several directions in which this research can be developed. Rules can be devised to determine simultaneously dynamic allocations to each of a group of CTAs. Aspects of CTA performance apart from the monthly mean and serial correlation can be introduced. Ad hoc human factors can be reconsidered now in the context of a formal dynamic allocation system. Consequently, the findings of this research are by no means an end, they are perhaps not even the end of the beginning. But they are the beginning of the beginning of a new approach to trading investments. What we have shown here is merely that such a thing can be done, but it is a thing which can offer investors improved performance, stability, objectivity and transparency.

Appendices

Appendix A: Comparison of Static and Dynamic Allocation using Rule 3

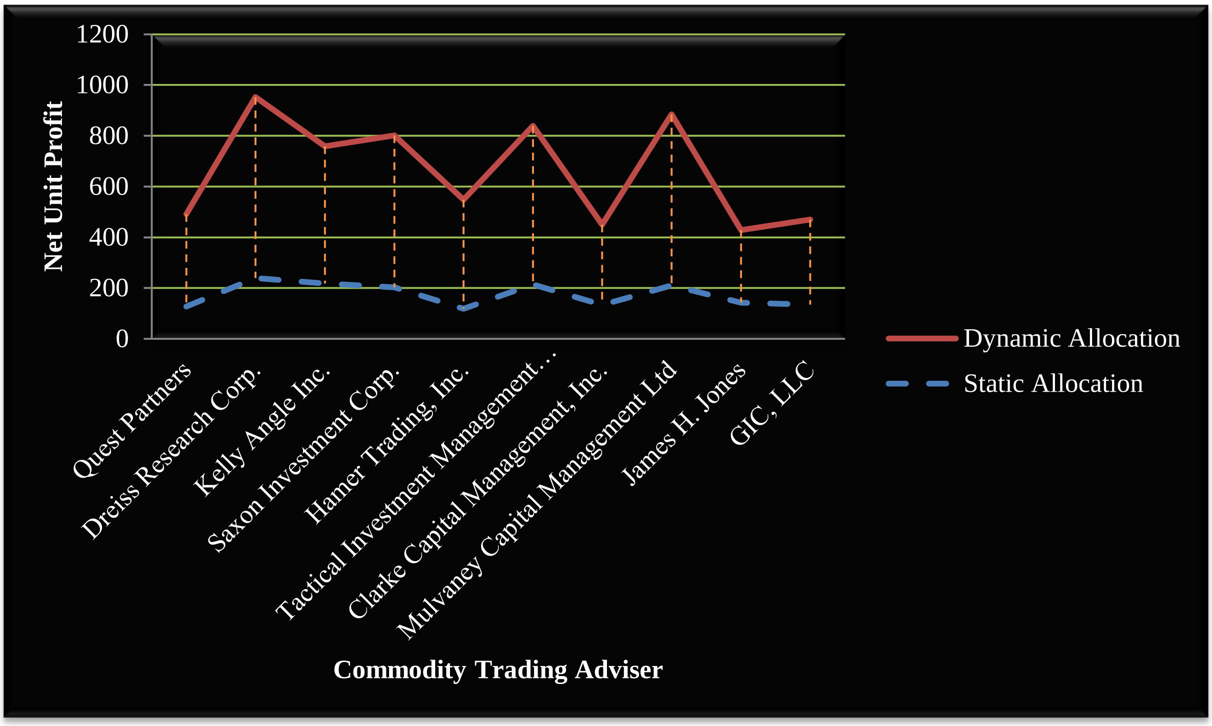

References

Anzai, Y. (1984). Cognitive control of real-time event driven systems. *Cognitive Science, 8*, 221–254.

Barberis, N., and Huang, M. (2001). Mental accounting, loss aversion, and individual stock returns, *Journal of Finance*, Vol. lvi, No. 4, August 2001

Barberis, N., Huang, M., and Santos, T., (2001). Prospect theory and asset prices, *Quarterly Journal of Economics* 116, 1–53.

Benartzi and R. H. Thaler. (2001). Naive diversification strategies in defined contribution saving plans. *The American Economic Review*, 91(1):79–98.

Berry, B. C., & Broadbent, D. E. (1984). On the relationship between task performance and associated verbalizable knowledge. *Quarterly Journal of Experimental Psychology, 36A*, 209–231.

Black, F. and Litterman, R. (1991). "Global Asset Allocation with Equities, Bonds, and Currencies." *Fixed Income Research*, Goldman, Sachs & Company, October.

Black, F. and Litterman, R. (1992). "Global Portfolio Optimization." *Financial Analysts Journal*, September/October, 28-43.

Brehmer, B. and R. Allard, (1991). 'Real-time dynamic decision making. Effects of task complexity and feedback delays'. In: J. Rasmussen, B. Brehmer and J. Leplat (eds.), Distributed decision making: Cognitive models for cooperative work. Chichester: Wiley.

Brehmer, B. (1992). Dynamic decision making: Human control of complex systems. *Acta Psychologica, 81*(3), 211–241.

Brown, S, J., Goetzmann, W, N., and Park, J., (2001). "Careers and Survival: Competition and Risk in the Hedge Fund and CTA Industry." *Journal of Finance*, 61, pp. 1869-1886.

Clemen, R.T., & Reilly, T. (2001). *Making hard decisions with decision tools*. Belmont, CA: Duxbury Press.

Collis, J. and Hussey, R. (2003). *Business Research. A Practical Guide for Undergraduate and Postgraduate Students*, 2nd ed., Palgrave Macmillan, New York

De Bondt. W. F. (1998). A portrait of the individual investor. *European Economic Review*, 42(3-5):831–844,

DeMiguel, L. Garlappi, and R. Uppal. (2007). Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *The Review of Financial Studies*, 22(5): 1915–1953.

Dienes, Z., & Fahey, R. (1995). Role of specific instances in controlling a dynamic system. *Journal of Experimental Psychology: Learning, Memory and Cognition, 21*(4), 848–862.

Diehl, E., & Sterman, J. D. (1995). Effects of feedback complexity on dynamic decision making. Organizational Behavior and Human Decision Processes, 62(2), 198–215.

Ding, L., and Ma, J., (2010), "Portfolio Reallocation and Exchange Rate Dynamics", Working Paper March 2010.

Dörner. D.. R. Kreuzig and T. Stäudel, (1983). Lohhausen: Vom Umgang mit Unbestimmtheit und Komplexitat. Bern: Huber.

Dörner, D.. T. Stäudel and S. Strohschneider. (1986). More-Programmdokumentation. Memorandum 23. Lehrstuhl Psychologie II, University of Bamberg.

Dwyer, M. B., Avrunin, G, S., and Corbett, J, C., (1998), Property Specification Patterns for Finite-State Verification. In M. Ardis, editor, Proceedings of the 2nd Workshop on Formal Methods in Software Practice, Clearwater Beach, Florida, pages 7–15.

Dwyer, M. B., Avrunin, G, S., and Corbett, J, C., (1999), Patterns in Property Specifications for Finite-State Verification. In 21st International Conference on Software Engineering, Los Angeles, California.

Edwards, W. (1962). Dynamic decision theory and probabilistic information processing. *Human Factors, 4*, 59–73.

Elman, J , L., Bates, E, A., Johnson, M, H., Karmiloff-Smith, A., Parisi, D., Plunkett, K., (1996). *Rethinking Innateness: A connectionist perspective on development*, Cambridge MA: MIT Press.

Elton, Edwin J., Martin J. Gruber, and Joel Rentzler. (1987). "Professionally Managed, Publicly Traded Commodity Funds." *Journal of Business*, 60, 175-199.

Endsley, M. R. (2006). Expertise and situation awareness. In K. A. Ericcson, N. Charness, P. J. Feltovich & R. R. Hoffman (Eds.), *The Cambridge handbook of expertise and expert performance* (pp. 633–651). Cambridge: Cambridge University Press.

Flake, S. , Mueller, W., and Ruf, J., (2000) Structured English for model checking specifications. In Proc. of the GI-Workshop "Methoden und Beschreibungssprachen zur Modellierung und Verifikation von Schaltungen und Systemen", Frankfurt, Germany.

French S. (1988) “Decision theory – An introduction to the mathematics of rationality”. Ellis Horwood.

French and J. Poterba. (1991). International diversification and international equity markets. *The American Economic Review*, 81(2):222–226.

Gibson, F. P., Fichman, M., & Plaut, D. C. (1997). Learning in dynamic decision tasks: Computational model and empirical evidence. *Organizational Behavior and Human Decision Processes, 71*(1), 1–35.

Getmansky, M., Lo, A., and Makarov, I., (2003). “An Econometric Model of Serial Correlation and Illiquidity in Hedge Fund Returns.” MIT working paper (2003).

Gonzalez, C. (2004). Learning to make decisions in dynamic environments: Effects of time constraints and cognitive abilities. *Human Factors, 46*(3), 449–460.

Gonzalez, C. (2005). The relationship between task workload and cognitive abilities in dynamic decision making. *Human Factors, 47*(1), 92–101.

Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science, 27*(4), 591–635.

Gonzalez, C., & Lebiere, C. (2005). Instance-based cognitive models of decision making. In D. Zizzo & A. Courakis (Eds.), *Transfer of knowledge in economic decision-making*. Palgrave Macmillan.

Gonzalez, C., Thomas, R. P., & Vanyukov, P. (2005). The relationships between cognitive ability and dynamic decision making. *Intelligence, 33*(2), 169–186.

Gonzalez, C., Vanyukov, P., & Martin, M. K. (2005). The use of microworlds to study dynamic decision making. *Computers in Human Behavior, 21*(2), 273–286.

Gonzalez, C., & Vrbin, C. (2007). Dynamic simulation of medical diagnosis: Learning in the medical decision making and learning environment MEDIC. In A. Holzinger (Ed.), *Usability and HCI for medicine and health care: Third symposium of the workgroup human-computer interaction and usability engineering of the Austrian Computer Society, USAB 2007* (Vol. 4799, pp. 289–302). Germany: Springer.

Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review, 118*(4), 523-551.

Goetzmann W. N. and Kumar. A. (2001). Equity portfolio diversification. NBER Working Paper Series.

Goetz, T., (2011). "Harnessing the power of Feedback Loops". *Wired Magazine* (Wired Magazine). http://www.wired.com/magazine/2011/06/ff_feedbackloop/. Retrieved 28 June 2012.

Huberman and W. Jiang. (2006), Offering versus choice in 401(k) plans: Equity exposure and number of funds. *The Journal of Finance*, 61(2).

Horswill, M. S., and McKenna, F. P. (2004). Drivers' hazard perception ability: Situation awareness on the road. In S. Banbury & S. Tremblay (Eds.), *A cognitive approach to situation awareness: Theory and application* (pp. 155–175). Aldershot, England: Ashgate.

Jangmin O., Jongwoo, L., Jae Won, L,. Byoung-Tak, Z., (2006). “Adaptive stock trading with dynamic asset allocation using reinforcement learning”, Information Sciences, 176, pp.2121–2147

Lee, E. B., & Markus, L. (1967). *Foundations of optimal control theory*. New York: Wiley.

Lo, A., (2008), “Hedge Funds: An Analytic Perspective (Advances in Financial Engineering)”, Princeton University Press, pp. 201 – 364

Lovett, M. C., & Anderson, J. R. (1996). History of success and current context in problem solving: Combined influences on operator selection. *Cognitive Psychology, 31*, 168–217.

Lurchins, A. S. (1942). Mechanization in problem solving. *Psychological Monographs, 54*(248).

Lurchins, A. S., & Lurchins, E. H. (1959). *Rigidity of behaviour: A variational approach to the effects of Einstellung.* Eugene, OR: University of Oregon Books.

Lovrić, M. (2011), “Behavioral Finance and Agent-Based Artificial Markets”, Ph.D. Thesis, pp. 29-32.

Marcus, G. F., (2001). *The Algebraic Mind: Integrating Connectionism and Cognitive Science (Learning, Development, and Conceptual Change)*, Cambridge, MA: MIT Press

Markowitz, H. M., (1952). Portfolio selection, The Journal of Finance 7, pp. 77-91.

Markowitz. H, M., (1959). *Portfolio selection: efficient diversification of investments*. Yale University Press, Martin, M. K., Gonzalez, C., & Lebiere, C. (2004). Learning to make decisions in dynamic environments: ACT-R plays the beer game. In M. C. Lovett, C. D. Schunn, C. Lebiere & P. Munro (Eds.), *Proceedings of the Sixth International Conference on Cognitive Modeling* (Vol. 420, pp. 178–183). Pittsburgh, PA: Carnegie Mellon University/University of Pittsburgh: Lawrence Erlbaum Associates Publishers.

McKenna, F. P, & Crick, J. (1991). Experience and expertise in hazard perception. In G. B. Grayson & J. F. Lester (Eds.), Behavioral Research in Road Safety (pp. 39–45). Crowthorne, UK: Transport and Road Research Laboratory.

McLeish. D. L. (2005), *Monte Carlo Simulation and Finance*. John Wiley and Sons, Hoboken, NJ.

McLeish. D. L. (2011), *Monte Carlo Simulation and Finance*. John Wiley and Sons, Hoboken, NJ.

Mees A.I., (1981) "Dynamics of Feedback Systems", New York: J. Wiley, ISBN 0-471-27822-X. p69

Mooney, C. Z. (1997). Monte Carlo Simulation (Sage University Paper series on Quantitative Applications in the Social Sciences, series no. 07-116). Thousand Oaks, CA: Sage.

Rachev S.T., Ortobelli S. and E., Schwartz, (2004). The problem of optimal asset allocation with stable distributed returns, in Volume 238 (Alan C. Krinik, Randall J. Swift eds.) *Stochastic Processes and Functional Analysis* Marcel Dekker Inc., 295-347.

Ramaprasad, A., (1983), "On The Definition of Feedback", Behavioral Science, Volume 28, Issue 1.

Rapoport, A. (1975). ‘Research paradigms for the study of dynamic decision behavior’. In: D. Wendt and C. Vlek teds.). Utility, probability and human decision making. Dordrecht: Reidel.

Rigas, G., Carling, E., & Brehmer, B. (2002). Reliability and validity of performance measures in microworlds. *Intelligence, 30*(5), 463–480.

Samuelson, P., 1965, “Proof that Properly Anticipated Prices Fluctuate Randomly”, Industrial Management Review 6, 41-49.

Thaler, R. H. (1980). "Towards a positive theory of consumer choice", *Journal of Economic Behavior and Organization*, 1, 39-60

Tobin, J., (1958). "Liquidity preference as behaviour towards risk," Review of Economic Studies, 25, 65-86.

Turkle, S. (1984). *The second self: Computers and the human spirit.* London: Granada.