Can Learning Classifier Systems Represent Competent Traders?

The Stock Markets Trading Case

Sonia Schulenburg

Doctor of Philosophy
Artificial Intelligence Applications Institute
Division of Informatics
University of Edinburgh
July 2003

Abstract

I seek an understanding of the dynamics: learning and evolution, of certain groups of artificially created agents – known here as *trader-types* – making decisions in a real stock market scenario. With this as the primary motivation, a specific problem has been devised and three different groups of agents have been modelled to learn, forecast and trade in a real stock market scenario given exogenously in the form of easily-obtained stock statistics such as various price moving averages, first difference in prices, volume ratios, etc. These artificial *trader-types* trade and learn simultaneously during – in most cases – a ten year period. They start with no prior knowledge about the market, i.e. they have no notion of what is a good or a bad approach to start with; all their market models are created randomly at the beginning of the period, with the idea that new models developed through experience will be formed and polished as time progresses. The life of such *trader-types* commences when they are given an initial wealth to trade over two assets (a risk less bond represented by the fixed interest rate given by the bank and a real risky stock) and ends in most cases after one decade.

First, in this problem I try to explore whether it is feasible to represent with Learning Classifier Systems (LCS) some of the key elements that play a role in the decision-making process of real stock market traders when viewing it from an evolutionary framework. Specifically, two fundamental questions are addressed under this first and broad topic: Are the *trader-types* able to (i) evolve and (ii) behave in similar ways to human traders under the real market conditions described above? Here the work is concentrated in LCS as the learning approach and in viewing the agent as part of a process where adaptation to a partially understood market environment is a necessary element for survival to occur.

Second, this thesis reports on a number of experiments where the forecasting performance of the adaptive agents is compared against the performance of the *buy-and-hold* strategy, a *trend-following* strategy, a *random* strategy and finally against the *bank* investment over the same period of time at a fixed compound interest rate. To make the experiments as real as possible, agents also pay commissions on every trade. The

results so far suggest that this is an excellent approach to make trading decisions in the stock market and that continual learning and adaptation not only play an important role but are also necessary elements in the decision-making process.

Third, the concept of continual learning is addressed, and to show how the model constantly adapts to new market behaviour, additional experiments which include a number of real stocks are presented, followed by a discussion section.

Acknowledgements

During the time in which I have developed these ideas I will present to you shortly there have been three very special people to whom I am extremely thankful: my husband, my supervisor and my friend.

To my husband Mark Calvert, for having such great faith in me and all I do. For his great support and strength. For wakening me up early in the mornings with a delicious capuccino, for bringing me to work on my thesis and for picking me up after work so many endless days.

To Peter Ross, *the* greatest supervisor, techer and mentor. For always being there when I needed him, for listening so patiently and for helping me understand what I was trying to say. For all the wise advice he gave me in all areas discussed throughout these incredible years and for making my PhD one of the most enjoyable times I have ever had.

To a special friend who shared with me such great moments and who showed me a great number of things about life, happiness and contentment. For showing me the joy of smiling without worrying when I needed it the most. Yes, to Enrique Filloy, my dearest friend.

This work has been partially funded by Consejo Nacional de Ciencia y Tecnología and Fundación Mexicana para la Educación, la Tecnología y la Ciencia, A.C. Many thanks for their financial support.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

To My Husband

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

> "How did reason come into the world? As is fitting, in an irrational manner, by accident. One will have to guess at it as at a riddle." Friedrich Nietzsche, *The Dawn*.

The main focus of the research presented in this thesis lies in the areas of economic modelling when considered from an evolutionary framework, and it is concerned with the process of creation and evolution of artificially-intelligent agents living in a real economic world. Therefore, the scope of the work presented here can be viewed as the integration of two ample categories: Economics and Artificial Intelligence. The area of economics will be covered from its most traditional hypothesis to its more modern approaches, adding to these a number of real and artificial applications. Regarding artificial intelligence, emphasis will be given to the more modern approaches, especially to those that involve applications of an economic nature. A number of different approaches discussed in this thesis include neural networks, fuzzy logic, genetic programming, genetic algorithms and classifier systems.

In this work, these two broad areas fuse together to create a simple economic setting in which different classes of predictive agents make decisions based on real facts about the market which are given exogenously to them. Throughout this process, two interesting factors give rise to new possibilities of exploration: learning more about (i) the characteristic behaviour of the market analysed, and (ii), about the dynamical properties of the *agent-types* being modelled.

In connection to the first factor, a number of artificial markets have been developed with the purpose of increasing our understanding of how certain properties of actual markets, such as bubbles and crashes, originate. This has been achieved through the development of artificial markets, which is one of the main components of this model. But regarding the second factor, unfortunately not much has been achieved in the development of new ways to explore the dynamic properties of the agents under such market conditions. For example, how these agents evolve and behave over time in an uncertain, real market environment? How can they learn – or even survive – from a world when they can not have control over? Can they improve and be proficient? If they are indeed proficient, is it because of good luck or because they did find some market inefficiencies that they took advantage off? These types of questions still remain unanswered and it is the aim of this thesis to try to address them.

Figure 1.1 describes the model when viewed from a global perspective, indicating how the elements that affect the system are interconnected. As it can be seen in the illustration, real *financial markets* are affected by world news of various types, such as *environmental, economical* and *political*, among many others. This, in turn, affects the *stock market*, which is composed of real market participants (for simplicity, called *traders* throughout this thesis) who affect the *market conditions*. Market conditions reflect the dynamics of specific stock prices, volume of transactions, etc., which are observed at any given time. As indicated by the straight arrow, a subset of these market conditions is fed into the model, which is composed of different types of *artificial traders*, called *trader-types Tt1, Tt2* and *Tt3*.

Some information from the market environment can be shared by the *artificial traders* (illustrated by the intersection between them). For instance, they all receive the difference of the current price with respect to the previous price (whether the price today is higher than yesterday's price). Other bits of information are specific to each *trader-type*. Such *trader-types*, with no prior knowledge about the market, start building and improving a pool of market hypothesis while analysing the market information they receive. This process (illustrated by the curved bi-directional arrow) will be explained in detail in Chapter 7. Their behaviour is then observed along with a set of explanations of their actions and several conjectures about real traders, the market and the continual leaning process involved can be drawn. These will be explained in Chapters 8 and 9.

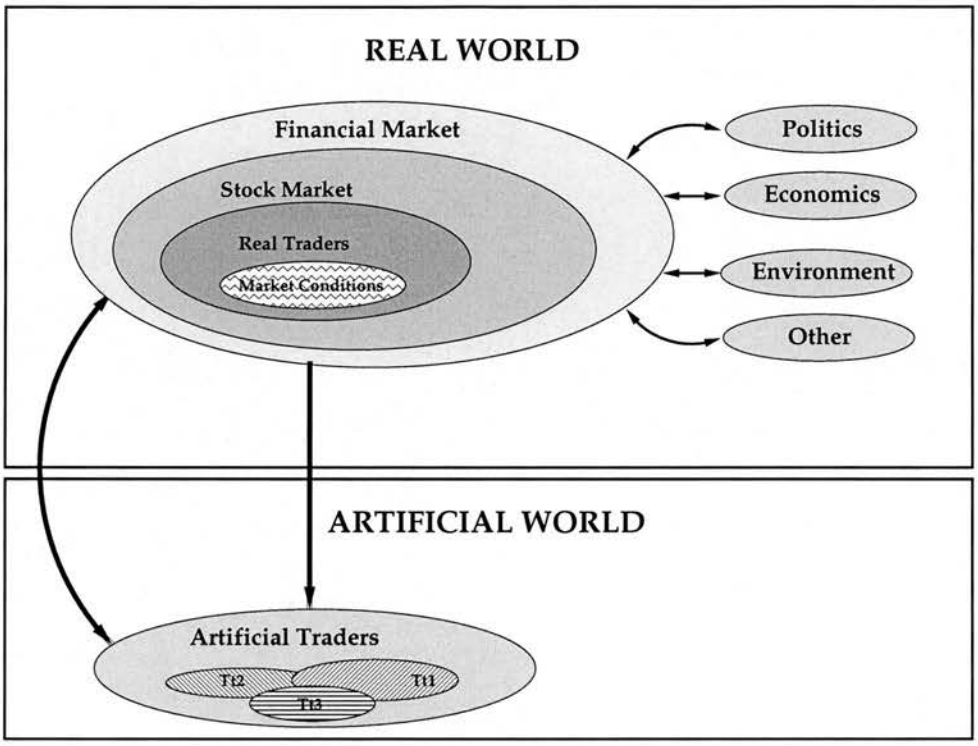

Figure 1.1: Diagram of Key Elements and Interconnections of the Proposed Model

Figure 1.2 shows the main elements that the proposed model deals with. The flow of activities goes as follows:

1. At the beginning of period of time t_i, the three *trader-types* receive simultaneously certain market information which they process accordingly.

2. At the end of t_i their actions are posted back into the environment.

3. With a delay of one period (or generation), at t_{i+1}, the traders receive a reward from the environment if certain conditions were met, otherwise they do not receive positive reinforcement.

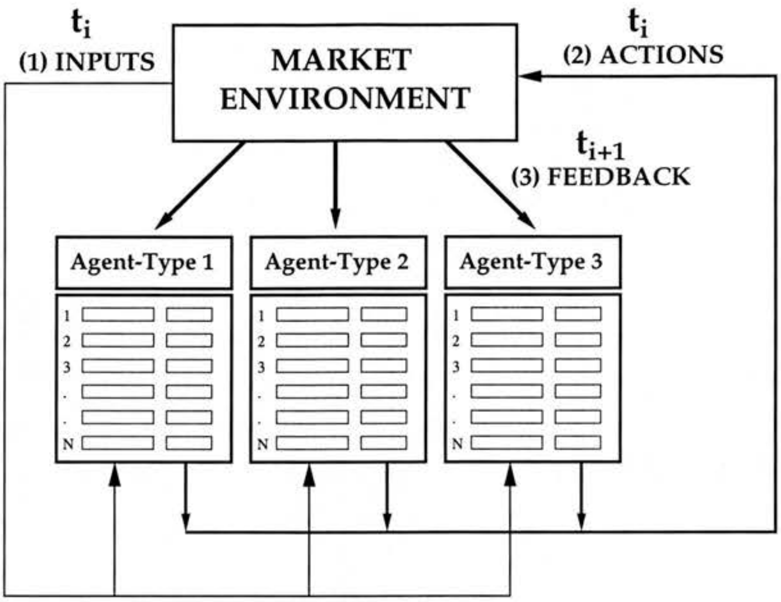

Figure 1.2: Flow Diagram Showing Interconnections Between the Elements of the Model

With the following two purposes in mind: (i) to develop better trader models and (ii) to increase our understanding and actually learn from these adaptive agents, an important deviation from the *artificial market* framework has been proposed here, one that takes *reality* into account. In other words, the model can be viewed as an artificial stock market, except that in this model the financial information is provided exogenously, therefore the actions of the *trader-types* do not affect the real world. Although it might seem awkward to propose an artificial market model that is partially real, there are several reasons for doing it this way.

Firstly, the attempt to provide better methods for learning about the evolutionary process of agents of this type, at the same time as learning about the market dynamics is the main interest of this research. Secondly, because of the use of real stock prices, in addition to these goals, the model can also be viewed as an on-line, automated trading system; which in fact, presents encouraging forecasting capabilities when considered to serve this purpose. This aspect is also addressed throughout this thesis, as they are all linked properties of the proposed system.

A striking question can arise after making these statements: On what grounds could one develop such a *semi-artificial* market model? Shouldn't models be contained within the boundaries of one framework or another, but not sharing bits and pieces from both worlds? To answer this question partially, it can be said at this point that the model proposed here was developed intentionally this way, and the reason is simple: to allow full concentration in the dynamics of the agents rather than in the global effects they have on the market. Artificial markets research focuses on the study of the time series generated to later be able to *infer* knowledge about the agent – that is, if it displays properties found in real markets. The view here is to perform this process directly via the agent, not through the artificially-generated market. Therefore, the agents of this model are not allowed to change the price sequence. This way, the ones that constantly change are precisely the agents, which according to the argument this thesis supports, provides a greater potential for learning more about these agents, which can perhaps even prove to be a useful method to provide important insights about real markets and traders, i.e., How do they behave, learn, operate, and adapt?

In this context, the aim is to analyse the dynamics and evolution of a simple economic system with a small number of decision-making agent-types and to see how this performs using real data – are such agent-based systems capable of trading successfully over any period of time? This is necessary, because in the real world, only human traders that are reasonably successful will stay in business.

There are some key elements of novelty identified in this model:

- *Trader-types* are modelled rather than individual traders, under the hypothesis that this is more plausible and tractable than trying to model an individual trader. Specifically, three different types of agents have been modelled for comparative purposes, but any number of *trader-types* can easily be implemented.

- There is no separate training and testing phases, this is a true on-line learning model in which market data is never seen more than once by the artificial traders, therefore there is no *in-sample* data: all learning is performed with *out-of-sample* data, following the second hypothesis of this thesis which states that learning in such environments should be viewed as a continual process and therefore there should not be a training phase with the unique purpose of testing in the future, the success of the strategies learnt in the past. Strategies continuously change and should not be frozen in time, nor should agents be tested on strategies they have previously developed successfully – and perhaps memorised – from being exposed to the same past data a great number of times.

- In this model real market and trader behaviour can be analysed from the evolved sets of market strategies in a variety of ways. The relevance of certain bits of information used by the agents during the decision-making process can also be tested in order to rank their level of importance against each other and to develop better *trader-types*.

1.1 The Analogy

In the work presented here I propose the use of learning classifier systems as a way of representing the cognitive process of financial traders living in a simplistic real stock market. The idea is to examine whether it looks promising to model real traders this way, which has been inspired by an intrinsic beauty of LCS which allows adaptation and learning coexist in two different contexts: the first one, which I call the "micro-adaptation" level, happens at a short time interval, through positive reinforcement of the good strategies in a continual process of adaptation to an environment which is completely unknown initially. In this scope, the agent can be seen as being in a sort of *do the best with what you've got!* mode. The second level, which I call "macro-adaptation", offers a more powerful evolutionary component where discovery of new, invented strategies, while maintaining the diversity of the pre-existing mixture of strategies, is essential in order to "view" the entire process from a global perspective. Here, the agent can be viewed in a ... *but keep always looking for different alternatives* mode.

I want to illustrate this intrinsic property that classifier systems offer with the following analogy, which crossed imagination while writing this thesis. Imagine you are walking in an infinitely large, dynamical virtual museum and you able to see only what your eyes can capture while walking very quickly through many corridors. Imagine, and imagine more of this. You are trying to locate, for example, only paintings of Kandinsky. But on the way you notice that there are millions of other beautiful paintings, including Cezanne, Klimt, Bouguere, Matisse, Monet and Van Gogh, among many others, located one after another. Whether they are organised in some specific ways or just placed randomly, one does not know. However, one's mind immediately tends to learn limited patterns of short distances from the angle where the eye is observing (such as *after* turning right in a door like this I will see a painting like that... and so forth). This type of processing corresponds to the micro-adaptation context. But at the same time an inevitable thing happens: one is always creating and trying new theories, as if one was constantly drawing and redrawing full diagrams of the shape of the entire building as seen from a higher level, one that suggests complementary – and perhaps different – ideas about where to find such paintings. One's mind is constantly connecting corners and rooms with one another at an imaginary level in order to form the "macro" perspective of the domain. The things learned here, could suggest, for example, that Kandisnky's paintings are never found in the north-west part of the building, where Bouguere's more angelical paintings appear to be more common at that specific instant in time.

With these two levels of adaptation working together, one is able to anticipate what the eye will see next and test whether that was right or wrong. It is not known whether there exists a unique governing mechanism arranging the paintings dynamically in certain ways, but at that high speed patterns are changing so quickly, that the idea of its existence appears almost impossible to conceive. One could think that perhaps there is more than one governing law behind this. Could it be that there are millions of formulas placing the paintings at a given moment in time? Will we ever be able to find them – or at least some – quickly enough to be able to use them to anticipate the next event? But perhaps, in this jungle of the unknown it is enough to be able to draw even the tiniest clues. How many times does one need to anticipate right where the next painting will be? Perhaps indeed, there are small clues to be found in what could appear as an amorph arranging technique.

Financial markets are far more complex to model than this analogy could suggest. It is possible that after reading this thesis, the reader's number of questions will perhaps be the same as it was before – perhaps even greater. But it should not remove any motivation to continue with the reading, as it is not the number of questions addressed what matters, it is the quality of the questions, which will gradually increase in granularity. This thesis can not answer all the questions addressed, the attempt is only to address issues to help us advance in our thinking about how to model complex systems. From this broad an ambitious perspective, the motivation is to try to contribute to our understanding (and highly controversial views) of how intelligent systems build internal models of their surroundings to improve their existence. Specifically, addressing the question: Can we model this kind of behaviours with a LCS? Recall that the goal here is not to get *the* best performance minimising computational complexity, but instead, it is to explore whether LCS exhibits a good repertoire of behaviours we humans display in order to face common problems of our daily life. The idea of using classifier systems to model adaptive systems seems very powerful to me, and they appear to be well suited to handle more complex systems that we originally thought they would be capable of. In these problems we are looking for possible explanations of behaviours that seem to be unique and quite simple in humans, but impossible or extremely difficult to describe using the exact sciences. The hope is that this approach could lead us to new insights about modelling complex adaptive systems.

1.2 Thesis Outline

General backgrounds in Economics and Artificial Intelligence are described in Chapters 2 and 3, where traditional theories related to the work presented in this thesis are explained, as well as how such theories emerged into new, more dynamical and flexible ideas which involve evolution, the primary component of this work.

In addition to the basic theories behind the model, other areas, which are not directly linked to this model, are also explained, in order to make it easier for the reader to understand descriptions that will be given throughout the thesis about other models and systems that have been developed, and which also share common elements with this model; for instance, while describing the current state in the areas of computational economics and forecasting.

For this purpose, Chapter 4 describes various approaches to computational economics, which try to tackle different problems that this thesis deals with. This Chapter starts with an introduction to evolutionary economics, followed by a description of two adaptive, artificial market approaches, whose main goal is to focus on market dynamics through the analysis of artificially created data. Finally, this chapter describes a number of predictive models, which have been mainly developed to specifically tackle the forecasting problem.

Chapter 5 gives a brief introduction to the financial issues involved in the proposed model, followed by Chapter 6, which describes a number of artificial intelligence applications in finance.

The model is fully described in Chapter 7, followed by Chapter 8, which shows the system's performance with a number of experiments of different stocks in the UK and the US. The concept of continual learning and a further analysis of the process of adaptation is exemplified with more stocks in Chapter 9. Finally, a summary of the analysis, conclusions and further line of research is provided in Chapter 10.

Chapter 2

Economics Background

> "[Anything which] is a living and not a dying body... will strive to grow, spread, seize, become predominant - not from any morality or immorality but because it is living and because life simply is will to power... 'Exploitation'... belongs to the essence of what lives, as a basic organic function; it is a consequence of the will to power, which is after all the will to life." Friedrich Nietzsche, *Beyond Good and Evil*, translated by Walter Kaufmann.

The basic theme of the problem devised for investigation in this thesis is concerned with the way an individual or society chooses to employ certain resources in order to survive, satisfy a desire or maximise some known quantity. In simple terms this is what economics is all about: **economics** studies how people choose to use limited resources to obtain the maximum satisfaction of unlimited human wants. It is, therefore, concerned with the description and analysis of the production, distribution, and consumption of goods and services.

While trying to study the way individuals and societies make rational decisions when confronted with scarce resources and an uncertain environment, the broad field of economics has been divided into two major disciplines: **microeconomics**, which deals with how individuals and firms make decisions within the context of an isolated market, and **macroeconomics**, concerned with issues such as how the economy as a whole behaves over time.

In parallel to this line of thought, because the problem domain described in this thesis is of such a broad nature as all economics is concerned, a decision had to be made regarding the direction of where to concentrate the efforts to narrow the number of elements explored. Therefore, in a "microeconomic" sense, the purpose of this work is mainly to focus on the way choices are made by agents by looking at these directly from the agent's point of view rather than indirectly via the analysis of the market properties created from such decisions. With this view in mind, the central theme is related to establishing a human-like decision-making process which revolves around a real and uncertain world driven by humans.

Looking more deeply into the chosen subject of this work, one discovers the need to address the greatly debated question of whether it would be possible to forecast price movements successfully in the short to medium-term time horizons. A reasonable (and common sense) answer would suggest that if a certain model is capable of capturing accurately some market inefficiencies, it would then be able to generate excess returns. But what does economics have to say about these issues? Would this be possible?

2.1 Efficient Markets Hypothesis

Classic economics, in the form of the Efficient Market Hypothesis (EMH), clearly answers no to this question by stating that markets are assumed to be efficient if all available information is reflected in current market prices [Fama 70, Fama 91]. But what exactly this means is a question that in many respects still remains unanswered.

EMH evolved during the 1960s from the publication of the famous PhD dissertation of Eugene F. Fama at the University of Chicago. His thesis, entitled "The Behavior of Stock Market Prices" and originally published in the Journal of Business [Fama 65b], caused a tremendous and almost instant impact in academia. In addition to this work, Fama also published ideas of this kind in a number of journal papers, such as "Random Walks in Stock Market Prices" which appeared in, among others, the Institutional Investor [Fama 65a], where the following excerpt is taken:

> "An 'efficient' market is defined as a market where there are large numbers of rational, profit-maximizers actively competing, with each trying to predict future market values of individual securities, and where important

> current information is almost freely available to all participants. In an efficient market, competition among the many intelligent participants leads to a situation where, at any point in time, actual prices of individual securities already reflect the effects of information based both on events that have already occurred and on events which, as of now, the market expects to take place in the future. In other words, in an efficient market at any point in time the actual price of a security will be a good estimate of its intrinsic value."

The great impact of EMT dispersed rapidly from academia to the investment community. Fama's 1970 paper on the topic "Efficient Capital Markets" [Fama 70] argues that on average, it is nearly impossible for an individual to consistently beat the stock market as a whole because of the broad availability of public information. Some people think this is equivalent to saying that an investor who throws darts at a newspaper's stock listings has as much chance at beating the market as any professional investor. In summary, the statement that stock price behaviour is random and not predictable by the so-called market forecasters had become a dominant paradigm used by economists to understand and investigate the behaviour of financial markets.

Looking for more evidence to sustain this theory, Fama studied all the funds that survived during a 20-year period starting in 1976. But because these funds were naturally biased from being survivors during the period, he reduced the sample by choosing only the 20 biggest winners from the first ten years and analysed their performance over the second 10 years relative to a risk-corrected model. Not surprisingly for him, he found out that exactly half were above average and the other half below. This means that the best performers had a 50% chance to succeed over the next period. With this new evidence he continues to favour passive management[1] over active management[2].

[1] Traditionally, investment managers have been categorised as active or passive types according to their trading profiles. Passive Management, also known as *index* management, focuses on the belief that people cannot get more than the market rate of return from the given category they are in because security prices are the best estimate of value, therefore no effort is needed to distinguish between one security over another to try to 'beat the market'. Portfolios are created to resemble the performance of well-known indexes and price adjustments are made in response to changes in the underlying universe or index.

[2] Active management states that there are people who can make valuation judgments that are superior to the market; portfolios are constructed by using a variety of strategies which are believed to offer excess returns. There are more costs associated with this investment strategy than with passive investment because portfolios are usually more dynamic.

But still, for those who do not follow the EMH, there were some winners over the whole period. (The biggest winner of all was Fidelity Magellan, credited to Peter Lynch.)

Still there are some open-end questions surrounding these concepts because EMH does not explicitly say that **you** cannot make money in the stock market (just the average of individuals cannot) or that it does not matter what you invest in – you will earn the same return in any case (again it is *THE* average investor). It doesn't say either that flipping a coin or throwing darts is an equally good method as any other for selecting stocks. It is implied that some methods will be good and others bad.

2.1.1 Types of Market Efficiency

Academics and financial professionals have debated this issue for a number of decades, conducting mainly three types of efficiency tests: weak, semi-strong and strong. Most of these studies conclude that the major financial markets are efficient and that all information is reflected in current prices. However, certain methodological questions addressing whether the observed departures from market efficiency are due to any genuine market inefficiency or due to a deficiency of the market pricing model being used as a yardstick to compare actual with theoretical prices [Olsen *et al.* 92] have weakened such conclusions.

The weak tests investigate whether market prices actually reflect all available information. These tests support the **"Weak form"** of efficiency, which asserts that in a market that assimilates information efficiently it is impossible to predict the future price of a security on the basis of its **past price**, so a weak-form efficient market is one in which past security prices are impounded into current prices and where technical analysis is of no use. This means that because asset prices fully reflect the information contained in the historical sequence of prices, there cannot be investment strategies which yield abnormal profits on the basis of past prices. All past market prices and data are fully reflected in security prices. Past stock patterns cannot be used to make extraordinary profits.

An argument in favour of this form of efficiency says that a short-term market-timing rule cannot make money because the opportunity goes away if everyone follows the same strategy. So the belief is that markets do not follow major patterns in stock prices, and the minor patterns caused by momentum are costly to explote. Evidence suggests that prices do appear to be random.

The **"Semi-strong"** tests analyse the degree of market reaction to events based on **public information**, which includes news announcements, annual reports, gossip columns, news clippings, etc. A market that impounds all of this information in its current price is known as semi-strong efficient. Most people believe that the U.S. equity markets by and large reflect publicly available information [Goetzmann 01]. But the question of what is considered public information is still unanswered: is the information obtained through the Internet considered to be of public nature? Is ALL of this already impounded in the stock price? If so, at what rate was this information discounted in prices? A common guess is that the easier it was to get certain type of information, the more likely it is to have already been traded upon. Therefore in this view it is impossible to predict on the basis of publicly available fundamental information. Evidence supports that earnings and dividend announcements are incorporated into stock prices within 5 minutes.

Note that since past prices are considered public information as well, weak form of efficiency could imply semi-strong form of efficiency and vice versa.

Against these two types of efficiency are a great number of experienced financial professionals and financial groups who usually take positions from public information, even though there is no proof that they can actually beat the market. Such evidence would be impossible to obtain for obvious reasons, but the fact is that it is hard to believe that they do not outperform the market somehow.

Finally, the **"Strong form"** of efficiency analyses whether investors have **private information** to take advantage of. Private information includes insider information such as a personal note passed from a CEO to the CFO regarding a major financial decision, which according to this form of efficiency, would suddenly have an impact in the stock price! Not many people believe that the market is strong-form efficient. In this view it is, therefore, impossible to predict from any information at all.

Evidence suggests that professional investors cannot beat the S&P500 consistently, but company insiders do appear to make abnormal profits when they trade their own firm's stock. Strong-form efficiency is not well supported, managers seem to know things about their firms that the stock market does not know.

To illustrate whether this form of efficiency exists in real markets, refer to figure 2.1. The graph shows the changes in stock prices due to an event such as a merger between two companies (example taken from [Goetzmann 01]). According to this type of efficiency, insider information will cause a rise in prices due to the fact that some people will start trading (even though it is illegal) before the public release of the event. Unfortunately prices typically do rise prior to a merger, indicating that there could be some tendency towards strong-efficiency. This could be due to private information being incorporated into prices, but also because of some other shared cause: that is, both the price rise and the merger may have been triggered by some events or circumstances pre-dating either. But surprisingly, the trend continues to go up after the merger is publicly announced, suggesting that the market is not strong-efficient.

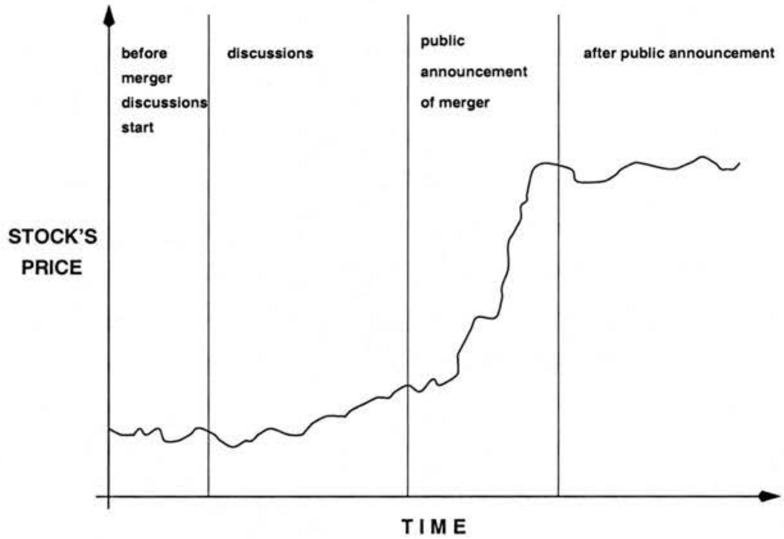

Figure 2.1: Effect of stock prices close to the announcement of a merger

If at any given time security prices fully reflect all available information, the implications of the EMT are profound. It means that if markets are efficient and current prices fully reflect all information, then buying and selling securities as an attempt to outperform the market is just a game of chance rather than skill.

To recap all these views of market efficiency, security prices adjust rapidly to the arrival of new information, which comes to the market **randomly** and prices adjust rapidly to reflect the new information. Price adjustments are imperfect, yet unbiased. But does this work in real markets? Are they efficient? Is the movement of the price really unpredictable?

2.1.2 Random Walk Hypothesis

As previously stated, EMH has become the dominant paradigm used by economists to understand financial markets.

Let's consider the following equation taken from [Zulauf & Irwin 97] which underlies the major concepts of the EMH:

$$P_{t+1} = \beta P_t + \alpha + \varepsilon_t, \tag{2.1}$$

where P_{t+1} is the stock's price at time $t+1$, P_t is the current price, α and β are parameters and ε_t is a random error term which is independently and normally distributed with mean 0 and constant variance σ^2.

To aid in understanding EMH, equation 2.1 is rearranged as follows:

$$P_{t+1} - \beta P_t = \alpha + \varepsilon_t \tag{2.2}$$

If $\alpha = 0$ and $\beta = 1$, then

$$P_{t+1} - P_t = \varepsilon_t \tag{2.3}$$

Taking the expectations from this equation,

$$E_t(P_{t+1} - P_t) = 0 \tag{2.4}$$

The price process described above is usually referred as the **random walk** [Campbell *et al.* 97, Tomek & Querin 84]. This means that the expected average change in price is zero and the best guess of price at time $t+1$ is the price at time t given the information set available at time t. Furthermore, since ε_t's are uncorrelated, changes in prices are uncorrelated as well. A commonly used analogy for the random walk is the flipping of a coin.

The principle of modern finance regarding higher return for higher risk is consistent with Fama's market efficiency by allowing a price bias, so equation 2.4 becomes:

$$E_t(P_{t+1} - P_t) \neq 0, \tag{2.5}$$

and provided that the bias is α, the compensation for risk $\alpha \neq 0$. This variation of the RWH is referred as the **random walk with drift**. In Keynes terms [Keynes 30], an $\alpha > 0$ implies normal *backwardation*, i.e. the expected price is lower than the realized price so future prices should increase over the course of a contract, resulting in positive returns to a long position. In a *contango* the opposite applies and $\alpha < 0$. Here the expected price is higher than the realized price and a short futures position will earn positive trading returns.

As pointed out by Eugene F. Fama in an interview by Peter J. Tanous [Tanous 97], it is important to note that EMH and RWH are not the same thing: the first one being much more powerful than the latter one, which only states that future price movements can not be predicted from past price movements alone.

Yet there is evidence strongly suggesting that the stock market is not a random walk as originally thought, it is believed that it tends to trend in one direction for too long at each time. Another idea against RWH is that some people reinforce trends by not buying until they see a price trending upwards to confirm their instinct to buy that asset, confirming the idea that "the more people share a belief, the more the belief is likely to be true." Therefore the behaviour of people second-guessing the expectations of others produces a self-fulfilling prophecy [Ridley 93]. In this sense people generate price movements, changing the dynamics of the market in ways that are not random at all.

The self-fulfilling prophecy works as follows: when it is predicted that the price of a stock will rise, it produces a purchase, which, as a consequence, will cause the price of the stock to rise. Such prophecies can deliver some elements of predictability that could be captured by some systems in practice. From this follows that the price behaviour of each day depends, up to a certain point, on the price of previous days, refuting the idea that prices follow a random walk.

2.1.3 Reactions to New Information

According to efficient markets, the reaction to news would immediately be adjusted and the stock market would fully reflect the new information. There would be no tendency for subsequent increases and decreases. To illustrate this, refer to the merger example given under strong-form market efficiency in section 2.1.1.

Figure 2.2 shows the effect of under and over-reaction to the new information (taken from [Viswanath 01]). The solid line indicates the efficient market reaction, which is characterised by the instant increase or decrease in stock prices, fully reflecting all new information in such a way that no further adjustments need to be made. The apparent delayed reaction supports that the price adjusts only partially to the new information; in fact, usually several days elapse before the price completely reflects the new information. The over-reaction occurs when the price over adjusts to the new information creating a bubble in the price sequence which then corrects itself to the new price.

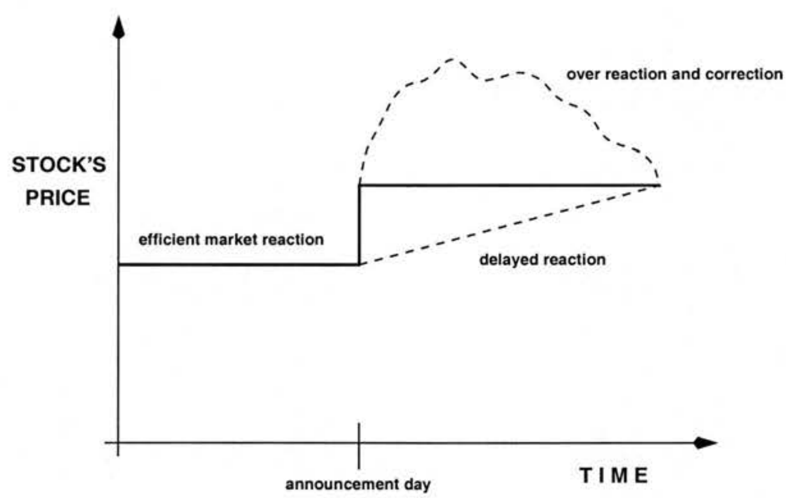

Figure 2.2: Effect of stock prices close to a public announcement.

There seems to be some evidence that investors over-react. Is it possible for the over-reaction (and perhaps even the delayed reaction) sequence to be efficient? In recent studies, Fama maintains that the evidence of the many long-term return anomalies that has been produced in the economic literature over the past three decades does not suggest that market efficiency should be abandoned; he still holds the same position in favour of EMH, as well as many economists and even financial institutions. There is also a growing body of evidence suggesting the opposite view.

In a recent paper, winner of the Fama-DFA Prize [3] for the best asset pricing paper in the Journal of Financial Economics, Fama points out the following:

> "Consistent with the market efficiency hypothesis that the anomalies are chance results, apparent over-reaction of stock prices to information is as common as under-reaction. And post-event continuation of pre-event abnormal returns is about as frequent as post-event reversal. Most important, the long-term return anomalies are fragile. They tend to disappear with reasonable changes in the way they are measured." [Fama 98]

Fama has also addressed the issue of under and over-reaction. In [Tanous 97] he explains it with the analogy of the market crashes of 1929 and 1987. He claims it is not that markets were not efficient in those times, prices were correctly priced previous to the crashes, but as sometimes prospects of growth do not go as expected, there are big declines. He adds:

> "the same thing can happen to the market as a whole. It can also be a mistake. I think the crash in 87 was a mistake. Take the previous crash in 1929. That one wasn't big enough. So you have two crashes. One was too big and one was too small.... Half the time the crashes should be too little, and half the time they should be too big..... Think of a distribution of errors. Unpredictable economic outcomes generate price changes. The distribution is around a mean -the expected return that people require to hold stocks. Now that distribution, in fact, has fat tails. That means that big pluses and big minuses are much frequent that they are under a normal distribution. So we observe crashes way too frequently, but as long as they are half the time under-reactions and half the time over-reactions, there is nothing inefficient about it."

2.1.4 Impacts and Controversies surrounding EMH

The following are some of Fama's original assumptions:

[3] The Fama-DFA Prize is an award funded by Dimensional Fund Advisors (DFA), given in recognition of the contributions the Journal of Financial Economics and Professor Fama have made in the areas of capital markets and asset pricing.

1. No transaction costs
2. Costless information
3. The implications of current information for both current price and the distributions of future prices are generally accepted by all market participants.

From these assumptions, we can see that at least (1) and (2) are unrealistic. First, when buying and selling stocks, transaction costs do exist in at least two forms, the commissions and the bid/ask spread. With the widespread use of the Internet commission costs have generally dropped over the last decades, but even at present they occur. Online brokers and investors can claim that commissions for relatively small transactions (such as blocks of 5,000 shares or less) are in the range of $15.00 U.S. dollars. Of course this varies from one online broker to another one and other rules may apply. And second, information is costly as well. Acquiring it usually involves a service fee and analysing it is another added expense that varies greatly.

Taking such costs into consideration, the idea is that in an efficient market the gross trading returns shall not exceed transaction costs. Although in a limited way, researchers have started to evaluate this market efficiency criteria by taking into account transaction costs such as commissions in their economic models. In other economic models such as the double auction models, bid and ask spreads are also being considered. These will be explained in more detail in Chapter 4.

2.1.5 Grossman and Stiglitz Model

The Grossman and Stiglitz model, also known as the **noisy rational expectations** model, addresses the problem of costly information more realistically than the EMH. In this model [Grossman & Stiglitz 80] information is costly, therefore it is impossible for prices to perfectly reflect all available information. Here, the constant β from equation 2.2 of section 2.1.2 will not necessarily equal 1 (as in EMH), allowing the market to analyse information slowly. In this sense, the market "learns" while adapting the incoming information.

The model allows for the possibility of profitable trading by using and analysing certain information to take a position in anticipation of price changes that will occur as the rest of the market learns about that piece of information. The profits obtained would cover the costs of acquiring and analysing such information. Therefore those who obtain the information faster or those who possess superior analytical abilities can earn positive returns.

This assumption is based in the fact that markets are human institutions that need to learn new information. Zulauf and Irwin point out that large traders are immersed in national and international information flows because they have access to more resources than small traders [Zulauf & Irwin 97]. According to Grossman and Stiglitz's model, these advantages imply that large traders should make most of the money from trading in futures and options markets. Therefore small traders should lose money as a group due to their inability of being the first ones in acquiring new information.

Zulauf and Irwin cite extensive studies over a number of markets which are consistent with Grossman and Stiglitz's model in [Zulauf & Irwin 97]. For example, Hartzmark analysed trader's returns of futures data from nine different markets from 1977 to 1981. In this study, large traders (hedgers and speculators) made over $850 million while small traders lost $853 million [Hartzmark 87]. In other markets, Leuthold, Garcia and Lu [Leuthold *et al.* 94], Phillips and Weiner [Phillips & Weiner 94] and Irwin, Krukemeyer and Zulauf [Irwin *et al.* 93] attribute the significant profits of larger traders to superior information and/or forecasting ability.

Unfortunately, there is mixed evidence in regards to whether information indeed aids in a consistent manner to permit superior forecasting ability. Excess profits due to advantages in information involve costs related to obtaining it and analysing it. Different views will be addressed in section 2.4.

2.2 Rational Expectation Theory

A considerable body of economic theory and research centers on the notion of "rational expectations," a classical economic theory formalised in the 1970s by Nobel laureate (1995) Robert E. Lucas, Jr. from the University of Chicago, who is believed to have the greatest influence in macroeconomic research since 1970.

The basic assumptions of the Rational Expectations Theory are:

1. Economic actors (such as consumers or investors) behave rationally by collecting and studying carefully current and historical conditions.

2. Markets are highly competitive

3. People act in response to their expectations

The idea that in an efficient market a trading model cannot generate any excess returns is based on the assumption that all investors act according to the rational expectations model. Broadly speaking, the Rational Expectations Theory (RET) suggests that, in a stationary world, if all investors have the same data, they will all entertain the same expectations and those expectations will be true (or rational). The theory suggests that most participants in the stock market are "smart investors," whose individual expectations, on average, anticipate the future correctly.

2.2.1 Limitations

While there is no doubt as to the merit of RET in explaining many economic relationships, one must admit that at times it appears out of synchrony with observed behaviours. The theory has two immediate limitations: first, we do not live in a stationary world but in one that is subject to structural changes; and second, all investors will not necessarily reach the same conclusions, even when acting under the same observables. In addition to these, there are other factors such as emotions which often seem to play a significant role in economic decisions by bringing a certain degree of irrationality to our true perceptions. If the assumption of rational expectations is wrong to begin with, then the validity of its conclusion in the form of the EMH could also be questionable.

Clearly, Rational Expectations Theory cannot explain everything that happens in investment markets, particularly the large distortions and swings in security prices seen during recent years. For this reason, other versions of the RET have also been addressed, such as the Rational Beliefs Theory, developed by Mordecai Kurz of Stanford University which will be explained in the following section.

2.2.2 Rational Beliefs Theory

This new theory expands the RET to cover other situations and responses, allowing investors to respond to some structural changes differently, even if they have exactly the same data. As a result, their expectations about future risks and returns might be different. This theory does not categorise some common behaviours as irrational. It explains that because certain market conditions are too fast to analyse, people might fail to understand fundamental changes in the world and their expectations can be wrong, not irrational.

The Rational Beliefs Theory (RBT) presents some important differences with respect to the RET; it allows two economic agents who are equally intelligent and who have exactly the same information to make two different rational forecasts by holding two competing theories which are compatible with the data. In this model the agents can disagree with each other in several ways. For example, they can assign different (i) weights to the possibility that the environment is stationary, (ii) probabilities during the time sequencing of events and (iii) opinions to important and rare events. Disagreement among rational agents arises form their different theories about the nature of the fluctuations of the system rather than from the behaviour of its long term averages [Kurz 94b, Kurz 94c].

The economy defined here is stable in the sense that it is possible to use statistical techniques to compute density functions of relevant variables. But this does not imply that it is stationary. In fact, structural breaks can happen at unpredictable future times. Agents might have different beliefs about the economy, they will not form rational expectations in general and therefore do not have to agree on using the same model of the economy.

Furthermore, Kurz defines Rational Belief Equilibrium (RBE) as a basis for a new theory of asset pricing. In his theory, Rational Beliefs are probability beliefs about future economic variables which cannot be contradicted by the data generated by the economy. RBE is an equilibrium in which the diverse beliefs of all the agents induce an equilibrium stochastic process of prices and quantities and these beliefs are, according to Kurz, wrong and rational in the sense that they are different from the true probability of the equilibrium process. Here agents use the wrong forecasting functions and their forecasting mistakes play a crucial role in the analysis. In [Kurz 94a, Kurz 97] he shows that these mistakes are the reason why stock returns are explainable in retrospect and forecastable whenever the environment remains unchanged over a long enough time interval for agents to learn the forecasting function. He then shows that these mistakes generate endogenous uncertainty: it is that component of the variability of stock prices and returns which is endogenously induced by the beliefs and actions of the agents rather than by the standard exogenous state variables.

There are three important implications of RBE, explained in detail in [Kurz 94a, Kurz 97]:

1. Common stock returns are forecastable within each environment but it takes time for agents to learn and approximate the forecasting functions. For some agents the time is too short so that it is too late to profit from such learning.

2. The equilibrium forecasting functions change from one environment to the other in an unforecastable manner so that learning the parameters of one environment does not improve the ability to forecast in the subsequent environments.

3. More than 2/3 of the variability of stock returns is due to endogenous uncertainty rather than exogenous causes.

This theory seems more robust than RET in the sense that it deals with a world that is subject to structural change and where the investors' expectations will not necessarily agree with each others even when they receive the same type of information. People weigh situations in different ways and come up with different interpretations. Properties such as these are captured in a more realistic way by this theory. In conclusion, as pointed out by [Beltratti *et al.* 96], the main contributions of RBT are that "this theory proposes a vision of the economic system which is much different from the one proposed by rational expectations, and it is compatible with a research project which assigns a central role to learning in the study of behaviour of economic agents."

2.3 A Step Forward: Learning and Adaptation in the Economy

The previous sections give a careful and subtle review of economic models and theories starting from the early 1960s. These are characterised mainly for considering static environments and assumptions that are, in general, useful for certain types of problems but for the scope of this thesis, far from real. From here on, a new economic view which starts to address "learning" as an important factor of the economic model will be given. So far only RBT in section 2.2.2 addresses learning, but it is done so in a very statistical way. This could be due to the issue of learning still being in its infancy in economics, as well as other factors such as how agents learn in practice and how their decision-making process works, including information processing. In this and the next sections, special attention will be given to situations of incomplete information and uncertainty.

In a talk during the Conference "Einstein Meets Magritte" given at the Free University of Brussels, W. Brian Arthur [Arthur 94b] introduces us to the subject brilliantly:

> "Much of what was real and machine-like and objective and determinate at the start of the century, by mid-century was a phantom, unpredictable, subjective and indeterminate. What had defined science at the start of the century–its power to predict, its clear subject/object distinction–no longer defines it at the end. Science after science has lost its innocence. Science after science has grown up.... There are indications everywhere these days in economics that the discipline is losing its rigid sense of determinism, that the long dominance of positivist thinking is weakening, and that economics is opening itself to a less mechanistic, more organic approach. In this talk I want to show my own version of this loss of certainty. I want to argue that there are major pockets of uncertainty in the economy. I want to show that the clear subject/object distinction in the economics often blurs. I want to show that the economy is not a gigantic machine, but a construct of its agents. These are not 'anomalies' to be feared, they are natural properties of the economy, and if we accept them, we will have a stronger, not a weaker science."

Arthur's views clearly opened new possibilities to explore other than the traditional, perfectly rational "economic man" that reasons deductively on problems that are well defined. The standard view no longer explains the true character of new economic problems; factors such as technology cannot be ignored or treated as exogenous anymore. With this new approach, the structure of the economy emerges from the aggregate effects of the agents' subjective beliefs, which in turn, arise, decay, coevolve, change, mutually reinforce and mutually negate. This what according to Arthur's views gives rise to the character of financial markets.

2.3.1 The Deductive Metaphor

Reasoning in economics usually starts with the assumption that agents are perfectly rational, in other words, they obey axioms of reasonable and logical behaviour.

Arthur exemplifies the deductive way of reasoning with a typical "micro-economic" problem taken from the industrial organisation literature of the mid 80s. The problem concerns a number of airline companies, which have to choose time slots for their planes without being too close to the others' time slots.

According to the RET with perfect, deductive rationality, airline X would choose and figure out its time slot knowing the arbitrarily positions chosen by the previous X-1 airlines **and** assuming that airline X+1 will also choose the optimal slot given the positions of the previous X airlines. When going from the last airline to the first one (from X = N ... 1), this can get complicated, but the logic can actually work out in reverse order by backward deduction. As every conclusion follows from the given premises -is deduced,- a decision process of this kind is called *the deductive metaphor*.

Now, let's analyse what sort of assumptions are involved in this solution; every airline **must know**:

1. exactly what its preferences are
2. the others' preferences
3. that every other airline accurately knows the preferences of every other airline
4. that every airline knows that every airline knows the preferences of every other airline and so on in an infinite regress
5. that it is rational enough to work out the solution

6. that every other airline is rational and will use perfect rationality to work out the solution

7. in an infinite regress again, that every other airline is using this rational way to work out the problem, because if one of these airlines messes up, it messes the solution up for every other airline.

In addition, the optimal placement of each airline using this backward deduction must be unique. If any link of this network of requirements breaks, the solution ceases to exist [Arthur 94b].

In this scenario people are clearly trying to predict a world that is created by their beliefs and everybody else's as well. A self-referential loop is found as each prediction depends on the expectations others might form. Or, in other words, predictions are forming a world those predictions are trying to forecast. There is no logical way to determine each prediction because there is a logical indeterminacy – without knowing how others determined their forecasts, each forecast is therefore indeterminate.

In summary (addressed in [Arthur 92]), a solution involving the deductive metaphor requires the following ingredients:

- full knowledge of the problem
- a unique solution
- perfect ability to compute a solution
- common knowledge that other agents operate under the conditions of fully knowledge of the problem and that there is a unique solution.

The deductive metaphor works well in cases in which the solution can be computed so everyone can figure out what to do, but clearly financial markets are definitely not the case. In financial markets these expectations become unstable: if some people think prices will rise because they think others think the same, they start buying the stock and the price indeed goes up until factors such as negative rumours make people revise these expectations again. Then prices might suddenly change direction due to the fact that others might be expecting a drop at this time. Because some reassess downward and others upward trends, predictions then become unstable and price bubbles start. Under the realistic assumption that traders may interpret the same information differently, expectations become indeterminate, unstable and possibly mutually self-fulfilling.

The economic theory of capital markets presents the same indeterminacy as the airline problem. Agents need to form expectations of an outcome that is a function of these expectations. With reasonable heterogeneity of interpretation of "information," there is no deductive closure. The formation of expectations is therefore indeterminate.

2.3.2 The Inductive Metaphor

The deductive logic works well for certain problems in economics where the computations needed to arrive to a solution are known, computable and available to all. But what about other problems where these deductive metaphor assumptions do not apply and yet their solution seems to come from a deductive logic process?

To answer this question, Arthur defines a *problem complexity boundary* as the point "beyond which arriving at the deductive solution and calculating it are unlikely or impossible for human agents; and beyond which other agents cannot be relied upon to carry out their part of the deducing process. Beyond this boundary where the requirements fail, there are often no instructions on how to proceed logically; so that away from the prescribed solution –'out of equilibrium'– the problem becomes ill-defined and the deductive metaphor cannot operate at all. Exactly where this boundary lies is of course fuzzy." ([Arthur 92], pages 5-6.)

Agents might use some behavioural process other than the deductive one when dealing with certain problems. Could we say that people use a process that brings them to a deductive type of solution without actually following any deductive logic step by step?

> "And yet... and yet... in every market, in every day, people do form expectations. How do they do this? If they cannot do this deductively, then should we model their behaviour in this area?" ([Arthur 94b], page 5.)

But how can we model people's behaviour? A good candidate might be *inductively*.

To address this issue firstly we need to understand more about the way people reason and behave when making decisions in complex problems where the goal is to behave optimally in a world where "behaving optimally" is also ill-defined.

One important factor that starts to be mentioned in this quest is *behaviour*. With this in mind, Arthur analyses the sophisticated human decision-making process using some fundamental concepts about cognition. These concepts do not appear to be fully understood from any of the sciences yet, but they play an important role in the process. The following are some questions that will be addressed in the following paragraphs under an inductive framework: how to use what we know? What is knowing and learning? If learning is viewed as an updating process of some kind, what then is knowledge?

In the inductive view, we tend to think in fixed categories, knowledge is the data or stored facts, all organised into these fixed categories – the information set. To illustrate this, let's think how a person would deal with the following situation: the price of a stock is £45.38. As this datum is a fact, it is stored in some organised way under that stock's name, a fixed category. It then becomes part of a person's knowledge base, the information set. The next day, when the price is £48.60, this information is revised and updated accordingly. In this context knowing means retrieving something stored in categories. Learning is revising it and updating it according to the new external information that has just been perceived. As a result, such activities change the person's behaviour. For example, the updating could cause a decision to buy, sell or hold certain amount of that stock.

This process involves building and rebuilding representations of the world by constructing some and discarding other old and obsolete ones. New sensory perceptions can easily be dropped into new well defined categories, but acquiring knowledge is not just acquiring facts, it often has to do with searching over our representations in order to find categories that best suit certain conditions we have encountered in the past. In this sense, the economic world under this type of learning is path dependent: two people facing the same situation can have different experiences associated with it and act over the same data differently.

Humans are very good at looking ahead, anticipating consequences, seeing, recognising and matching patterns. As addressed by John Holland, we appear to form internal models or mental models of the situations we are dealing with and rerun these scenarios in our minds [Holland *et al.* 86]. We are able to search very easily, use analogies, transfer experience from one problem to another and are able to generalise: reason from the particular to the universal. We use these methods to help us to infer from partial information of the problem to a whole solution. However, we do so using finite resources of memory, time and processing; Simon [Simon 55, Simon 59] refers to this as bounded rationality.

Beyond the complexity boundary where perfect deductive behaviour fails, we operate heavily in the **inductive mode**. Induction can always be improved by learning, and in this sense, induction calls for a dynamic approach. This does not mean that there is no deductive process present at all, actually both are present at times, but because the deductive is much more limited in the complexity it can cope with, it becomes overloaded and gives way to the inductive mode.

These are some of the issues that contributed to the birth and rise to Evolutionary Economics. Modelling behaviours from a dynamical perspective forms the main topic of this thesis.

2.4 Who knows what really happens? – Views from the experts

Let's review the situation first: under EMH, today's price discounts in – or reflects already – all new information that could be used to forecast the price of the stock tomorrow. From here it follows that it would be impossible to predict tomorrow's prices at all, and technical traders with all their charts and calculations are of no use, nor fundamentalists with all their ratios. As a result, the average investor can not beat the market and a prudent strategy to follow would be to buy some index funds or good stocks and hold them for a long period of time. This view clearly favours what is traditionally known as passive investment. However, many investors disagree arguing that "new information is absorbed differently by different investors at different rates; thus, past price movements are a reflection of information that has not yet been universally recognised but *will* affect future prices" [Ross 99].

Table 2.1 shows the names and current positions of great contributors of various practical and theoretical investment methods. The purpose of this section is to share their personal views about the relationship between economic theory and practice in the financial field, given with no specific order.

Table 2.1: Some Experts in Finance

Andrew W. Lo	Harris & Harris Group Professor of Finance at the Sloan School of Management, Massachusetts Institute of Technology
A. Craig MacKinlay	Joseph P. Wargrove Professor of Finance at the Wharton School, University of Pennsylvania
Burton G. Malkiel	Chemical Bank Chairman's Professor at Princeton University, Economics Department
Eugene F. Fama	Robert R. McCormick Distinguished Service Professor of Finance, University of Chicago
Richard Olsen	Chairman and CEO of the Olsen Group
Peter Lynch	Former Manager of the Fidelity Magellan Fund
David E. Shaw	Chairman and CEO of D. E. Shaw & Co.
George Soros	President and Chairman of Soros Fund Management LLC, and Chief Investment Advisor to Quantum Fund N.V

The first to start with is Eugene Fama, whose ideas are clearly summarised in an article from an interview which originally appeared in an Ibbotson newsletter [Ibbotson 01], by answering the question **how do you invest?** as follows:

> "I'm a passive investor. All stocks. I believe in efficient markets. I know nothing about stock picking. I don't trust anyone else to interpret the data better than myself so I don't believe the opportunities are there to beat the market."

Fama is not the only passive investor in the market; other investors share these beliefs as well, such as Rex Sinquefield, who points out his views in another interesting interview [Tanous 01].

The second one to mention is George Soros, whose fund is recognised as the best performer in the world during its entire history and who is still remembered by many as the man who broke the Bank of England back in 1992 with a single currency transaction which was worth more than $1 billion:

> "The prevailing wisdom is that markets are always right. I take the opposite position. I assume that markets are always wrong. Even if my assumption is occasionally wrong, I use it as a working hypothesis" [Soros 01].

He has extended his beliefs about economic theory as follows:

> "The theory of rational expectations, the theory of efficient markets, the random walk theory, are all highly respectable scientific theories but they are based on a somewhat outmoded, deterministic view of the universe. To be specific, classical economics has modeled itself after Newtonian physics and it is trying to determine the equilibrium position. In order to do so, it must postulate assumptions which may or may not prevail in the real world. To be specific, it must assume away the phenomenon of reflexivity. As a result, its conclusions may not fit the real world" [Soros 95b].

> "Economic theory has managed to disregard the gap [between perception and reality] by taking demand and supply as given and focusing its attention on the relationship between supply and demand. It has construed an elaborate interpretation of reality which is, at least in one case, namely in the behavior of financial markets, far removed from reality. Economic theory, for instance, is valid as a hypothetical construct in which some of the consequences of imperfect understanding are assumed away. A distortion is created only when we apply the conclusion of economic theory to the real world. This is particularly noticeable in financial markets. The theory of rational expectations and efficient markets yields highly misleading results" [Soros 95a].

> "Economic theory needs to be fundamentally reconsidered. There is an element of uncertainty in economic processes that has been largely unaccounted for" [Soros 01].

Burton G. Malkiel in his book entitled "A Random Walk Down Wall Street" [Malkiel 99] agrees with a more relaxed market efficiency, under the assumption that the transaction costs reduce any of the advantages a given strategy could offer, so that a buy-and-hold strategy of index funds produces higher returns. He shows that a broad portfolio of stocks selected by chance performs as well as one that has been carefully chosen by the experts. In this book he also compares the *holes* of the EMH to proverbial \$10 bills lying in the gutter. The view of some economists is that a given person cannot find \$10 bills in gutters because someone else have already picked them up. This again seems like a contradiction: "someone else already picked them up!". Being *the* one is the hope of millions of investors out there, investors which in turn are affecting the market dynamics by their decisions and expectations. But this argument, although interesting, is not the purpose of this thesis.

As a response to these claims, Andrew W. Lo and A. Craig MacKinlay have provided important evidence showing that financial markets are not completely random. They have edited a number of their papers in the book entitled "A Non-Random Walk Down Wall Street" [Lo & MacKinlay 99]. In this volume, they basically put RWH to the test, finding that predictable components do exist in recent stock and bond returns. In addition, they also explain various techniques for detecting predictabilities and evaluating their statistical and economic significance and offer some of their views of the financial technologies of the future.

With more or less the same line of thought, Richard Olsen founded a research firm which specializes in prediction with high-frequency data. The organization has made enormous contributions supporting the growing body of evidence that financial markets deviate from EMH and RWH: their claim is that markets show some internal structures due to the ways participants act based on their differing profiles such as risk aversion, time horizons, etc. They offer a number of financial time series forecasting products. D. E. Shaw is another major player in financial markets with a very modern quantitative approach that makes use of both, technology and finance to find better investment opportunities, portfolio and risk management, and the reduction of transaction costs.

2.5 Concluding Remarks

There exists mixed evidence regarding market efficiency and trading profitability. With all these different views, could we ever take any one of these positions as dogmatic? I believe not.

One could argue that markets are becoming more efficient in their handling of information, which is becoming easily and rapidly available, but the rate in which it is compounded into the prices is still the subject of many debates. Even more important, the way people react to events can not be factored out immediately, their expectations are based on what they expect that others expect and therefore many differing actions affect the price dynamics in different ways. Some of these price changes could be predicted with some success as Olsen Group, D. E. Shaw & Co., Prediction Company, Parallax, and many other research boutiques believe. In this sense this is where this thesis stands.

Inconsistent with the EMT, one could also argue that not everybody would pick up the $10 bill even if they are the lucky ones to find it. Along these lines, there are a number of reasons why EMH may not be correct, because investors have a wide variety of reasons for trading, e.g. for a short-term profit or a steady profit with long-term stability, or some even to lose money. Transactions are also performed by reasons other than to make a profit; for example, the Government of France, buying francs to support their price rather than to make a profit. Another very interesting evidence suggesting that the stock market is not a random walk is that it tends to trend in one direction for too long at each time. Some people believe that many investors wait until they see a price trending upwards to confirm their instinct to buy that asset. They, therefore, tend to reinforce trends, producing a self-fulfilling prophecy by second-guessing each other's predilections [Ridley 93]. So people's expectations can generate price movements or change the dynamic of the market. In general, when it is predicted that the price of a stock will increase, it produces a buy, which will then produce the price of the stock to rise. Such prophecies deliver some elements of predictability that are captured by some systems in practice. This is why it has been demonstrated that price behaviour of each day depends, indeed, up to a certain point, on the price of past days, refuting the idea that prices follow a random walk.

By looking at history one can see that it has been possible to make money by *anticipating*. Whether it is economic shifts around the world or the possible outcome of a stock's price in a short period of time, trying to *anticipate events right* is a crucial factor in financial markets. Peter Lynch said in a recent interview "In this business if you're good, you're right six times out of ten. You're never going to be right nine times out of ten. This is not like pure science where you go, 'Aha' and you've got the answer. By the time you've got 'Aha,' Chrysler's already quadrupled or Boeing's quadrupled. You have to take a little bit of risk" [PBS01].

This Chapter has covered the basics of economics from its most rigid and classic theories to some of its most modern approaches. When considering financial markets, it is striking to see the dichotomy in the lines of taught in the whole industry, ranging from major financial institutions, large public and private investors and academics to even the smallest practitioners who buy and sell a few shares over the Internet. If at least **all** academics agreed in the EMH view and **all** practitioners were against the theory (trying to beat the market) things would be easier to understand before taking either one of those positions. But even with academia itself, a great number of academics believe that prices do not follow a random walk (see [Lo & MacKinlay 99]) and many others still believe that prices do follow a random walk ([Malkiel 99]). It appears that both worlds have enough evidence to last them for a life-time supporting their views. The same dichotomy occurs with market investors. The bottom line is that the market as a whole is composed of both types of investors and both are changing the market dynamics according to their behaviours and that is exactly the position taken in this thesis. One that accepts the possibility of both views, even if it appears to be a contradiction.

On one hand, it accepts that it is likely that the proper use of certain types of information could help in the discovery of a trading rule performing better than the average: up to this point this view looks like being against the EMH. But the deviation starts here: it is hypothesis of this thesis that such profitable rule could NOT be profitable for a long period of time, therefore the motivation is to develop sets of strategies continuously in order to check the consistency of results of a forecasting model over a long period of time. This thesis will address this issue by providing an important insight when looking at real financial markets from an evolutionary framework in which the main actors are heterogeneous agents.

At the same time, it is not the purpose of this work to be totally in favour or against traditional economic theories such as EMH or RET, but rather to give an alternative evacuation route from that long debate in which perhaps even addressing the question should not be a relevant issue anymore because the main concern here is to investigate whether it would be possible to design and if so, how could a completely unexperienced artificial agent gradually learn and adapt its behaviours in a real market environment, and most important of all, in a human-like fashion? Secondly, would it be possible to use such market model to tackle the ancient problem of forecasting? This is a problem millions of people have tried from various perspectives and it is also the concern of this thesis to try to find out whether it would be feasible to learn to forecast real markets by modelling groups of market participants in a real environment.

Chapter 3

Artificial Intelligence Background

"Over immense periods of time the intellect produced nothing but errors. A few of these proved to be useful and helped to preserve the species: those who hit upon or inherited these had better luck in their struggle for themselves and their progeny. Such erroneous articles of faith... include the following: that there are things, substances, bodies; that a thing is what it appears to be; that our will is free; that what is good for me is also good in itself." Friedrich Nietzsche, *The Gay Science*, translated by Walter Kaufmann.

3.1 Introduction -Artificial Intelligence?

Besides the well-known classical perspective of Artificial Intelligence (AI), developed in the 1950s, based mainly on symbol manipulation and first-order logic with some accomplishments in the areas of expert systems, game playing and natural language processing, there are other more biologically oriented approaches that have recently emerged and rapidly grown, some of which include various ways of brain modelling, evolutionary computation and even immune systems design. These methodologies try to simulate some of the mechanisms that are believed to be involved in human intelligence and tackle common sense, taught processes and concept formation in more feasible ways. They follow a trend which accepts imprecision and uncertainty, a common characteristic of real world problems. Extensive research has been made in fields characterised by the use of numerical methods rather than symbolic ones and good solutions have been reported to complex real-world problems involving process control, signal processing and optimisation in areas such as neurocomputing and fuzzy theory.

Under the view point of a new approach called Soft Computing (SC), complex real-world problems require Intelligent Systems (IS) that combine knowledge, techniques and methodologies from various sources [Jang *et al.* 97]. The idea is that such intelligent systems will possess human-like expertise in a given domain and will also be able to adapt and learn to improve in the changing environment they are part of. Soft Computing favours the use of hybrid intelligent systems to cope with complexity.

It is beyond the scope of this thesis to examine philosophical questions of whether such more modern methods lie within the boundaries of AI, but it will be assumed here that the main goal of AI resides in a computing approach to machine intelligence, i.e. the creation and better understanding of machine intelligence and therefore it is considered throughout this thesis that these new methodologies do belong to a more modern AI approach, one that focuses in learning from a dynamical-systems perspective.

3.1.1 Intelligent Systems

Intelligent Systems adaptively estimate continuous functions from data without specifying mathematically how outputs depend on inputs [Kosko 92]. Humans are particularly good at this type of task. We tend to associate responses to stimuli in various ways, by successfully mapping (i) actions with different scenarios, (ii) effects with causes or (iii) labels with patterns. This can viewed mathematically as a function of the form $f : X -> Y$ which transforms inputs to outputs. Systems can be defined according to these transformations. An (IS) is generally speaking one that formulates "appropriate, problem-solving responses when faced with problem stimuli" [Kosko 92]. In addition to this, such systems have the following properties:

1. Generalisation. Appropriate responses are obtained when faced with unseen circumstances due to their behavioral repertoires being far greater than their actual experiences.

2. Learning and Adaptation. Behavioural changes occur as a result of attempting

to improve their responses – some people regard it as simple parameter adjustments.

3.2 Fuzzy and Neural Systems

After a brief introduction to these systems is given in this section, similarities and differences will be given in the following section (3.2.3).

3.2.1 Fuzzy Theory

> "Fuzzy Theory holds that all things are matters of degree.... [It] also reduces black-white logic and mathematics to special limiting cases of gray relationships. Along the way it violates black-white 'laws of logic,' in particular the law of noncontradiction *not-(A and not-A)* and the law of excluded middle *either A or not-A*, and yet resolves the paradoxes or antinomies [Kline 80] that these laws generate. Does the speaker tell the truth when he says he lies? Is set *A* member of itself if *A* equals the set of all sets that are not members of themselves? Fuzziness also provides a fresh, and deterministic, interpretation of probability and randomness." [Kosko 92]

Fuzzy sets are not defined as having crisp boundaries. The gradual transition of being part of one set to being part of another one is expressed as a membership function that allows the modelling of linguistic expressions commonly used such as defining the age of a person as *young* or *old*. Classifying a person under one of these two attributes will most likely yield to a different outcome depending, for example, on who is doing the task, a child or an older individual. Modelling fuzzy systems allows the use of ambiguous concepts by first converting crisp variables (such as height = 1.73 m. and weight = 89 kgs.) into linguistic variables (tall, fat), calculating the degree of membership in each fuzzy set for all crisp inputs. The following step is to perform some sort of reasoning mechanism (usually known as inference procedure) upon the rules and given facts to derive reasonable conclusions. A process of defuzzification will follow in order to extract crisp values that best represent the fuzzy sets. There are a number of inferencing and defuzzification techniques, a good review of them and an introduction to the whole process can be found in [Cox 94].

3.2.2 Artificial Neural Networks

The growing discipline of Artificial Neural Networks (ANN), also known as Artificial Network Systems (ANS), neurocomputing, network computation and connectionism among others, deals with a type of model that must be taught or trained in order to learn different tasks such as associations, patterns, dependencies, etc., in a way reminiscent of how biological neural networks operate. The biological inspiration for the development of these models is the human brain, which contains millions of interconnected neurons, operating in complex and parallel ways. But ANNs are far less complex layered structures of simple computational elements called nodes, which are connected by links with adjustable weights. The process of "learning" is simply represented by changing parameters, i.e. adjusting the connection weights while passing a finite number of observations from the input to the output nodes of the network. The idea is that after this set of limited examples has been seen by the neural network for a considerable number of times, the net should then be able to infer strategies of behaviour and as a result, it should be able to deal with similar situations even when it has not seen them before. A widely used learning (or training) method is the back-propagation algorithm. For further details of this and other training methods refer to [Zurada 92].

> "A *neural network* is a parallel, distributed information processing structure consisting of *processing elements* (which can possess a local memory and can carry out localized information processing operations) interconnected via unidirectional signal channels called *connections*. Each processing element has a single output connection that branches ('fans out') into as many collateral connections as desired; each carries the same signal – the *processing element output signal*. The processing element output signal can be of any mathematical type desired. The information processing that goes on within each processing element can be defined arbitrarily with the restriction that it must be completely local; that is, it must depend only on the current values of the input signals arriving at the processing element via impinging connections and on values stored in the processing element's local memory." [Hecht-Nielsen 90]

Unlike traditional computer programs that perform specific tasks through sequential instructions, neural networks do not explicitly require an algorithm to be executed by each computing node. Instead, when designing a neural network, one must choose and specify other factors that will play an important role when measuring performance of the model, such as: (i) the training algorithm to be used, (ii) the type of architecture, (iii) the initial values of parameters such as weights, learning rate, etc., (iv) the type of neuron computations to be performed, etc. Then the idea is that the network will simultaneously explore many competing hypotheses and through a training process will map significant states of the analysed system into more compact internal representations.

Neural Networks have successfully been used in a wide variety of classes of problems such as pattern recognition, classification and control, with applications ranging from manufacturing, quality control and optimisation to fraud detection. The interest in such applications has been mainly motivated by their excellent mapping abilities and good generalisation. In addition, a number of other advantages have also been reported: "The power of an ANS approach lies not necessarily in the elegance of the particular solution, but rather in the generality of the network to find its own solution to particular problems, given only examples of the desired behavior .... the inherent ability to deal with noisy or obscured patterns is a significant advantage of an ANS approach over a traditional algorithmic solution" [Freeman & Skapura 91]. However, it is equally important to recall at this point that ANNs usually require not only different architectures and training approaches but also different data stored in the network in the form of weights. These drawbacks will be discussed in the following section.

3.2.3 Similarities and Differences

A property that neural and fuzzy methodologies have in common with intelligent systems is that they offer a model-free function estimation[1]. Besides the similarities of fuzzy sets and neural networks when compared to other more traditional approaches, due to their nature, they represent structured knowledge in different ways:

- The fuzzy system directly encodes the linguistic sample in a finite numerical matrix usually by describing inputs in a fuzzy distribution. But as fuzzy systems map fuzzy sets to fuzzy sets, the required output must be defuzzified to obtain

[1]Here, the function is simply taken as the one described in section 3.1.1 and model-free means that there is no need to explicitly state how outputs depend mathematically on inputs.

numerical values as output. This can be done in several ways, described in detail in [Jang *et al.* 97], along with various fuzzy inference system types.

- A Neural Network needs a great number of input-output samples to be trained and might need to recycle them thousands of times during the learning process. In addition to the problem of training, it has been difficult to know what the neural network is actually encoding, what it is recalling and forgetting and therefore we might not even know if it is encoding the original structure at all. We can characterise the neural network's behaviour only by exhaustively passing all inputs by through the black box and recording the recalled outputs [Kosko 92]. This characterisation involves a computational dilemma: as in most problems, where the number of input-output pairs we need to check is intractable, the neural network will be unreliable, and in the few cases of having a finite set of pairs it would be unnecessary to use a neural network because we could instead store these pairs and retrieve them directly and without error as a look-up table.

In this sense fuzzy systems provide a combination of purely numerical approaches of neural networks and mathematical modelling with the symbolic, structure-rich approaches of artificial intelligence.

With all their advantages and disadvantages, neural methods and fuzzy systems offer an alternative approach to decision making over traditional problem-solving techniques. However, there are other modern methods which rely on interesting ideas of evolution that have not only been promising when tackling complex real world problems, but have also provided with important insights of how adaptive processes occur in natural systems. These will be described in the following sections. First, it will be useful to describe genetic algorithms.

3.3 Genetic Algorithms

Genetic Algorithms (GAs) were developed by John Holland, his colleagues and students at the University of Michigan during the 1960s. Their main goals were to study the phenomenon of evolution in natural systems and to incorporate the basics of such biological adaptation ideas into the development of artificial systems. But it was not until 1975 when John Holland introduced the Genetic Algorithm more formally, in his book entitled *Adaptation in Natural and Artificial Systems* [Holland 75], where he provided a more theoretical and biological background of the genetic algorithm. For an introduction and biological terminology widely used in GAs consult [Mitchell 96].

3.3.1 Canonical GA

A GA can be simply described as a series of steps designed to convert a *population of chromosomes* (e.g. binary strings) into a new one – and hopefully better – by means of natural *selection* mechanisms and genetic operators of *crossover* and *mutation*. There are a number of different variants of this process, here only the canonical GA [Whitley 93] will be explained. The implementation is also known as Simple Genetic Algorithm (SGA) [Goldberg 89]). For more on GA implementations refer to [Michalewicz 96].

The basic process can be visualised in diagram 3.1, where the initial block consists of a single step where a random initial population of binary strings of length L is generated, evaluating every member of the population (measuring its performance with respect to a particular set of parameters) in order to assign them a fitness value. Here, a fitness function transforms that measure of performance into an allocation of reproductive opportunities. The evaluation of a given string is independent of the evaluation of other strings, but the fitness of one string is taken with respect to the others, defined as: $fitness = \frac{f_i}{\overline{f}}$, where f_i is the evaluation associated with string i and $\overline{f}$ is the average evaluation of all the strings in the population. This population immediately becomes the *current population*.

The second block shows the execution of one *generation* of the canonical GA, but usually a large number of generations is needed before reaching an appropriate or optimal solution to a given problem. As shown in the diagram with the arrows, these steps are iterated until a particular stopping criteria has been achieved, such as reaching the limit of a fixed number of generations or when an optimal solution has been achieved. This block consists of three main steps: the first one is to evaluate and assign a fitness value to every chromosome in the *current population*, as explained in the previous block. The second step is to apply a selection mechanism to the current population in order to create an *intermediate population* that should contain more of the better individuals (chromosomes) that were in the current population. And finally, the third step is performed by applying to this population the operators of (i) recombination, in order to obtain some mixed properties of the good elements and (ii) mutation, to add some new features to this mix. The new chromosomes will now be part of the *next population*.

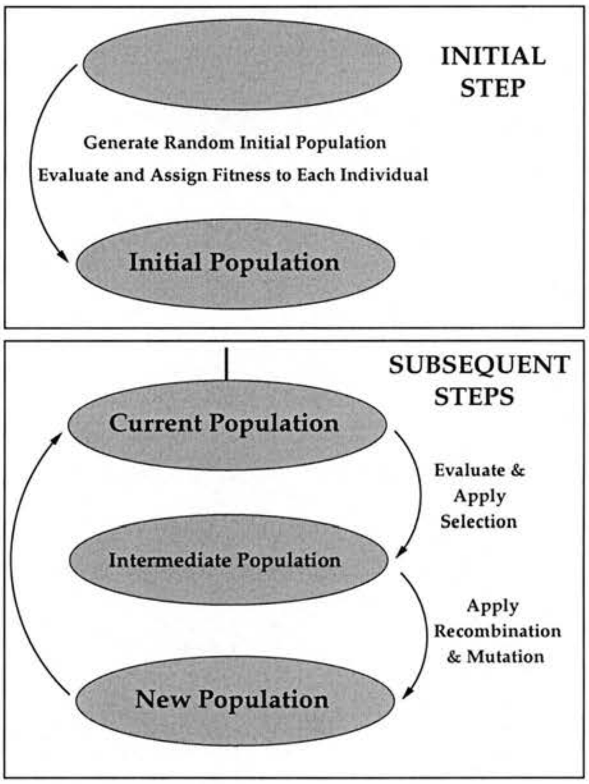

Figure 3.1: Diagram of the Canonical GA

3.3.2 Representation

The idea is that a chromosome represents one point of the search space of candidate solutions, and the process of converting one population into another one is, as explained earlier, achieved by assigning a score or *fitness value* to each chromosome as a way of measuring how well it solves the problem in question. The fitness landscape is a plot that represents the space of all possible chromosomes along with their fitnesses, and the GA can be viewed as the process that moves the population along this landscape which might be composed of hills, peaks, valleys, etc. depending on the complexity of the problem. Evolution and adaptation will tend to move the individuals of a population in specific ways, with the main goal being to find and approach the regions of higher fitness – the peaks, which can be local or global.

In the standard process just described, the individuals in the population are usually binary strings, i.e. each one of its loci (string positions) has two possible alleles (feature values): 0 or 1, but the representation does not always have to be binary. For certain problems, other representations such as many-character or real-valued encodings might be more practical.

3.3.3 Selection

After a careful evaluation of the goodness of each member of the current population, selection of the parents for reproduction is the first of the iterated steps that represents an artificial version of how natural selection operates: Darwin's theory of survival of the fittest is implemented in one of its simplest forms as a roulette wheel (fitness-proportionate selection), where the size of every slot is proportional to every individual's fitness divided by the population's average fitness. The selection process starts by spinning the roulette wheel N number of times, once per individual in the population. Note that because there is a random element in this process, this method does not necessarily guarantee that the selection will be proportional to every element's fitness, specially for small populations. Ideally, the more times the roulette is spun, the more the selected members would represent their true fitness proportion in order to overcome the luck of the draw effect. Better methods have been devised, such as Stochastic Universal Sampling (SUS) [Baker 87], where due to its structure being made of N equally spaced arms, all N members are chosen by spinning the wheel only once, closely guaranteeing that the individuals will be selected for reproduction a number of times equivalent to their fitness proportion.

Other methods include steady-state, elitism, tournament and rank selection.

Steady-state selection works as follows: in every generation only a small number of individuals (the least fit ones) is replaced by the new members that were created after crossover and mutation. "Steady-state GAs are often used in evolving rule-based systems (e.g., classifier systems; see [Holland 86] in which incremental learning (and remembering what has already been learned) is important and in which members of the population collectively (rather than individually) solve the problem at hand" [Mitchell 96].

For an extended description of selection mechanisms and genetic operators refer to [Holland 75, Mitchell 96, Goldberg 89, Michalewicz 96, Whitley 93], this thesis only covers topics relevant to the development of the proposed model.

3.3.4 Crossover and Mutation

Following selection, the intermediate population is created and recombination occurs by applying crossover to the selected pairs of chromosomes according to a reasonably high probability denoted by P_c (usually set to 0.7, often, not a very sensitive choice). The simplest case of crossover is *one-point crossover*, where one location of the chromosome is randomly selected, dividing both parents into two parts which are exchanged to form the offspring. When two children are created, all the genetic material from the parents will be preserved, but when one child is created, only half is preserved. Note that with probability $1 - P_c$, the selected parents will be transferred exactly as they are to the following step carried out, which is mutation. Other variations of crossover include two-point crossover and uniform crossover. Again, these will not be covered in this thesis.

The value of each bit of the population is flipped according to a very low probability denoted by P_m (usually set to 0.01 or less, or 1/N, where N is the number of genes in a chromosome) after crossover has been performed. The hope is that mutation could help "prevent the irrecoverable loss of potentially important genetic material" [Goldberg 89]. The considerable difference in these probabilities follows from the idea that "in biological systems crossover is much more frequent than mutation, often as much as a million times more frequent" [Holland 95].

3.4 Learning Classifier Systems

We now come to the alternative method of decision making modelling, an adaptive one, that was alluded to in section 3.2.3.

As defined by [Goldberg 89], a "classifier system is a machine learning system that learns syntactically simple string rules (called classifiers) to guide its performance in an arbitrary environment", where machine learning primarily refers to systems which acquire and improve knowledge by using input information.

In Learning Classifier Systems (LCS), the general idea is to learn concepts through decision rules that account for positive examples in order to "predict a classification of previously unseen examples, or suggest (possibly more than one) classifications of partially specified descriptors" [Michalewicz 96]. The learning takes place by adjusting certain values associated to the rules according to the environmental feedback they receive and by discovering new and better rules. Considering that chromosomes are good ways of representing knowledge as well as good candidates to be manipulated by genetic operators, the GA community quickly responded with two distinctive approaches that were labelled according to the Universities where they were developed: Pittsburgh (or 'Pitt approach,' led by De Jong and his students) and Michigan (mainly led by Holland, Reitman and Booker).

In the Michigan-style approach (also called population of rules approach), it is pointed out in [DeJong 88] that "the members of the population are individual rules and a rule set is represented by the entire population (e.g., see [Holland & Reitman 78, Booker 82])." In this view, individuals compete via fitness for reproductive rights. The competition of rules takes place within the set of rules, and highly fit individuals have the opportunity to match with other highly fit individuals thereby increasing the chance that their progeny will have better survival characteristics.

In the Pittsburgh-style classifier approach (also known as population of rule sets approach), the competition occurs between entire rule-sets. The idea here is "to represent an entire rule set as a string (an individual), maintain a population of candidate rule sets, and use selection and genetic operators to produce new generations of rule sets" [DeJong 88]. So the population is composed of multiple rule-sets competing to mate other rule-sets and exchanging individual rules with the hope to combine rules over many generations to form an effective rule-set.

The main difference (apart from their State of origin) in these two approaches resides in the nature of the members of the population created (single rules or rule sets). More recently, other two approaches have emerged: the Strength Based Classifier System and the Accuracy Based Classifier System. In the former one, the measure of fitness used by the GA is a value associated to each rule representing its goodness, called *strength*, while in the latter, the measure used is the rule's prediction of reward that the system will receive when the rule is used, called *accuracy*. These two approaches are generally of the Michigan type LCS. In the following sections basic concepts of these approaches will be reviewed, other general and more advanced topics can be found in [Lanzi *et al.* 00].

3.4.1 Strength Based Classifier Systems

Strength Based Classifier Systems represent the classic Holland's approach [Holland 75], characterised by using the rule's strength as the main guidance during the learning process. As noted earlier, strength can be defined as a measure of the payoff received by each rule when it matches certain characteristics of the external environment and when its action is chosen. The implementation of this type of classifier system is also known as Simple Classifier System (SCS) which is fully explained in [Goldberg 89]. SCS is the main LCS used throughout this thesis, and in addition to the concepts reviewed in this section, many others will be explained in more detail in Chapter 7.

In general, SCS uses a population of a finite number of individual rules and reinforcement algorithm (usually the Bucket Brigade algorithm, which will be explained in section 3.4.1.1) in conjunction with a GA. Figure 3.2 shows the main components of the Canonical LCS. As it can be seen from the diagram, information flows from the environment through the *detectors*, where it is decoded into fixed-length *messages*. These messages are compared (checking if there are matches) to all the *classifiers* contained in the list, where some might become active. Because a number of rules of this list can match a particular environmental state, such rules must enter a competition in order to be selected to fire their actions. This is done by making bids proportional to their strengths (net worth); perhaps plus some random noise. Using only strength in bidding would tend to preserve the status quo. The winner of the auction posts its action through the *effectors* and pays its bid to previous winners and based on its action taken, a reward might be received from the environment increasing the strength of the rule. The environmental feedback received by the system is used to assign credit to those rules in the population which led to the payoff. Then using the classifier's measure of strength, rules are reproduced, crossed and mutated as dictated by the GA.

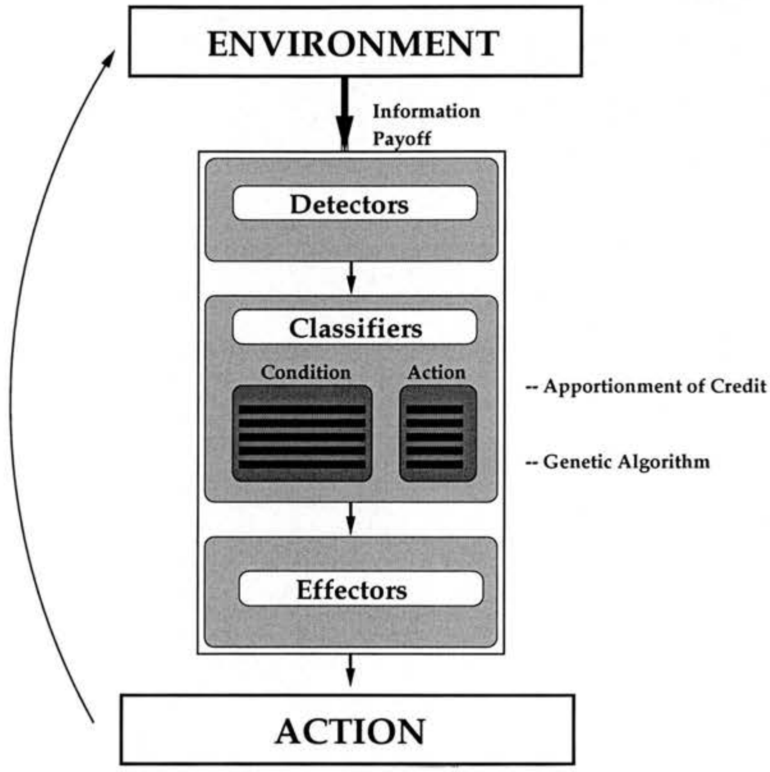

Figure 3.2: Diagram of a Simple Classifier System

So far, in this architecture, three main components have been mentioned: a system of rules, an apportionment of credit system and the GA. But one often wonders how exactly these rules look like and what they represent. In this context, the system of rules is represented as a population of rules of the form *if* $< condition >$ *then* $< action >$, where if the rule is "fired", the action attached to it is taken. It is important to note that in classifier systems rules are restricted to have a fixed-length representation, but this restriction also offers two advantages, stated by [Goldberg 89] as: "First, all strings under a permissible alphabet are syntactically meaningful. Second, a fixed string representation permits string operators of the genetic kind." In addition, "a single rule or a small set of rules can represent a complex set of thoughts compactly."

On the implementation side, the condition is a bit string of fixed-length with each position taken any of the three values 0, 1 or #. The 1 and 0 match corresponding bits in the state vector, while # is a wild-card symbol which matches either of the conditions, allowing some state variables to be ignored. The action part of the classifier has a more specific character and can only take the values of the binary alphabet $\{1, 0\}$.

3.4.1.1 Apportionment of Credit Component

In SCS, the *Apportionment of Credit* or *Credit Assignment* is usually carried out by the *Bucket Brigade Algorithm*, which is encharged of properly allocating funds in what can be viewed as an *economy of classifiers*. The matching classifiers make bids proportional to their strengths to post their actions in the environment in an auction process. In the most simplistic form of bucket brigade, the winner pays its bid to the previous winner. Also, as this economy is no different from any other, there are unavoidable tax payments to all those classifiers that are currently bidding in the auction. This tax is calculated according to the classifier's strength, and it is called *bidtax*. Yet there is another more strict scenario where an extra existence tax could be charged to all classifiers, called *lifetax*. The result of this system is that good rules will have higher strength than those performing badly. Two reasons can explain why such taxes exist. First, they prevent "free-loading" classifiers that do something inessential to earn reward, or sets of classifiers that merely support each other. Second, reward flows into the whole system. If strengths grew, proportional selection would get skewed over time. That, of course, raises a danger – if tax causes strength to flow faster than reward flows in, the system can simply die.

3.4.1.2 The Discovery Component – Niching

Finally, in any LCS, the *rule discovery component* is governed by the GA. In SCS, a standard steady-state GA is used, like the canonical described in section 3.3, except that it has some variants in order to promote the formation of well-adapted sets of rules instead of a best single-rule, through a special type of *crowding* scheme, where replacement candidates are chosen from a low-performance subpopulation on the basis of similarity to the children being inserted. This will be explained in more detail in the following section.

An LCS searches for the smallest number of rules that solve the task being modelled in the best possible way with the idea of generalising well enough the learned behaviour when similar conditions occur. In summary, it tries to look for a "concise, accurate, and robust concept description, where concept description is a group of rules" [Horn *et al.* 94]. In this sense, any LCS, even in its simplest form, offers an implicit niching mechanism due to its multi-objective nature. If no niching method is applied, the GA operator *selection* used alone over the entire population (i.e. roulette wheel as commonly used in single-criterion optimisation problems) would ruin this goal by promoting the creation of a single rule (probably a very general one – with many # symbols) to be used when several criteria occur, making it impossible to obtain the well-adapted sets of rules (more specific to certain conditions) we are searching for. Although several ways have been reported to accomplish this goal, what exactly is best to do, is still an issue of debate and experimentation in the LCS community.

For example, one of the first niching mechanisms proposed was *preselection*, where the children created simply replace one of their parents [Cavicchio 70, Mahfoud 92], followed by a more elaborate two-fold process known as *crowding* [DeJong 75], where on one hand, a proportion of the population is chosen by roulette wheel selection (RWS) to undergo crossover and mutation, creating children which will be inserted back in the population after a second process is made for every child that will be inserted: a *crowding factor* (CF) number of times, CF individuals are selected randomly from the full population, and each child will be compared to them, taking the place of the most similar to it. For a more detailed description of these two niching techniques refer to [Mahfoud 92].

A third approach [Horn *et al.* 94], incorporates a niching mechanism known as *implicit fitness sharing* into the LCS by forcing competing rules to share their reward (similar to the explicit fitness sharing used in GAs). With this method, high quality and diverse niches were maintained virtually indefinitely. One important challenge LCS face today is the need to innovate the GA component. In [Goldberg 00], David Goldberg comments that "the competent GAs that have been developed over the last decade should be adapted to LCS usage and this should benefit the search for appropriate rules in difficult problems."

3.4.2 Accuracy Based Classifier Systems

The most widely used accuracy based classifier system is called XCS (eXtended Classifier System), developed by Stewart Wilson, in response to some drawbacks that were considered important about the traditional LCS. It is outside the scope of this thesis to provide a full description of XCS, however, in order to understand the main differences between the strength-based and the accuracy-based classifier approaches, a brief description will be given in this section. For a complete review of XCS refer to [Wilson 95], and for a full comparison between the two approaches see [Kovacs 00].

3.4.2.1 Motivation for Accuracy

The main reasons for shifting from strength to accuracy based LCS is given in [Wilson 95]. They will be summarised in this section.

The classifier's strength value, apart from being useful in the performance component by estimating the payoff that the classifier will receive when its action is chosen - helping the system to learn to choose the most remunerative action, – strength is also used as the fitness measure by the GA in the LCS's discovery component. However, this double function of the strength parameter allows the best performers to lead the search, as they are more likely to be reproduced and selected for crossover and mutation. As a result, the classifiers of high-payoff niches will tend to take over the rest of the population. But why is this undesirable? Because, for example, classifiers that are vital but fire rarely will earn little strength, and so will be candidates for replacement.

In order to tackle this problem, a sharing technique can be implemented by dividing the payoff by all active classifiers instead of giving them the full value. The payoff is predicted by the total of such shared strength among the classifiers with the same action. However, as a classifier will match in several distinct occasions, the single value of strength associated with each classifier will become unclear. Also, sharing still favours the stronger niches, which causes a problem in sequential tasks where rewards are not given in every situation. In this case the classifiers which tend to help others to receive a reward on a later stage will become weaker and eventually disappear, not allowing the formation of long chains. To alleviate this problem, Booker [Booker 82] restricted the GA to the match set instead of panmictically using the whole population. This means that the differences in payoff between match sets will not affect the classifier's selection chances and therefore competition will be restricted to classifiers within a niche (in this sense, niche refers to the matching set).

However, even this solution does not seem to stop a common problem most traditional classifier systems share: the vast creation of overgenerals – classifiers with more # symbols than the ones required for a good generalisation to occur. Overgenerals ('guessers,' called by Wilson) will tend to reproduce more as they tend to match more often by guessing right more times than other classifiers that could be more accurate but that do not match that often. Furthermore, Wilson states that "under payoff-based fitness, there appears to be no clear tendency or, indeed, theoretical reason for accurate generalizations to evolve" [Wilson 95].

The development of XCS was also inspired from the Reinforcement Learning (RL) community [Sutton 91], where the goal is to build relatively complete mappings X x A => P from the product set of situations and actions to payoffs. With this view, traditional LCS as stated by [Wilson 95], tries "to discover the best rule in each niche without worrying too much about knowing the payoff consequences of every possible action. However, should a sub-optimal rule be converged upon as a consequence of incomplete exploration, it may be difficult for the standard system to discover and switch to a better one. If, on the other hand - as in reinforcement learning – the system were oriented toward learning relatively complete maps of the consequences of each action in each niche, then determining the most remunerative action would be straightforward. For this, it seemed logical to base fitness on some measure of accuracy."

Given these problems and inspirations, Wilson first offered an improved version of the traditional approach, called ZCS ("Zeroth level" Classifier System, [Wilson 94]), which closely resembles the traditional approach. However, later on he proposed a second one, XCS, which formally develops the idea of accuracy based classifier systems (fitness-based accuracy) motivated by the problems explained above, including the inability of estimated payoff to distinguish between accurate and overgeneral classifiers. In this view, "the system might need to be bigger because the number of accurate classifiers could exceed the number of highly remunerative ones. However, overgeneral rules would be suppressed" [Wilson 95].

3.4.2.2 Description of XCS

Similarly to SCS, in XCS information flows from the environment through the detectors, and the classifiers (rules of the form $< condition >$: $< action >$) constitute population [P]. Instead of having a measure of *strength* associated with each rule, there are three main parameters (taken from [Wilson 95]):

1. *Prediction, p*: average of the payoff received when that classifier's action controlled the system.

2. *Prediction error*, ε: average of the error in the prediction parameter.

3. *Fitness, F*: measure of the prediction's accuracy, defined as the inverse function of the classifier's average prediction error.

The first two of these parameters are used in the performance component of the LCS. The third one is used in the discovery component, by the GA, which is executed in niches defined by the match set [Booker 82], instead of panmictically as described above.

3.4.2.3 Performance Component

The system checks the input coming from the environment and a *match set* [M] is created with all the classifiers that match that input, in the standard way, by checking the condition of every bit. With this set's associated actions, a *system prediction* $P(a_i)$ is formed according to a fitness-weighted average of the predictions of classifiers advocating a_i and placed in the *prediction array*. Normally the system selects deterministically the action with the highest prediction, but it can also be selected probabilistically, according to P(a_i) or even randomly in some cases. Once an action is selected, an *action set* is created with all the classifiers in [M] that have the same action. The action is performed by the effectors and a reward is returned by the environment.

3.4.2.4 Reinforcement Component

This component updates the parameters p, ε and F from the previous step's action set. The p values are adjusted with the Q-learning technique [Watkins 89], resulting in a quantity called P which is used to adjust the predictions p_j of the classifiers in the previous action set using the standard Widrow-Hoff delta rule [Wilson 94]. For more details in these procedures, refer to [Wilson 95]. The Widrow-Hoff delta rule is also used to update ε and part of the adjustment of F.

3.4.2.5 The Discovery Component

The GA acts on the match set [M], first by selecting two classifiers with probabilities proportional to their fitness and then by performing crossover and mutation as usual. The classifiers to be removed from [P] are selected either stochastically by (i) choosing the classifiers proportional to the match set size estimate, with the aim of keeping similar niche sizes, or by (ii) doing the same as the previous step, except that if the fitness is lower than certain value, this classifier with lower fitness is deleted with a higher probability. As the aim here is to allocate resources equally amongst the niches, the GA acts on the match set only when the number of time-steps after it was previously run exceeds a certain threshold. The purpose of this threshold is to ensure a constant reproduction rate per match set.

In addition to the GA, there is a *covering* mechanism [Wilson 85] in the discovery component, which is used to add a new classifier when there is no actual match of the current situation. This new classifier is normally added in [P] with a condition matching the input and a random action, replacing a classifier in the same way as described with the GA insertions. In practice, covering occurs more often at the beginning of the run, as [P] is normally empty to start with and in this way new classifiers are being added on every time-step until more matches occur.

3.4.2.6 The End Product

The idea behind XCS is that the end product of combining accuracy and a niche GA results in a complete and accurate mapping X x A => P from inputs and actions to payoff predictions. Further, as Wilson stated "XCS tends to evolve classifiers that are maximally general subject to an accuracy criterion, so that the mapping gains representational efficiency. In traditional classifier systems there is in theory no adaptive pressure toward accurate generalization, and in fact accurate generalised classifiers have rarely been exhibited, except in studies using payoff regimes biased toward formally general classifiers (e.g. [Wilson 87])".

3.4.3 Applications: Strength Versus Accuracy

The classic paper entitled "A Critical Review of Classifier Systems" [Wilson & Goldberg 89] gives a good summary of the unsolved problems and new challenges that LCS faced in the late 80s. Since then, there have been great accomplishments in theoretical aspects (mapping performance and generalisation), of XCS solving a variety of single-step environments such as the boolean multiplexer and sequential environments (multi-step) like the woods-type of problems. However, most economic models involving the use of classifier systems still use SCS (some with a small number of variants in order to cope with the problem being modelled, but the major learning device is of SCS type). I believe this could be attributed mainly to the following reasons:

1. Before stepping back for a decade from LCS research trying to solve first the GA paradigm, David Goldberg made a very good job by providing not only an excellent introduction, but also a clear and efficient way for implementing SCS in Pascal. In addition to Goldberg's efforts, in 1994, Jörg Heitkötter made available a C translation and extension [Heitkötter 94a, Heitkötter 94b] of Goldberg's SCS, as presented in Appendix C of his seminal book [Goldberg 89]. Unfortunately, by that time, neither a good introductory book, nor any source code, were provided by XCS practitioners; at least until 1998, when Alwyn Barry developed his first version of Wilson's XCS [Barry 01]. A full-working (and excellent) im-

plementation of XCS was not made public until 1999 by Martin V. Butz at the IlliGAL laboratory [Butz 99, Butz & Wilson 00, Butz 00]).

2. In parallel to this, John Holland was involved with economists and other scientists in the development of a truly artificial economy [Arthur *et al.* 97, LeBaron *et al.* 99, Palmer *et al.* 94, Palmer *et al.* 99] with ideas of [Arthur 92], in which the market participants' learning process was modelled with LCS, in the very *Holland* way (this model, known as the SFI artificial stock market, is described in more detail in section 4.2). Due to its multi-disciplinary nature, this might have also brought the attention of researchers from other areas who were interested in the applications arena to use SCS in their work. Now, over a dozen papers describe the SFI stock market, some of which are more modern versions which describe new experiments and interesting evolutionary aspects such as the effect of changing the system's learning rate, etc. These are reported in [Joshi & Bedau 98, Joshi *et al.* 98, Joshi *et al.* 99].

3. The reported difficulties of traditional LCS made during the late 80s did not appear to have a negative effect in the development of real-life applications, perhaps because it was not fully understood how exactly these applications would be affected by such difficulties, or because they were not applicable (or easily solved) to these types of problems. Chapter 5 provides the reader with a number of applications involving LCS. The "earlier difficulties", summarised by Wilson in [Holland *et al.* 00] as competition vs. cooperation (specifically with chain formation) and performance vs. generality (uneven distribution of payoffs causing overgeneralisation), were the main motivation to the development of the accuracy based classifier system XCS. However, as it was just mentioned, such difficulties could perhaps be easily overcome and do not necessarily apply to the development of economic and other type of real-life applications. This argument will be brought forward later on, when describing the proposed model and technologies in Chapter 7.

3.5 Genetic Programming

Despite being one of the most relevant areas to the work presented here, due to space limitations, only a modest description of genetic programming (GP) will be given in this thesis. However, a number of references covering general aspects of GP will be given to the interested reader in the present section.

John Koza [Koza 92, Koza 94] developed a method called Genetic Programming, designed to evolve Lisp programs encharged of solving certain problems. As programs in Lisp can be expressed as parse trees, a GA can work on them to improve the solutions obtained.

The basic Koza's GP algorithm works as follows (taken from [Mitchell 96]:

1. Choose a set of possible functions and terminals for the program. In this step a reasonable set must be chosen, for example, the function (also called non-terminals) set could be as simple as {+, -, *, /, SQRT} and the terminal set could be just the variable $\{A\}$.

2. Generate a random population of trees using the set of possible functions and terminals. Such random trees: (i) must be syntactically correct, e.g. the number of branches extending from function must equal the number of arguments taken by that function, and (ii) can be of different sizes (having different number of nodes and levels in the trees).

3. Calculate the fitness of each program in the population by running with a set of input/output pairs, known as "fitness cases".

4. Select probabilistically in proportion to fitness (i.e. the best) 10% of the population and copy them into the next population without modifying them. Apply selection, crossover, and mutation to the remaining 90% of the population. Crossover works as in GAs, by exchanging the subtrees of each parent. Note that here the size of the trees in this case will be variable, not fixed as in GAs. Mutation works in a slightly different way. A random point of the tree is chosen and any subtrees derived from it are replaced by a random subtree.

Koza's original method [Koza 92] uses large initial populations with the idea that the members should contain a good set of building blocks so that only crossover is applied, eliminating the need for mutations.

After following these steps, 3 and 4 are repeated for a number of generations, until, like with GAs, a certain criterion is met that stops the program. This method is also known as generational GP. Another type of GP which is gaining ground in the research community is the steady-state or tournament selection model, characterised by not running in a generational manner. This method should also be simpler to implement, the basic steady-state algorithm, taken from [Banzhaf *et al.* 98], is as follows:

1. Initialise the population
2. Randomly choose a subset of the population to take part in the tournament (the competitors)
3. Evaluate the fitness value of each competitor in the tournament.
4. Select the winner(s)
5. Apply genetic operators to the winner(s).
6. Replace the losers in the tournament with the results of the application of the genetic operators to the winners.
7. Repeat steps 2-7 until termination criterion is met.
8. Choose the best individual in the population as the output from the algorithm.

Both of these methods offer a number of advantages, for example, GP allows the size and complexity of the evolved trees to increase, rather than keeping the individuals fixed as GAs do. However, as stated by [Mitchell 96], the original method does not include a "mechanism for automatically chunking parts of a program so they will not be split up under crossover, and no mechanism for automatically generating hierarchical structures (e.g., a main program with subroutines) that would facilitate the creation of high-level primitives from built-in low-level primitives." The area of GP research is growing rapidly and new methods have been devised to overcome these drawbacks.

After these brief descriptions, Chapter 4 follows with a number of interesting applications involving GP in forecasting and decision making.

Chapter 4

Approaches to Computational Economics

"I have done that," says my memory. "I cannot have done that," says my pride, and remains inexorable. Eventually-memory yields. Friedrich Nietzsche, *Beyond Good and Evil.*

4.1 Evolutionary Economics

From pure analytical to heavily computational approaches, general interest has grown lately towards the creation of methods and theories that combine adaptive, complex and evolutionary elements trying to understand at least some of the many empirical puzzles financial markets present. For a good review of these methods refer to [LeBaron 95]. Such approaches differ from each other in the way they tackle three of the main factors involved: the economic environment, learning process and trading mechanism. As a result, various directions of research have emerged and although other models will also be explored, the main focus here will be given to models whose economic environment is a type of financial market, such as the stock or foreign exchange markets. Models of this type follow two broad directions, depending on the environmental data they use: the first one navigates entirely in an artificial framework and its main concern resides in creating artificial data, analysing it and comparing it with phenomena present in real markets; and in the second one real data is examined with the purpose of creating better forecasting models.

The first direction is based in models of the economy that are composed mainly of a number of agents (of the same or different types) that interact and adjust to each other over time by choosing their actions consciously and according to the other agents' expectations about the future. In a number of these models, agents learn and adapt to solve the hard problem of forecasting by creating, exploring and improving market strategies and the emergence of various economic elements such as institutions, behaviours and technologies is observed. The activity generated is then compared with actual financial data. It has been observed that under some circumstances, parts of the economy being modelled are attracted to an equilibrium state, while others never reach this state and therefore continually evolve and never settle.

This thesis is concerned with modern computational approaches where artificial intelligence techniques are applied to model adaptive behaviour in financial markets. In this category, a vast number of adaptive models have been developed recently, most of which use different approaches to model adaptive behaviour. For example, the adaptive element of such models has been commonly represented by neural networks, genetic algorithms, classifier systems, fuzzy systems, various machine learning techniques, genetic programming, etc., and although it is outside the scope of this thesis to describe the majority of artificial intelligence applications in finance, a brief description of some of them is given in Chapter 5. An interesting point of view of many of these models is that they share a dynamic way of looking at the market. Throughout this thesis, emphasis will be given to them, with special attention to the LCS models.

As mentioned earlier, a great number of new approaches have been rising lately. In some of them, labelled under the title *behavioural economics*, the main assumption, as described by ([Bass 99], page 270), is that "there are patterns in markets because traders behave in groups. They obey universal laws of human psychology. They share common elements in their backgrounds and biases; common information sources; and common emotions of fear and greed. They are not fully rational. They are overconfident. They have poor intuitive grasp of statistics. They act in herds. They make decisions in a world endowed neither with perfect knowledge nor perfect competition

.... Financial evolution does not proceed to an optimum point and then just sit there. Traders exploit patterns in the market to make profits, causing old patterns to diminish and new patterns to form. As in biological systems, this results not in the attainment of some static endpoint, but in the evolution of successively richer and more complex strategies."

Doyne Farmer's model lies in this evolutionary economics category. His model, described in [Farmer 98], consists of interacting heterogeneous agents of different nature: trend followers (technical), value investors (fundamentalists), all of which, by their characteristic behaviours and interactions, create interesting price dynamics in the economy such as bull and bear markets. Like in real markets, some of these traders profit from using certain strategies, while others go broke from too much losing. The evolutionary component resides in that the population of strategies evolves with time and therefore the market is an evolving, self-organised process. The agent behaviours displayed by value investors depends on what value these agents assign to the security, according to some fundamental information. If they believe it is under priced, they buy, otherwise, sell. The technical trading agent bases its decisions on historical data, such as past prices. As the model allows different traders to perceive different values, they can cause excess volatility, creating good opportunities for trend followers. Motivation of this work and more advanced topics about the model can be found in [Farmer 99], [Farmer & Lo 99] and [Farmer & Joshi 00].

4.2 Artificial Stock Markets

It is important to give special attention to the work done in artificial markets as this is the area where most concepts are borrowed from, as well as the main motivation of this research. This section is divided in two subsections describing important artificial markets that have been developed and which share a common objective: incorporating learning into the artificial agents in the economic system being modelled, which is viewed through an evolutionary framework rather than as a stationary world – both using artificial intelligence techniques to implement such learning process. The first one to describe is inspired by John Holland's ideas, it represents the adaptive behaviour through classifier systems and genetic algorithms, as explained in [Holland & Miller 91]; and the second one, inspired by ideas of John Koza, models human learning through genetic programming. Although both are interesting in their own way and deserve special attention, more detail will be given to the first one due to the similarities in the way reasoning is modelled – inductively – to the work presented here. For an excellent review of the origins of these two approaches in the development of artificial adaptive economic agents in evolutionary economics, see [Chen 01].

In addition to these models, there is a great number of computational market approaches composed of a large number of heterogeneous interacting agents. Blake LeBaron gives a good summary of six early agent-based economic models which form the foundations of artificial stock markets in [LeBaron 00], which he complements in [LeBaron 99] with promising directions of related research along with full descriptions on building the adaptive agents, the trading mechanism and the level of rationality in an ASM.

4.2.1 The Santa Fe Artificial Stock Market

In 1988, John Holland and Brian Arthur started to develop an artificial stock market, one of the very earliest artificial, agent-based financial market models by exchanging ideas about economics and adaptive systems while they were at the Santa Fe Institute. Later on, a physicist, Richard Palmer and a finance expert Paul Tayler joined the project, followed by Blake LeBaron who is a financial theorist in economics. Given that the ideas of their model, widely known as the Santa Fe Artificial Stock Market (SFI model), are the main motivation of the work presented in this thesis, they will be explained in some detail (most of the descriptions that follow were taken from [Arthur 94b].

An interesting assumption of the SFI model is that in this market there is no incoming information from the real stock market, it is a truly artificial market evolving inside the machine which is composed of a number of artificial agents buying and selling a stock from one another and at any given time. Therefore it is possible to see how the stock's price is changing, the dividend, who is buying and selling, who is making money and who is not, who is in the market and who is out, and so on. The price is calculated by a market specialist who sets the price to clear the market, through bids and offers as it happens in actual stock markets.

According to the authors, the initial modelling question to tackle was: if the agents cannot form their expectations deductively, how can they form them? They decided to follow some aspects of modern cognitive theory about how actual human beings behave in such situations. With this in mind, they allowed the agents to look at various hypothetical models for forecasting and to be able to test them constantly. Each of these hypotheses has a prediction associated with it. At any stage each agent uses the most accurate of its hypotheses, and buys or sells according to these suggestions. The agents learn in two ways: they learn which of their forecasting hypotheses are more accurate, and they continually toss out ones that don't work and replace these using a genetic algorithm. So they are learning to recognise patterns they are collectively creating, and this in turn collectively creates new patterns in the stock price, which they can form fresh hypotheses about. This kind of behaviour – bringing in hypotheses, testing them, and occasionally replacing them – is called induction, as defined in Chapter 2. These agents, therefore, use inductive rationality, a very realistic form of behaviour.

Then the second question was to investigate whether such an artificial market would converge to the rational expectations equilibrium or whether it would show some other behaviour. In response to this question, the authors found that two different regimes emerged. One, which started with sets of predictive hypotheses close to rational expectation, was consistent with the rational expectations regime, creating price sequences near rational expectations. The second emerging behaviour, characterised by a faster rate of learning, was called the "complex regime." In this regime the behaviour of the market changed, and mutually reinforcing predictions were observed. In [Arthur 94b], Bryan Arthur explains:

> "Imagine we have a 100 artificial agents each using 60 different prediction formulas, so that there is a universe of some 6,000 predictors. Some of these predictors that emerge are mutually reinforcing, some are mutually negating. Suppose many predictors arise that say the stock price cycles up and down over time. Such predictors would be mutually negating because they will cause agents to buy in at the bottom of the cycle, and sell at the top of the cycle, mutually negating profits, and therefore eventually disappearing from the population of predictors. But if a subset

> of predictors emerged by chance that said 'the price will rise next period if it has risen in the last three periods,' and there were enough of these, they would cause agents to buy, which on average would cause the price to rise, reinforcing such a sub-population. Such subsets could then take off, and become embedded in the population of predictors. This was what indeed happened in the complex regime, endowing it with much richer set of behaviors. Another way to express this is that our artificial traders had discovered forms of technical trading that worked. They were using, with success, predictions based upon past price patterns. And so technical trading was emergent in our artificial stock market. This emergence of subsets of mutually reinforcing elements, strangely enough, is reminiscent of the origin of life, where the emergence of subpopulations of RNA in correct combinations allows them to become mutually enforcing."

This complex regime moves from periods of high to low volatility. But how can this property emerge? An explanation is that in the artificial market, every so often some number of investors discover a new way to do better in the market and therefore change their buying and selling behaviour, which, in turn, causes the market to change, even if slightly, possibly causing other investors in turn to change. "Avalanches of change sweep through the market, on all scales, large and small. Thus emerge periods of change triggering further change, periods of high volatility, followed by periods when little changes and little needs to be changed, periods of quiescence. This is GARCH behavior" [Arthur 94b]. In addition to [Arthur 94b] and [Arthur 92], more descriptions of the model have been published in [Arthur *et al.* 97, LeBaron *et al.* 99, Palmer *et al.* 94, Palmer *et al.* 99], and newer versions describing results of experiments and related work can be found in [Joshi & Bedau 98, Joshi *et al.* 98, Joshi *et al.* 99]. These papers will be described in more detail throughout this thesis due to the similarities with the proposed model.

4.2.2 The Taiwan Artificial Stock Market

In the area of financial markets, there is a vast area of research using GP as the basis for adaptive behaviour. Similarly to LCS, these models using GP can be divided into two groups, the artificial market models and the predictive systems. The latter group will be explained in section 4.3. Within the artificial framework, the work of Shu-Heng

Chen and his research group at the Economics Department of the National Chengchi University in Taipei, Taiwan, together with the work of Thomas Lux and Michelle Marchesi, is very extensive, it encompasses a great number of publications, and is constantly growing. This section only covers the basics of such models.

The Lux and Marchesi model focuses on financial market analysis through a model composed of interacting agents of two different types: *fundamentalists* (who trade on the stock if they believe that the price is below or above its fundamental value) and *noise traders* (who identify trends and patterns and consider the behaviour of other traders as their source of information), the noise traders can also be of optimistic and pessimistic nature. The model represents an artificial market where stock prices are created endogenously, that is, driven by the agents rather than coming from external sources, allowing the agents to shift from one group to the other depending on the dynamics of the market. In [Lux & Marchesi 99], the authors have shown, against the EMH (that the changes in prices are consequences of only unbiased, incoming news about the future earnings of the stock), that the market exhibits periods of high volatility which is caused by the interactions of the agents from having different beliefs and market strategies, not only by changes in the inputs that normally influence prices. Furthermore, as Lux and Marchesi pointed out in [Peterson 99], "in periods of high volatility we also find a large fraction of agents in the noise trader group. Theoretical analysis shows that a critical value for the number of noise traders exists where the system loses stability."

More detailed descriptions of their model can be found in [Lux & Marchesi 99], [Lux & Marchesi 00]. In addition, [Chen *et al.* 01], present an investigation of the time series behaviour of simulated data of the model, extending their earlier findings that the artificial market exhibits real market properties, generated by the trading process itself rather than by the exogenous news arrival process. While this work (using the prototype of Chen's artificial market) concentrates more on the time series properties of the artificial market, the work of Shu-Heng Chen and Chia-Hsuan Yeh resides in the area of agent based modelling, where the adaptive component is mainly represented by GP.

Chen and Yeh have done extensive research in the areas of artificial intelligence and economics, but due to space limitations, only three lines of their research will be addressed here. The first one, motivated by the SFI model (described in [Chen & Yeh 97]), involves modelling speculations about the speculations of others using GP. In their model, the authors found patterns that are consistent with findings in experimental markets with human subjects, including a destabilising, noisy behaviour, arising from the *GP-speculator*. They observed that in the early stage of evolution, speculators can help stabilise the economy when appropriate financial regulations (constrains) are allowed.

The second line of research to mention here deals with a new type of architecture to study artificial stock markets: an agent-based model of a *school*, which was devised in response to a criticism about the difficulty of implementing, in social processes, the differences between a biological system's phenotype and genotype – in this context, actions correspond to phenotype and strategies to genotype. The authors argue that it would not be interesting enough if agents could only learn to imitate the actions of others (but not the strategies) as it would represent only a minor part of the interactions between the agents. The important issue would be to learn the strategies hidden behind those actions.

In their model, *school* is a collection of studies (strategies) of market participants and it is treated as a single population GP. It can be viewed as a procedure to map phenotype to genotype, or as the authors describe it, a procedure "to uncover the secret of success" [Chen & Yeh 99] (called the *business school*). The agents of the model, or *traders*, are allowed to consult the *school*, just as if it was a library in order to improve their strategies, and both, traders and school, coevolve in this artificial market.

Analysis of the outcome of this model reveals that prices and returns are not normally distributed. In addition, this artificial market works in accordance to their null hypothesis that prices do follow a random walk and that returns are independently and identically distributed (iid), corresponding to the classical EMH. The authors point out that "what is interesting though is that this iid series was generated by traders, who do not believe in the EMH at all. In fact, our study indicates that many of our traders were able to find useful signals quite often from business school, even though these signals were short-lived" [Chen & Yeh 99].

Finally, the third line of research addresses the EMH from a biological point of view. With a GP-based approach, in [Chen & Yeh 96b] the authors formalise the notion of *predictability* in the EMH by forecasting three chaotic time series and a small sample of the S&P 500 and the TAIEX indexes. Their results show that GP can beat the random walk by 50%, while other linear methods are unable to do so. However, according to the authors, these results also suggest that "while the short-term nonlinear regularities might still exist, the search costs of discovering them might be too high to make the exploitation of these regularities profitable, hence the efficient market hypothesis is sustained" [Chen & Yeh 96b], which brings the question of whether *predictability* implies *profitability*.

To answer this question, the authors addressed the notion of *profitability* in [Chen & Yeh 96a], by measuring search cost in terms of computational cost as well as risk in an experiment involving longer series of the the S&P 500 index. They provide interesting findings:

- Training with small samples shows there is room for profits, but they found that as the rule needs regular updates (in a weekly basis rather than in a monthly or yearly basis), the **computational cost** associated with these updates "might be so high that, in terms of net profits, there is no significant difference between long-range and short-range forecasting" [Chen & Yeh 96a]. In overall, results with the large sample show that the EMH is sustained, while small sample results show the EMH is not competitive.

- Considering another type of cost, a measure of **risk** was defined as the difference in the post-sample MAPE [1] between the best and worst GP-search. They show that the short-range forecasting risk is much higher (up to five times higher) than the risk with large samples.

With respect to the EMH, it has been made clear in the paper, that biologically inspired approaches can help to re-formalise the EMH by investigating more deeply

[1] Mean Absolute Percentage Error (MAPE) is defined as follows (n is the sample size and $\hat{R}_i$ is the prediction value of R_i):

$$MAPE = \sum_{i=1}^{n} \frac{|\hat{R}_i - R_i|}{|R_i|} \tag{4.1}$$

some of the following issues: (i) whether patterns can be discovered at reasonable costs, (ii) whether the retraining or updating of certain advanced techniques beating the random walk is competitive and (iii) whether this can be done both, at a reasonable cost and continuously – for example, by a generic technology like neural networks rather than just by *one* specific recurrent network.

As explained in Chapter 2, instead of proving whether the EMH is correct or not, these are the types of questions this thesis intends to address.

4.3 Predictive and Decision Making Models

As real situations such as stock markets are almost intractable analytically due to the highly non-linear underlying processes involved, the proposed model involves computational modelling when viewed from an evolutionary perspective. The idea here is to develop a model to try to tackle difficult elements such as people's *non-linear behaviours* [2] because it is well known that in financial markets people act in a non-linear fashion. In other words, the same-seized cause can have different-sized effects, depending on the circumstances.

This property of non-linearity is present because at any given time, traders receive vast amounts of information, which arrives most likely in numeric form and which can also be noisy. This information is then used to extract relevant pieces of information before making any investment decisions. Depending on the type of investor, these decisions frequently rely on the integration of various statistical measures compressing much of the data with news events and other information such as analyst recommendations. Visualising techniques such as graphs and bar charts also help them understand current situations the market is going through.

Considering that such decisions involve taking into account significant nonlinear relationships among the components of the data, evolutionary methods such as neural networks and genetic algorithms have been employed to build systems to improve upon the judgement of traders, to derive better automated trading strategies, and to optimise portfolio management. This is a very complex task because it is well known that even

[2] Meaning that effects (actions) are not proportional to causes (motives).

though traders can factor in such relationships in the analysis, usually they are not able to explain their decisions or point out the significant factors and relationships they consider. In many occasions, specially when they are successful, they do not share such information due to secrecy matters. Also, due to the fact that the mathematics of cause and effect are complicated in extreme, the ability to forecast markets depends, up to a certain point, on the success of non-linear methods.

For example, as pointed out by Guido Deboeck in [Deboeck 94], one promising idea is "to combine expert know-how, which is fuzzy, with findings from empirical time series, which are patterns that can be detected by NNs. As economists say, the true probability distribution for future outcomes is based on 'convoluting' together the probability distributions from experts and from empirical knowledge. So we can do better if we make use of both kinds of knowledge."

In addition, a number of conferences report many papers on financial time series analysis and prediction. Specially, during the early nineties, when the use of neural networks appeared to be an area of great potential in securities forecasting, dozens of papers using neural networks were published. A good review of these results can be found in [Tri93, Ref95, Ref96, Abu99], and for a good summary of techniques, see [Masters 95]. Although later on the revolution came to an end, what that era brought to the surface was perhaps the need to address the fallacies encountered with traditional approaches and the need to search for non-linear models capable of handling the complex task of forecasting. Too much was expected from NNs, and soon people were going to realise that, as pointed out by Halbert White when reporting his results, "in either case the implication is the same: the neural network is not a money machine" [White 88].

More recently, as a result of the fast growth of the Internet, other approaches started to emerge, such as the work on text in finance [Wuthrich *et al.* 98, Thomas 00, Thomas & Sycara 00]. Much of this effort concentrates on finding profitable trading rules from the analysis of text information of financial interest publicly available on the Internet – chat boards or news stories.

A review of specific artificial intelligence applications in financial markets will be given in Chapter 5. This chapter concentrates in the most relevant approaches.

Because neural networks have been used very successfully in commercial applications for predictive purposes (see, for example products offered by Olsen Group and HNC Software), as well as being widely used in financial research, the following section shows an example that covers the basics of time series analysis and predictions using neural networks.

4.3.1 Time Series Analysis and Predictions

The change in time over stock prices, foreign exchanges, interest rates, and other factors interchanged in financial markets, present a very interesting prediction problem, mainly because it is impossible to solve it perfectly. If one tries to solve this problem with a NN, it is important to assume that there is a dynamic system that governs the global evolution of the market. At this moment, our understanding of the mechanisms that generate this time series is very limited. The generator of the series is not static, but changes over time in a much lower scale. In other words, the rules change over time, but not as fast as the data we want to predict. This addresses the impossibility of long time predictions, even if we didn't have the problem of error accumulation.

Given the peculiar nature of the problem (non-linearity and noisy transactions), defining the data pre-processing phase becomes crucial. It is necessary to start with data of high quality, and in enough quantities, and also using useful transformations so they can be useful to the NN, otherwise a common problem could arise when training the net with non-transformed data: they tend to predict the last value of the input node ($X_{t+1} = X_t$), concluding that the best prediction of tomorrow's price is today's price, which is consistent with the EMT.

A stationary time series is defined as one that has a mean that fluctuates consistently about some fixed level over time. Further, it has an essentially constant variance and a constant autocorrelation function over time. However, literature indicates that most of the actual time series of interest are non-stationary. In fact, most series actually encountered in industry or business exhibit homogeneous non stationarity, which is characterised by a fluctuating mean, but the broad behaviour of the series (the variance) is stable over time [Reynolds *et al.* 95]. Most time series that arise in economic and business applications are homogeneous non stationary processes.

Also, time series analysis and predictions in different areas of application varies significantly. If one is analysing weather time series, the causes producing weather changes are roughly the same now as they were 100 years ago. Sun spots too. But the difference with financial markets is that they change so rapidly over time that there is no current method reliable enough to take over other alternate methods. One advantage about NNs is that they are trainable, allowing to add more information relevant to the changes in prediction.

At this point it is important to note that from the EMT is derived the well known *buy-and-hold* strategy. This means, as explained by its name, that you buy a certain amount of stocks and wait. The most likely event that will happen is that the stock will behave as the movement of the rest of the market, and that the earning will be proportional to the movement of the global market index such as the Dow Jones Industrial Average Index (DJIA) in the New York market.

Due to the fact that many problems require analysis of one or multiple variables changing over time, as pointed out by David Lowe, "one of the central roles of science in the study of naturally occurring phenomena is in forecasting: given knowledge about a system and its past behaviour, what predictions can be made regarding its future evolution?" [Lowe & Webb 91]. The idea is to create a model with the known behaviour of a variable with the hope of explaining its behaviour under other unknown circumstances, for example, the future.

Sometimes the general model of the behaviour of the variable is known. It can be linear ($y = ax + b$), logarithmic ($y = a + blogx$), or any other. Yet in other circumstances it is not known. It is not always possible to generate a precise model of the data because we don't understand or know anything about the governing rules. In these cases of uncertainty with respect to the model, it is recommended to analyse the data directly with a more generalised technique.

4.4 Forecasting Using Neural Networks – A Simple Example

It is not possible to know exactly when and how NNs have been used in time series of financial markets analysis, but we know that the range of applications is very wide. NNs have been used for predictions in different time horizons (from intra-day to monthly decisions) over different data (foreign exchange, interest rates, stock prices, etc.) and for different objectives (portfolio optimisation, to prove economical theories, to decide if buying or selling, etc.). The various systems developed also give different precisions, some give the output in binary form (up or down), while others' output are real numbers, approximating the series to the future. The decision over the input data given to the NN depends upon the application. Many applications choose a commonly used method known as the *time delay method*, which can be followed precisely or with some modifications.

In this method, the network can be fed with several inputs, each one representing one specific characteristic of the variables to analyse. For example, it can be a financial indicator, or the percent earnings of a company in a certain period. No doubt the most common problem tackled using NNs is that of Time Series Prediction a few periods over the future. To increase the accuracy of the net, it is always good to use all the additional variables that can help to get the prediction of the price of the stock such as volume of transactions, a general Index (DJIA, S&P, Nasdaq, etc.), and also other indicators (which are transformations of these variables).

The transformation of the input data can help considerably to improve the prediction accuracy. This means pre-processing the input data with some sort of transformation which can be as simple as $sin(x)$, to a more complex functions used in financial analysis as indicators. The idea behind this pre-processing of data is to provide the NN with information about the data which otherwise it would be very difficult to distinguish. For example, to eliminate noise so the net can work with a cleaner series.

Sometimes NNs are used as a tool along with other statistical methods. They are combined with expert systems, or are used to support some other techniques. For example, a NN trained to recognise triangles in time series (a type of oscillation) with examples given a priori by an expert, could be able to recognise this trend over the future, but in the analysis of the triangles other methods are used.

Modular NNs have also been used. These are composed by sub-nets which jointly contribute to a better solution than the individual net. A simple model averages the results of identical nets (but with different initial conditions), trying to avoid the problem of noise originated by inadequate initial weight values. Another model, known as hierarchical, is based in training each net with the error function of the previous net of the hierarchy [Deppisch *et al.* 91]. There are other models in which each net works on an specific task (for example, predicting direction, magnitude and different time horizons) to get to a more precise solution.

Forecasting the future values of sequences of observations is, in many ways, ideally suited for NNs. Large amounts of data may be available, and the underlying relationships are often nonlinear and unknown. NNs, mostly of the standard Back-Propagation type, have been used with success in many forecasting applications, depending on the characteristics of the process being forecasted.

NNs may be viewed as a parametrised statistical methods, where the parameters are adjusted in function of the data. In other words, they are trainable models with the ability to adjust to a variety of functions. It has been shown in many cases that NNs approximate well under noisy conditions (noise is a random component added to the data). In addition to this, NNs exhibit nonlinear behaviour (the effect is not proportional to the cause). These, among others, put NNs as an ideal candidate for problems in which the internal model is unknown. NNs are capable of approximating transformations with any number of input parameters and any number of outputs.

It is important not to think that NNs are applicable to all kinds of problems. For instance, if the data exhibit a linear behaviour, then it would be better to use just a traditional linear method such as linear regression. This method will make a better prediction and it will take less computational time. In contrast, for many complex processes such as stock prices, a good prediction of the near future is obtained by using an appropriately weighted combination of recent measurements of the variable being predicted and other correlated variables. The most widely used approach is the Auto Regressive Moving Average (ARMA) model. When the process is nonlinear and sufficient data are available, NNs provide a more accurate model than the linear ARMA model (see [Box & Jenkins 76] for extensive descriptions of conventional ARMA models and the Box-Jenkins modelling approach).

In addition, NNs can be used in conjunction with conventional forecasting methods. For example, one can often produce more accurate forecasts by providing several conventional forecasts as input to the NN. In this case, the networks serves partly as a combining method producing a weighted average of the different forecasts.

4.4.1 Time Delay Method

Consider we are trying to predict the behaviour of a variable that changes over time. We must then start with a sequence of data taken in several intervals, creating a discrete time series to analyse. For example, if the variable of interest are the prices of a stock during the month of November with daily measurements, we would have a series of 30 data points: P1, P2, P3, ... , P30. The Time Delay Neural Network (TDNN) consists of some inputs that get segments of the series, the segments are of fixed length.

Let's consider a TDNN with 3 inputs, 2 nodes in the hidden layer, and 1 output to predict prices after the 30th of November. When training the net, we would present first P1, P2 and P3 as inputs, and P4 as output, indicating that P4 is the desired output for this segment of the series. The desired output is used to adjust the weights of a Back-Propagation network or simply to verify the performance of the net. The next training segment of the net consists of P2, P3 and P4 as inputs, and P5 as output. This approach continues until we get to P27, P28, P29 as inputs and P30 as output. Each segment of the series is called a *window*, because each segment represents a view of fixed size as if it was a window moving along the series.

After training with the 27 time windows, the NN must have an internal representation of the function. To know P31, the estimated price for the 1st of December, we must now present the window P28, P29, P30 to get P31 as output. This is the prediction one step ahead. The prediction of the next step is obtained with P29, P30, P31 to obtain P32, etc, etc.

As it can be seen from the example, predictions of more than one period ahead could become unreliable when the predictions from previous periods are used as inputs.

If the first prediction was wrong, when it is used for further predictions, the error might increase (as noise is added to noise), and it could accumulate exponentially. The fact that the error is accumulated over time is common to many prediction methods of complex systems. For example, weather predictions can sometimes be reliable for up to 2-3 days, but when trying to predict a longer period such as one week in advance, the error might become unacceptable, and this is problematic specially where a high degree of accuracy is expected. In the case of using correct predictions as inputs, the situation does not necessarily get worse due to error cancellation.

Let's recall two important considerations about neural networks:

- The NN's feature for which it exhibits the ability to infer a function that approximates to the data points, is known as *generalisation*. Even NNs with relatively simple architectures are able to present non-linear complex behaviour.

- When trained excessively, they tend to *memorise* the data, in which case the learned internal representation of the function is lost, and so the ability to generalise. The probability of getting correct predictions is reduced.

The goal then is to build a NN that is able to generalise without memorising the entire series. Memorising can be achieved through long periods of training. The problem is, when the net is exposed to unseen data, the prediction is not going to be accurate. In other words, ideally, we want a net that can learn the underlying mechanisms of the series, not an encyclopedia that will give a perfect output to a given input that has fully memorised. Rather than memorising (which would be cheaper to implement as a look-up table), we are interested in *understanding* what is the most likely thing to happen under the given circumstances. A common way to achieve both is to separate the data (in this case the moving windows) into 2 sets. The first one, the *training set*, is used to train the network and the second one, called the the *testing set*, helps to test the overall performance of it over the data that was never used to train. For example, in a series of 1,000 data points, one can use the first 950 points for training and the remaining 50 for testing. The point of greater generalisation is found at the minimum error over the testing set. This is also the point in which the training period must end. If during training we continuously measure the error of both sets, we would find ideally that the error of the training set decreases asymptotically to zero. On the other hand, the error of the testing set starts to decrease and then starts to increase again.

The architecture and training algorithm selection involves another important decision when using NNs. Different combinations are best suited for different problems. The variety of prediction problems that has been solved is very wide, it includes predictions of sun spots, pollen, electric charges, and stock prices, among others. Researchers in the past have used a number of different methods, ranging from the Boltzmann Machine to Unsupervised Neural Networks. However, the vast majority comprehends Feed-Forward Neural Networks used with Back-Propagation learning algorithm. Most commercial packages include this combination.

4.4.2 Problem Description

The task is to use a NN method such as the Time Delay method just described to predict the next value of a given variable, typically the price. The problem has just been defined, it is a simple one. The next stage is to investigate more about such variable; for example, how does it behave over time? A good suggestion is to start by looking at real markets. Let's consider what they have to offer.

Each day, when the market opens, current information about all stocks is available through a variety of sources such as Telerate, Infosel, or Reuters. This information is constantly refreshed. At the end of the day, the most representative values for each stock that are recorded are the following:

- Opening Price
- Closing Price
- Maximum Price
- Minimum Price
- Volume
- Global Market Indicators

The following represents a typical format in which most data is freely obtained through the Internet. This data was taken from `http://quote.yahoo.com` and it corresponds to the Coca Cola stock during April 2001. As it can be seen, it starts with the Date. Date formats vary between sources. The following four columns describe the various prices that are recorded daily, starting from the opening price of the day, then the highest and lowest of the day, followed by the closing price. Finally, the last column displays the volume of transactions relative to that stock. Note that dates are given in descending order, they should be re-ordered in ascending order. This is an easy process which can be implemented in the NN system as part of the preprocessing phase, or it can be done separately by using almost any spreadsheet package.

```
Date,Open,High,Low,Close,Volume
30-Apr-01,46.75,46.75,45.85,46.19,3837000
27-Apr-01,47.50,47.50,46,47,3480900
26-Apr-01,47.60,47.98,46.90,46.90,4296400
25-Apr-01,47.50,48.40,47.40,48.20,3247700
24-Apr-01,47.45,48.45,47.17,47.39,4343400
23-Apr-01,47.50,47.90,47.03,47.45,3330500
20-Apr-01,47.05,47.25,46.28,47,4583600
19-Apr-01,46.05,47.77,46.04,47.50,5418200
18-Apr-01,46.30,47.08,45.71,46.75,7383700
17-Apr-01,45,45.70,44.81,45.70,4407900
16-Apr-01,44.57,45.50,44.15,45.41,4143500
12-Apr-01,43.50,44.90,42.65,44.57,5214600
11-Apr-01,43,44.75,42.37,44.13,6874200
10-Apr-01,43.91,44.30,43.21,43.65,5948200
9-Apr-01,45,45.59,43.90,43.90,4115800
6-Apr-01,44.43,45.35,43.85,45,3953900
5-Apr-01,46.25,46.60,44.62,45.30,4621100
4-Apr-01,44.52,46.05,44.17,45.25,4399300
3-Apr-01,45.60,45.70,44.18,44.66,4672700
2-Apr-01,45.40,46.92,44.86,45.85,5212400
```

The problem is essentially to analyse such historical data trying to predict the value of the variable for the next day. In other words, the goal is to predict today the price of tomorrow. Figure 4.1 shows a typical time series, in this case the series represents stock prices of the Coca Cola Company from January 30, 1996 to January 14, 2000 (Closing Price in US Dollars vs Time).

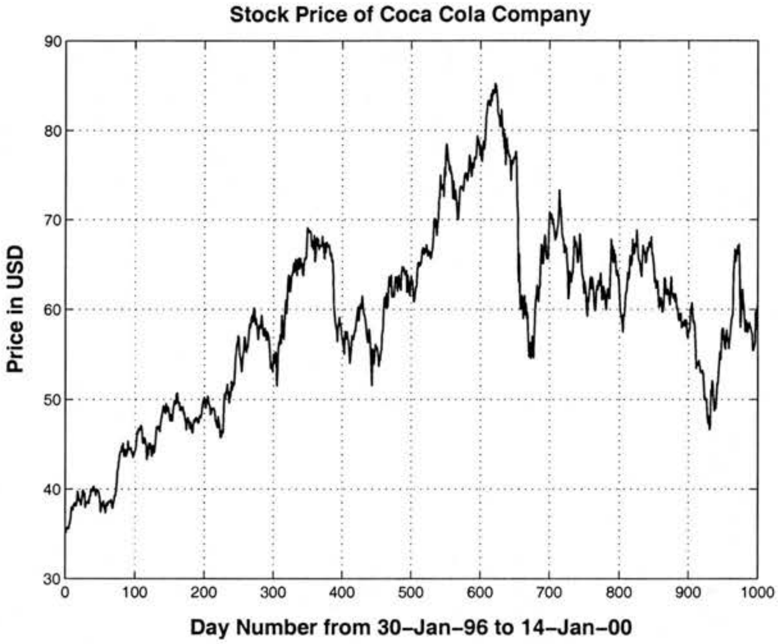

Figure 4.1: Example of a Time Series, Coca Cola Company stock.

4.4.3 Proposed System and Technologies

Having described the problem, the question to analyse is the following: Given historical data, is it feasible to use NNs to come with reasonably good predictions? The goal then is to design a system based on NNs to try to discover the behaviour of the shares. Price is only one of the many variables that can be considered. A typical system of this type is illustrated in Figure 4.2, and will be described in more detail in the following sections. In addition, Chapter 6 describes a number of real applications of systems such as this one.

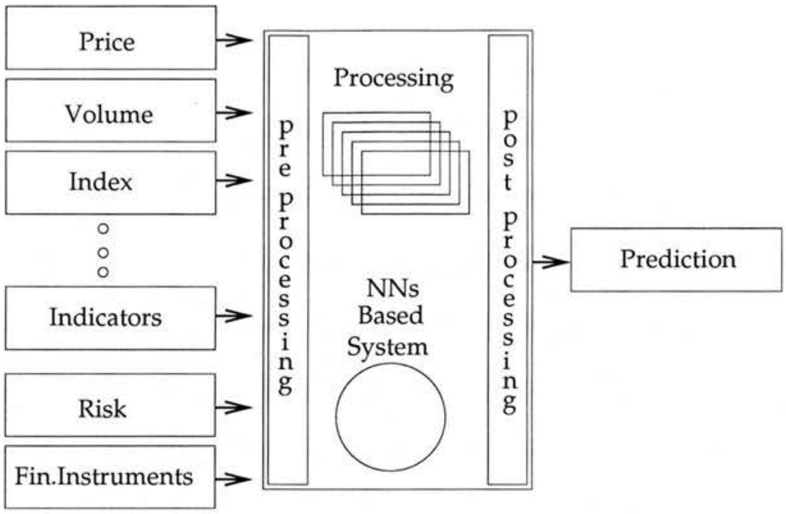

Figure 4.2: Diagram of a Typical Neural Network System. Historical data is processed, fed into the NN and the output is then processed again to obtain the prediction value.

4.4.4 Preprocessing Phase

This phase involves the data collection and analysis. As previously described, daily data can easily and freely be obtained from sources such as Yahoo (`http://quote.yahoo.com`), and high frequency data for research purposes can be obtained from Olsen Group (`http://www.olsen.ch/research/using\_our\_data.html`). The basic data consists of five daily values for each stock: opening, maximum, minimum, and closing prices, as well as the volume. Also, daily values of other global market indicators such as the New York Stock Exchange (NYSE), interest rates, foreign exchanges, etc can be obtained.

4.4.4.1 Data Transformation and Normalisation

When the system receives the original time series, it will first do a transformation of the data before they are used by the NN. The transformation could be defined by the user. Such transformations are functions of $R_n \to R$. For example, the percentual difference from one day to the next would be a $R_2 \to R$, a 10 day moving average would be

$R_{10} \rightarrow R$, etc. The idea behind the preprocessing phase is to get information about the series that can easier be used by the net.

Data must then, be normalised. In other words, the values must be altered so they lie within an adequate input range for the net, i.e. from $0- > 1$.

4.4.4.2 Preparing Test and Train Data Sets

The last step of the pre-processing phase is to separate the series into two different sets in order to avoid over-fitting: training and testing sets.

4.4.5 Processing Phase

The second phase deals with the design and construction of the Processing module, which integrates the following:

1. Back-Propagation and evolutionary training algorithms for weight adjustment.

2. A modular net topology composed of feed-forward or recurrent nets in a hierarchy. Once the data is ready, the modular NN must be generated according to the pre-determined parameters. Each sub-net must be of one of the types explained earlier. The NNs are made with 3 or more layers of neurons. The first layer is only connected to the second, and the second to the third, and so on until the output layer is reached. Neurons are fully connected, so every node of layer N-1 is connected to every node of layer N. The output layer is always of size 1. For recurrent evolutionary networks, the output layer is connected to the input layer as a window containing the last M results from the output. The organisation of the sub-nets of the modular NN must be hierarchical. The first level will perform the prediction of the series. The second tries to predict the prediction error of the first. The third estimates the error of the second, and so on.

3. The time series is presented to the net in the form of a time window of a pre-determined size W. In a series of data $X_1, X_2, ..., X_n$, suppose $W < n$, the input data is presented in this order (each set is known as a time window): $(X_1, X_2, ..., X_w)$ $(X_2, X_3, ..., X_{w+1})$ ... $(X_{n-w+1}, ..., X_n)$

4.4.5.1 Back-Propagation Network

It consists of 2 phases: (i) *Forward Activation*, where data is read (one window $X_1, X_2, ..., X_n$) and the output values of each subsequent layer are calculated until the output layer is reached, and (ii), *Backward Propagation*, in which the calculated error of the output is propagated backwards altering the weights of the output layer first, then the intermediate layers, and finally, the input layer. The weights of the hidden layers are initialised randomly according to the limits defined by the user.

4.4.5.2 Population of Evolutionary Networks

This is a simple evolutionary algorithm to find the necessary weights of the NN. First, a population on NNs is created. Then, they are evaluated according to a fitness function defined. The error against the original series is calculated. They are sorted from best to worst and the worse part is substituted by new ones similar to the good ones. Then the process is repeated for so many generations. After N generations of evaluate-sort-reproduce, there must be a set of weights with a minimum error.

4.4.6 Post-processing Phase

In this phase the output from the net is de-normalised again. Basically, the inverse process of the preprocessing phase to be able to get a real number meaning the predicted price of the following day. Note that in this example the output will not be just the direction of the price (up or down), but the magnitude (e.g. $1258.97 USD).

A great number of NN systems with varying success have been implemented in industry and academia. Some results of models such as this will be discussed in the applications section of Chapter 6.

4.4.7 Some important considerations

Financial data possess unique problems in constructing predictive models. A large part of the effort in using a neural net for predicting market movements is determining which indicators to use and how their values are to be preprocessed. For instance, about the raw indicators one must consider whether to include some of the following:

- closing price, the difference between bid and ask, high and low prices, real-time price, etc.
- additional related indexes such as DJIA, S&P 500, Nasdaq, etc.
- volume of transactions.
- macro-economic variables such as inflation rate (INF), price earnings ratio (P/E), interest rate (r), money growth rate (m).
- fundamental information such as the company's turnover, analysis of risk and yield of instruments used by the corporation, etc.

Then special attention must be given when choosing how these values are to be processed. This includes:

- choosing whether to use changes of prices rather than their absolute values,
- whether to telescope the data to look at certain historical values,
- whether to use sensitivity measures,
- volatility measures,
- moving averages, change in moving averages from one period to another and determining how far back in time the indicators should be examined.

In addition, the search space of parameters for training the nets (for example, learning rate, number of hidden nodes, momentum, eta, step size, etc.) is usually very large. A promising approach is the use of a GA to do the search instead of other approaches such as the one of multiple nets (trained using different parameters with the same data), where the output of the nets is combined. Even a simple average of the outputs of the individual nets appears to work well in practice.

The number of input values could also make the search space very large. For this reason some research has been devoted to the use of a GA for input selection. (See, for example, [Pictet *et al.* 96] for a discussion of designing trading models with relevant indicators optimised by a GA.)

Chapter 5

An Introduction to Financial Issues

> "Cause and effect: such a duality probably never exists; in truth we are confronted by a continuum out of which we isolate a couple of pieces, just as we perceive motion only as isolated points and then infer it without ever actually seeing it. The suddenness with which many effects stand out misleads us; actually, it is sudden only for us. In this moment of suddenness there are an infinite number of processes which elude us. An intellect that could see cause and effect as a continuum and a flux and not, as we do, in terms of an arbitrary division and dismemberment, would repudiate the concept of cause and effect and deny all conditionality." Friedrich Nietzsche, *The Gay Science*, s.112, translated by Walter Kaufmann.

This chapter begins with a brief explanation of some important factors about markets in general. Namely, the section addressing basic concepts about markets is followed by addressing some of the many factors that are believed to originate price changes and trader's perceptions. Following this financial background, more detailed descriptions of traditional and non-traditional approaches in financial market analysis will be addressed.

5.1 Basic Concepts About Markets

It is well known that real markets are complex structures that follow coordination mechanisms where an element of "exchange" must take place in such a way that a specific product's aggregate supply must equal its aggregate demand. This single fact complicates things even more when trying to model markets. As a consequence, there are a number of possible trader behaviours observed, which are characterised by different rules and organisations. For instance, some markets allow traders to buy and sell simultaneously, while in others, traders must have a clear side, creating important changes in the decision-making dynamics [Tordjman 98].

In a large number of markets, "prices are set up by some exogenous instances and maintained more or less constant until a big change in supply and demand conditions occurs which necessitates a change in the terms of trade"[Tordjman 98]. An example of this type is in the area of agriculture or certain fisheries (e.g. tuna), where the prices of products are controlled by the European Commission or some other external organisation.

Other type is a *price-making* market, where prices are determined through supply and demand relationships through buyers and sellers rather than controlled by external institutions. Common examples of this type of market are the fish, asset and commodity markets. Regarding the former one, there is extensive research done in France by Alan Kirman and Gérard Weisbuch with actual data gathered at the Marseilles fish market. Unfortunately it is outside the scope of this thesis to describe their work, however, for the interested reader, some pointers on their market model (an adaptive economic model) and further comparisons of their model with empirical data of the wholesale fish market in Marseilles can be found in [Weisbuch *et al.* 98, Kirman & Vignes 91].

Asset markets include the FOReign EXchange (FOREX) market of the London and Paris Stock Exchanges, among others. Both of these markets are described to be *price-driven*: there is no institution collecting orders, traders must find their counterpart, as opposed to an *order-driven* market such as the Paris Stock Exchange, where orders are collected by a centralised institution which calculates equilibrium prices for every stock (for details on how a typical transaction of these two types proceeds, refer to [Tordjman 98]).

FOREX markets are amongst the most efficient markets found so far. Immediately after a transaction is made, it is recorded and the new price is updated instantly by centralised communication systems such as Reuters or Telerate, making volume of transactions as well as prices publicly available instantly.

Finally, commodity markets are very similar to asset markets, except that the products traded are a variety of raw materials. The Chicago Mercantile Exchange (CME) and the London Metal Exchange (LME) are examples of this type of market, which now offer very sophisticated products such as options and futures.

5.2 Traders

A financial market can be viewed as a collection of recorded activities affecting it, creating and recreating it in different ways. It is a process directly affected by the actions of its most relevant component: *humans*, the real players involved in the very process of trading; humans making a broad range of decisions, buying and selling different securities with different purposes and time horizons. Most of them, at the end of the day, intend to profit from such transactions by using one of the most ancient ways of making money, "buy cheaply and sell dearly." However, the implications of this are not that simple. If everyone knew the price of a security is going to go up in the following day, everyone would buy today and sell tomorrow.

It is largely believed that as traders are diverse, there is some unexpected predictability in markets which might be explained with the heterogeneity of traders. People reason differently about the information they receive, they have different time horizons (short, medium, long term) and so are their attitudes to risk. This is why there is so much unpredictability in the markets. Traders who exploit some market inefficiency thereby cause it to disappear. It has therefore traditionally been very difficult to develop models which consistently outperform simple ones such as *buy-and-hold*.

By definition, not every trader can win, in any given transaction, the chances of winning are about the same as those for losing. But evidence suggests that some people win more often than they lose. This can be due to many factors apart form the obvious one which is luck, more wins can be the result, among others, of having chosen better portfolios, performed better market timing transactions and better decisions involving certain negotiations.

In addition to this, it is important to mention the difference between information and what people do with it. When receiving the same information at the same instant, two people can come to opposite conclusions about its likely impact on prices. Neither investor is irrational, although one might be proven wrong. Therefore, a market may well misprice assets at some point, but the mispricing is not necessarily caused by inefficiencies in distributing information but also because of different ways of reacting to it.

In reality, the fact is that there are some good traders and some bad ones. But who is a good trader? How can we measure success in traders? In financial markets a good trader is not necessarily the person who makes the "right decision" most of the times. It could be someone who only makes three good decisions in a year, probably obtaining larger profits than another one who "got it right" most of the times, but with a smaller profit average. So as long as they keep earning a profit, they might stay in the market. But what about the others? Well, many others who report losses are also in the market. There are other reasons to stay in the market, as explained in Chapter 2 such as a government trying to support the value of its currency. In a way, they are all survivors – market participants.

Thus the market is composed of all sorts of traders (experienced and unexperienced, mathematical, intuitive and gamblers, etc.) competing and hoping to get the most out of their investment at the expense of whomever is on their way. Some do not only rely on their own tactics, but instead they might copy what others, that they consider good ones, are doing. They imitate and possibly learn something by trial and error experimentation. In this pool of traders, some just leave the system whenever they want, and others are killed (traders are kicked out of the market usually by punishment, negative criticisms, etc). They can have a very short lifespan or a long one.

5.3 Price Moves

First, a broad and general introduction of some important factors that play a part in the dynamics of financial markets will be given. This includes a description of the factors that contribute to the changes in prices, followed by the the decision making process itself as defined by experts – is there such a single process? or is it that some people follow only some rules and are governed mainly by their instincts and own beliefs? –

and finally, a description of the different personalities involved in the process.

Although in the financial world it has been difficult to isolate a single factor that originates moves in the markets, it is in the interest of this work to find out how much the psychology of traders — people's expectations and reactions – affect the process. We see that prices go up and down, dealers panic (some first, others later). Every move means someone loses money and someone else is making it. What determines these moves? These reactions?

The main factors in determining the level of prices at any point in time are: (a) supply and demand, (b) external influences such as political or geographic, i.e. an earthquake, (c) economic statistics and (d) technical considerations. The latter one will be explained in section 5.4.1 and the former ones as follows, taken from [ADT01g]:

5.3.1 Supply and Demand

- If there are more buyers than sellers, prices will rise. Conversely, if there are more sellers than buyers, prices will fall. As buyers/sellers are satisfied and become less and less inclined to buy, the rate of increase/decrease in prices will slow. In the absence of other factors, prices will become static at the point when buyers and sellers are in balance, or there is no particular interest in either direction.

- Sentiment plays an important role in supply and demand. Markets are too large for any one authority, or even a group of authorities, to be able to do more than act out a reasonable impression of King Canute. This is not surprising when you consider the fact that $5 billion of transactions would represent a serious attempt at intervention by the central banks - but would represent only around 1% of the average daily FX volumes in London alone. Official intervention is only really effective when the market is already heading in the desired direction, or at least has significant doubts about its current level.

- Financial markets are closely inter-linked. For example, movements in bond markets often impact on equities and vice-versa.

5.3.2 Economic Statistics

Again, expectations are at least as important as reality, but "the figures" (the release of economic statistics) are eagerly awaited events, with all dealers and brokers at their desks, avidly watching their screens as they await the announcement of important numbers. There is often a decided "fashion" as to which statistics are regarded as the most significant at any point in time.

One of the most important numbers is the *money supply*. The greater it is, the easier it is to obtain money, thus affecting the level of prices of goods and services, the wages and salaries paid, and finally, causing a higher level of inflation (the "monetarist" philosophy). The natural extension to that theory is that inflation may be contained by controlling the supply of money and it is for this reason that the target rates of growth and, in due course, the actual rates achieved are announced [ADT01g].

Other important figures are the Government Borrowing Requirements (GBR) and the Consumer/Retail Price Index (CPI or RPI), which are common indicators of government spending and inflation. Trade figures are one component of the overall picture of a country's economy in terms of its trade and other financial transactions with the rest of the world.

5.4 Traditional Methods of Forecasting

Traditional approaches for modelling financial markets such as rational expectation models usually assume that agents have complete knowledge of the structure of the market as well as other agents' beliefs and objectives [Routledge 94]. Such theories state that a market remains in equilibrium if the agents display a rational behaviour in anticipating changes in supply and demand. Simply by looking at stock market data, though, one can observe that there is a type of volatility that does not match the rational expectations model. In addition to this, it is also assumed that traders react to the same information in a linear fashion and that in the aggregate all traders are rational (even if they do not behave rationally individually). In fact, the main assumption is that trader skills and experience are equally distributed.

Considering real financial markets, traditional methods have not been successful in explaining how the decision making process works in any type of agents, whether they are real or artificial. Many assumptions already mentioned are considered by these models, such as that there is common knowledge in all the agents and that there is a unique solution to the problem. Also, the strategies developed are usually derived from common procedures; this one being the basis of the rational expectations hypothesis.

Other assumptions have to do with the information contained in current stock prices. EMH argues that there is no way to predict the future price by looking at past data. It is a random process, therefore it is impossible to beat the market. Under this theory, because markets are perfectly random it is impossible to make abnormal profits and people are perfectly rational.

In the context of this thesis, economic methods based on classic economic theories such as the EMH and the REH are referred as *traditional methods*. However, financial analysis, which will be explained in the next section, is also considered a traditional approach.

Forecasting is widely known as being the weakest part of economics as a science. Investment is the most influential and capricious fuel to any economy, and individual investors size their trades reacting on what they expect to sell in the future. By means of a complicated social information structure, consumers and investors talk each other into some level of expectations. The way they do this is difficult, not to say impossible to model.

If, for instance, it becomes expected that interest rates will increase, potential lenders wait, and potential borrowers rush. As a result, the interest rises, may be even faster than the original expectation (self-fulfilling forecast). If it becomes expected that next year, oil will be scarce and hence oil investments will yield substantially above average returns, oil investments now may well cause that this will not happen (self-destroying forecast). Nevertheless, people do engage in the risky business of forecasting.

Two schools of stock analysis dominate: technical and fundamental, however, a tentative school which might lie in the middle of these two, known as quantitative analysis, will also be addressed in the following sections.

5.4.1 Technical Analysis

Some traders believe in the techniques to a great extent, while others are skeptical. What is clear, however, is that these methods are widely used by practitioners and academics. Forecasting is based on the rationale that history will repeat itself and that the correlation between price and volume reveals market behaviour. Prediction is made by exploiting implications hidden in past trading activities, by analysing patterns and trends shown in volume charts. This micro-level scrunity does not consider any external factors such as news about wars in the Middle East [Murphy 86].

Note that this type of analysis pays no attention either on to how the company is being run, the decisions are based on what the market already knows about the company as indicated by figures such as its share price and volume. The investment decision is made purely on the basis of future expectations of stock price within the market. These expectations, in turn, are believed by chartists to be a function of the story of price fluctuations alone [Essinger 90].

Some people believe that the principle behind technical analysis simply relates to the psychology of markets: all traders must be in some form of crowd – bullish, bearish or neutral. At the beginning of an uptrend, the bullish crowd will tend to be very small. As prices begin to rise, however, and the market becomes more bullish regarding future prospects, the size of the bullish group will grow at the expense of the bearish and neutral groups. When the market reaches a peak, this is quite simply because the bullish group has reached its zenith and there are no other willing buyers to be found [ADT01d]. Once the market has run out of buyers, the only direction for it to go is down. Whereupon the cycle begins again, only this time in reverse with sellers flooding into the bearish group, until it too reaches the limit of its cycle and a further reversal of prices is inevitable.

Technical analysis can be particularly beneficial for smaller traders, or dealers in small and medium-sized institutions, as all the information required is contained in the market price, negating the need to constantly trawl through volumes of conventional economic data in a quest for enlightment [ADT01d].

5.4.2 Fundamental Analysis

Basically, fundamental analysts tend to make long-term investments in stocks of companies based on facts such as how these companies have been run, the underlying value of their assets, their market segment growth and facts such as whether the CEO is competent or not. For this type of analysis, there is no need for great amounts of data, computers or fancy systems to select their stocks.

Fundamental analysis relates more specifically to the particular corporation or industry, and adds to the analysis two levels of human judgement: first, the opinions and forecasts of company management, and then, the investment manager's own subjective interpretations of those opinions and forecasts. Forecasting is based on macroeconomic data, such as exports and imports, money supply, interest rates, foreign exchange rates, inflationary rates, and unemployment figures, and the basic financial status of companies [Murphy 86].

Similar to quantitative analysis, fundamental analysis also involves the study of the actual corporations in which the investment manager is interested. The difference is that quantitative analysis is based on actual data only and tends to use large samples of it.

5.4.3 Financial Modelling

This broad category can also be referred as quantitative analysis or numerical analysis systems. There is no crisp distinction of this category. However, in general, it can be said that this type of analysis does not try to imitate any aspects of the human subjective judgement-forming process, but rather to apply a specific mathematical theory to current market data, and draw inferences regarding the likely future development of the market or financial instruments under scrunity.

It differs from technical analysis in that, rather than simply making assessments and forecasts of abstract price movements within a particular market, *quantitative analysis* focuses on basic data about particular corporations, and then analyses the risk and yield of investment instruments issued by those corporations, often going to apply the analysis to specific theories for how the risk and yield of that instrument or asset class are likely to vary in the future.

One particular form of *numerical analysis* that gained popularity during the latter part of the 1980s is known as Modern Portfolio Theory (MPT). This involves the system assessing variables such as historical risk and yield of particular stocks and producing optimisation forecasts (i.e. scenarios where the expected return is maximised for a given level of risk, and where the risk is minimised for a given level of return). Other forms of numerical analysis can be used in conjunction with any of the two classical types of stock analysis, technical and fundamental. The two main types of risk and return usually identified are *beta-risk* (risk existing by virtue of the market or industrial sector in which the corporation operates) and *specific risk* (risk deriving from the individual corporation).

5.4.4 Other Methods and Market Considerations

Some traders usually take a look at relevant pieces of information described in the form of charts. The most simple and popular ones are the bar chart, the tick chart and the point & figure chart. There are others involving more complex computer algorithms such as the candlestick chart, which basically searches for candlestick patterns in data such as an index and measures the frequency of occurrence (refer to [Wagner & Matheny 97] and [ADT01d] for additional descriptions of these types of charts). In some of these charts, for example, traders might analyse areas where demand for particular instruments has previously been apparent. Such *support* is generally visible where there is a cluster of "lows" which has not been substantially penetrated by sellers. Similarly, areas where sellers have been evident, as represented by clusters of "highs" on a chart represent *resistance*.

In an effort to discover support and resistance, there are trend lines drawn on charts. When a market does break into a trend, it may start to open up gaps in the chart, which can be an important signal. Recognition of patterns in charts is central to technical analysts. Some of the most important formations are the head and shoulders silhouette, flags, pennants and triangles.

In the late stages of a trend, markets often see a rapid acceleration in that trend - essentially caused by panic covering of contra-trend positions and because others are belatedly trying to get in the trend. This can cause what is known as *climax*. After accelerating, the market suddenly reverses with equal vigour. Climax formations can be very dangerous to stubborn traders who, having entered a market late in a trend, hang on only to find a rapidly accrued profit being decimated and a substantial loss occurring within a very short period of time [ADT01e].

It is important to mention that markets only trend at the absolute most 40% of the time and possibly as little as 20% of the time. Regardless of the precise numbers, it is fair to say that trends are the exception rather than the rule [ADT01e]. When markets are not trending, they are described as bracketed markets. In such markets the volume may tend to dwindle and traders become apathetic about the prospects. But short-term traders often buy or sell at levels whereby they will be triggered once the market moves outside the bracket, but will generally keep tight stops on positions in case the market cannot sustain the break out.

There is a *bull market scenario* when the market moves below the most recent minor market low and breaks down. Sometimes it does so as a knee-jerk reaction to an item of news or data release. Therefore the market can be prone at this stage to "gap down" below the recent minor high. On the other hand, there is a *bear market scenario* when the market rises above a previous short term minor bear market high, making it prone to gap upward [ADT01e].

The *long term trader* usually places a stop above the all time high or below the all time low, depending upon whether it is a bull or bear market stop. This type of traders allows for some element of retracement, just in case the market breaks down, retraces, stagnates and then pushes lower. The *short time trader* wants to avoid – to as great degree as possible – being in the market during stagnant periods. They usually place the stop in the middle of the area of congestion at extreme [ADT01e].

Another important factor is *volume*, when used in conjunction with other data,it is a useful determinant in identifying whether a continuation of or a reversal in the prevailing trend is likely to happen. When volume is low, it reflects uncertainty regarding the future direction of the market in question. If the volume is relatively high while the market is going up and remains relatively low during corrections, the inference is that the market is in a strong uptrend which should continue. In the other case, when the volume is high while the market is going down and relatively light during upward retracements, then the market is weak with a continuing downward trend likely [ADT01e].

Risk management is an important responsibility for financial institutions as well. It usually involves many different areas of the company, significant resources, and it is very complex and time consuming. It is vital, however, that it is carried out as effectively as possible; otherwise losses will inevitably be incurred – possibly to the extent that the institution will not be able to meet its obligations and thus be forced to close [ADT01b].

Although there are many types of risk (i.e. position, credit, liquidity, systemic and translation risks), only the *equity risk* will be mentioned here, which is part of a position risk, along with the exchange rate, interest rate, commodity and basis risk. The price of a share or a group of shares may move in a manner which incurs losses. For investors, the risk is normally that of a decrease in share prices reducing the value of a holding, or their overall portfolio. In the case of market makers who have "sold short", i.e. sold shares which they do not own in the expectation that they will be able to buy them at a lower price and thus secure a profit, however, the risk is that prices will rise before they have covered their position [ADT01b].

As it was mentioned earlier, financial institutions must, of course, place *limits* on their activities to ensure that they do not incur undue levels of risk. They would be aware of the level of profits they want to achieve, possibly in terms of both the actual amount and the return on capital [ADT01c]. So, one of the roles of dealers is precisely to allocate limits between different areas, setting appropriate profit targets. Any dealer who cannot make a profit within this limits, will almost certainly incur larger losses if he exceeds them.

5.5 Trader's Decision Making Process

The purpose of this section is to address some opinions about a trader's decision-making process. Although these ideas appear to be of great empirical and subjective nature, the results of these types of behaviours are also some of the causes that originate price moves in financial markets. Therefore they will be mentioned in this section.

Most trader's ultimate goal is to maximise profits. In doing so, according to the experts, an effective route to follow is forming a view first, and then trading according to this view of the market.

There are many important issues to consider while *forming a view*. The starting point is to decide whether the trend of the market "is going up" or "is going down". In this stage they examine factors such as (a) how well related markets are doing (i.e. forex, bond, stock, commodity)?, (b) economic arguments and statistics, (c) political situations (their impact should not be underestimated), (d) geographic considerations such as the relative location of the market concerned and the forces of nature; (e) the mood of the market, (f) technical indicators (as many traders strongly rely on them, they also have a great impact on markets), and finally, (g) information sources such as newspapers, periodicals, screen services, charts, etc. [ADT01a].

After deciding whether the market is rising or falling, traders then need to address a whole variety of considerations before putting on a position. These are their considerations, taken from [ADT01a]:

- The likely size of the move, when will it take place and how will it occur?

  - Small profits made fairly quickly a number of times (jobbing) or a large profit made from a significant position over a longer period?

  - What disciplines should be imposed as regards time horizons?

- Will the market stabilise or merely pause for breath?

  - If there is a rapid reverse of the market precise timing is vital.

- Considering the possibility of "doing nothing" if not really convinced of the view is a good alternative as well. There is no point whatsoever in trading just for the sake of it. Some of the most able traders make up to 90% of their profits on a limited number of occasions – perhaps as few as 5 days each year. So it is not as important to be right 70% of the time if an overall loss is incurred. Instead, it is far better to be right only 20% of the time and still make a profit.

- Risk aversion. Proper consideration to the maximum potential profit and possible losses ratio must be given.
- Limits. Will the trade being considered change the market? When could the trade be cut out if the market is moving against?

The final step is *trading that view*, i.e., determining how to get the best profit from the picture of the market. Some traders concentrate on one instrument, while others do not limit their investment and diversify trading futures, options, commodities, etc. At this stage they consider factors such as the currency availability, the period of the trade, balance sheets of banks or financial institutions, capital backing, credit risk, the size and time horizon of the move and the liquidity of the market. Regarding their degree of conviction and risk, they must now decide whether to put the whole position on in one trade, or gradually build it up over a period of time, considering that they can loose flexibility if they go beyond their own limits [ADT01f].

So far, most of the external factors traders consider when "making up their mind" have been explained. That includes the type of information they obtain – either publicly available or transformed by them or others in some way – and the process to follow. But the question as to why is it that people react differently when given the same information still remains unanswered. The underlying rule appears to be that traders use many different criteria in their decision concerning what to trade or whether to use a particular system or methodology, which seems to be part of a complex process. The following are three opinions about the subjectivity of the trading process as a whole:

First, some people strongly believe that "successful trading is essentially a simple activity. There are very few rules to follow and the process itself is not difficult to understand. The reason we find it so hard to succeed is that it is difficult for us to follow those simple rules. Doing so goes against the grain of our human nature which has had thousands (perhaps millions?) of years to evolve and cement its influence over us" [Jaeger 97].

Second, others believe that an interesting concept that plays an important role in the process involved is the so called "pendulum effect". In this discipline loop in which traders often go through, as explained by John Piper when talking about the book "Trading Chaos" by Bill Williams, "we acquire the discipline to follow our trader methodology, we start to make good money, we then get over-confident, we start to diverge from the discipline, we start to do badly, we are humbled, we re-learn the discipline once again, we start to do well, we again become over-confident, and so it goes on". He claims this pendulum effect is seen in the real world again and again, and that to succeed traders have to go beyond that state [Piper 97].

Third, studies made with the writings of past traders show that the same mistakes made fifty and one hundred years ago continue to be made every day. According to this view, "technology may change, but human nature never does" [Kleinman 97]. Some recommendations from experts, also found in [Kleinman 97] can be summarised as follows:

1. Determine the major trend and then to follow it. "The big money is made by going with the trend, not against it,"

2. "Stay in the trend until your indicators suggest the trend has changed, and not before,

3. Be aggressive on entry and cut losses quickly when the market is not acting right,"

4. "Act without hesitation, ... never regret your decision when you liquidate a trade based on sound reasoning, ..., money management is the key, you don't necessarily need a high win to loss ratio, but your average win must be higher that your average loss if you want to succeed," etc.

After reviewing a variety of views about the process of trading, a number of trading techniques and considerations will be addressed:

- Usually, the *scenario* is fundamental in content. Traders can get caught up in a scenario in which they believe the market will unfold. Becoming a neutral observer may be the best strategy. "*Scenario trading* is trading in accordance with a dominant idea or story. It can be very dangerous – or it can make you a fortune. The scenario acts as an integrating theme and it can dominate your thinking. That thinking could be fundamental, technical or personal in nature, and it colours your thought processes. And it definitely affects your trading." [Roosevelt 97]

- " ... *Experience* allows us to recognise a mistake when we make it again. Well, I want to minimise my mistakes, and so I use a notebook to help identify patterns before they become significant problems" [Bulkowsky 97]. The notes include the reasons for buying a stock and comments of the success or failure of a trade. They are divided into investment styles: trend channel trade, the 1-2-3 trend change, long-term holding, fundamental value, and so forth. The fundamentals are considered first, followed by the technicals, which basically are used to time the purchase and sale.

- " Keep in mind the risks of over-optimisation, however; finding a particular *set of parameters* that worked outstandingly in the past does not guarantee good future results" [Vakkur 97].

5.5.1 Issues Affecting the Decision Process

- The way traders *deal with change* is an important issue while making decisions. Usually they are very sensitive, and minor or subtle changes can result in the disruption of their performance equilibrium. While for some regular people changes such as moving into a different location, or things as simple as having to change parking lots does not directly affect their performance at work, they can produce a great variety of feelings invading a trader's focus and attention. These range from the feeling of inadequacy, depression, anger or resentment, to major discomfort, anxiety and fear. Other factors related to physical changes include loss of appetite, insomnia, heart palpitations and high blood pressure, as well as an inability to cope with cold, noise, or any other kind of stress [Toghraie 97a].

- When traders engage in negative self-talk (i.e. the inner voice that interprets what is happening and their feeling about it) they can fail to achieve their ultimate potential. Many traders can talk themselves into disastrous losing scenarios [Toghraie 96].

5.6 The Psychology of Traders

This is a difficult issue to address on its own, many types have been described in financial literature. A limited number of types of traders will be briefly described in this section.

As every person possesses a different level of tolerance for working within certain instruments, time frames, levels of volatility and risk, etc., the psychology of traders will be addressed by the following scenarios, according to: (a) personalities, (b) instrument vehicles, (c) logistics, (d) time frames, (e) source of money, (f) advice and finally, according to (g) industry. These types are taken from, and fully described in [Toghraie 97b], except where indicated.

5.6.1 According to Personality

- The Mathematical Trader. This type of trader believes in the rule of mathematics and views money as numerical figures, without attaching any emotion or judgement to it. He knows he will make money by following the signals. He can trade a trend following or oscillator type of system. His personality promotes precise, unemotional, linear thought.

- The Intuitive Trader. These traders use their intuitive signals to make money and are able to harness their emotional power to trade more effectively. They often use a mechanical system but add to it an intuitive indicator. His personality promotes illogical leaps of understanding, intuitive knowledge and emotional/sensory awareness.

- The Gambler Trader. This personality type needs to feel an instantaneous gratification of adrenaline and excitement. He is too excitable to maintain an emotional distance from trading, on the one hand, and too excitable to detect the subtle intuitive signals, on the other hand. He is able to tolerate systems that are loaded with high risk, high volatility, brief time frames, and high draw downs. Although he is capable of extraordinary moments of trading brilliance, he is usually unsuited to the profession of trading because of his lack of inner discipline and his need to create enormous wins rather than a steady steam of moderate gains. All traders can fit into this category for short periods of their trading life. No system is the right one for this trader because he is undisciplined.

5.6.2 According to Instrument Vehicle

- Futures Trader. He is willing to assume a higher level of risk than most other traders. High level of self confidence and they have more imagination for trusting in paper and ink value than equity traders do, because they are trading something that really does not exist until a point in the future.

- The Equity Trader. Instead of dealing with abstractions, he needs to deal with what he thinks is concrete reality. He visualises the value making research, going to the history of what they hold. Equities trade more on fundamentals than futures, which trade more on psychology and expectations [Young 97a].

5.6.3 The Upstairs/Downstairs Traders

- Upstairs traders are behind the computer screen working for their own, handling details and maintaining organisation and self discipline.

- Downstairs traders are closely connected to what is happening, to the rhythm and emotion of the crowd. They need instant gratification, but have to be disciplined enough not to get caught in a case of gambling fever. They have to have the patience to wait for a signal; and not to try to push action that is not there.

5.6.4 Time Frames

- Day Traders need to have a feeling of control over the day-to-day activities of the market. They need quick profits and lots of activities to sustain them, although their need for instantaneous gratification and activity is not as great as that of the average floor trader.

- Position Traders often have more confidence and money than any other traders have. They are willing to go through the ups and downs of the risk of long-term trading and usually have an opinion about the market. These are usually trend-following, moving average and oscillator traders. Long term trend is usually not computational intensive because the techniques themselves are not esoteric formulas. The good traders can explain the process in a single sheet of paper [Mullins 97].

5.6.5 Source of Money

- Their Money. These are traders who only feel comfortable trading their own money. They must think of their capital as inventory for their business and not relate it to personal funds.

- Other People's Money. This trader must disconnect himself from the stories behind the investor and only deal with the task of following risk adverse money management rules. These traders can invest the money of an individual, an institution or their employer's money.

5.6.6 Advice

- Taking their own advice by following a clear-cut, time tested system might be the best approach because they do not have to rely on anyone else.

- Giving advice to others. As a broker, an analyst, or through a newsletter that makes it possible for others to make money.

- Taking the advice of others works if you totally believe in the person you are taking the advice from.

5.6.7 Other Investment Types

- Stockbrokers are usually extroverted with people-orientated qualities.
- Money managers. Must posses the selling personalities of the stockbroker, but also some qualities of a businessman.

5.6.8 Another Classification

- Discretionary Traders: usually described as being emotional and subjective, they construct and use various mathematical indicators. Even though highly criticised by some, it is widely known that these traders follow many rules varying their key indicators from time to time, giving them control over a few markets [Young 97b].
- Systematic or Mechanical Traders: Are quite the opposite to discretionary traders. Objective and unemotional, whose trading concepts are governed with only a few rules and key indicators being always the same, this type of trader is able to control many markets. "As a systematic trader, my job is to make sure that my position using the system matches the hypothetical one, so my performance can match the hypothetical performance" [Hartle 97].

Although it is common for most traders to belong to one of these categories, there is too much controversy as to how to distinguish between the two. For example, for some, it is precisely the capacity to be objective and unemotional which marks out the great discretionary traders of all time. In the same sense, to say that a discretionary trader has many rules while the mechanical trader has a few, is not true, simply because in order to construct a system across a large series of markets cannot be achieved with only one or two parameters [Young 97b].

The genuine financial alchemists of the contemporary era are George Soros and various hedge fund managers. They tend to be widely subjective, noted for a certain degree of emotion, and certainly not averse to having very complex trading concepts with dozens of rules. Note also that these non-automated traders can quite easily think across (both analyse and trade) a series of different markets [Young 97b].

5.6.9 Working Environment and Behaviour of Traders

As pointed put by Eric C. Bettelheim during the 1996 Burgenstock Conference of the Swiss Commodities and Futures Association, (published in [Bettelheim 97]), the common place behaviour of traders is surrounded by physical and verbal aggression, low tolerance for frustration, and emphasis on oral gratification, shouting, swearing, excessive eating, drinking, drug-taking and childish behaviour, conspicuous display and showing off, practical jokes, gang-like relationships with adherence to a dominant leader and simplistic treatment of each other as sex objects. Admiration for toughness, recklessness and contempt for authority, characterise this market sub-culture.

It is also mentioned that most traders work either on an open plan trading floor, or on the floor of an exchange, allowing everyone to see and hear everyone else in the room. The author claims that this type of environment allows no privacy and an atmosphere of physical proximity characterised by unpredictable periods of intense stress and uncertainty. In addition, there is no chance for deliberation or reflection, nor for any form of thought aside from rapid arithmetical calculations and spontaneous, usually angry, outbursts.

Traders are terrified of failure and focused only on money as their reward. Greed and fear dominate the atmosphere. Punishment in financial markets is a common treatment given to those who "get caught." There is a constant effort trying to kick them out of the markets, to bankrupt them by issuing fines and legal fees, to stain their reputation and to generate bitterness (not remorse) and a sense of grievous injustice in many of their peers. As a necessary consequence of this narrow view of "success", and the short, life span of a trader even under normal circumstances, few if any of them evolve into responsible citizens of the financial world [Bettelheim 97]. Therefore many people think that traders should receive professional training instead of letting them loose in a psychological regressive environment of this kind; that way the system would not suffer some of these problems.

Chapter 6

AI Applications in Finance

"We have arranged for ourselves a world in which we can live - by positing bodies, lines, planes, causes and effects, motion and rest, form and content; without these articles of faith nobody could now endure life. But that does not prove them. Life is no argument. The conditions of life might include error". Friedrich Nietzsche, *The Gay Science*, translated by Walter Kaufmann.

6.1 Preliminaries

Because of the breadth of the financial industry's scope, the possibilities for AI applications are enormous, but for the purpose of *learning market and agent behaviour*, this thesis only concentrates on the problem of forecasting in the stock market, which alone involves a high degree of complexity, due primarily to the sudden changes it exhibits. This chapter is a follow up from Chapter 4 on Computational Economics, where Predictive Systems were described, except that here the focus is on the analysis of actual systems in practice, within various AI domains such as NN, GP, GA and LCS, concentrating on the results reported of such systems. Previous to these topics, Chapter 3 described the various AI techniques largely mentioned in the following sections.

Due to the interdisciplinary nature of the topic, this type of applications are difficult to organise in categories. However, following the order in which the various areas of this thesis have been divided, an attempt has been made to group them according to the AI technology that has been applied in such prediction problems. They are divided into the following: neural networks, learning classifier systems, genetic algorithms and genetic programming. But first, this section starts with a brief description of computational markets, which are relevant to this thesis from both, a computational and a learning point of view.

At the University of Michigan, Michael P. Wellman, John Q. Cheng, Peter R. Wurman, Tracy Muller, Tad Hogg and other researchers have made great contributions in the areas of computational market mechanisms for distributed decision making and the dynamics of multi-agent systems. In this view, a *computational market* is any collection of software agents interacting through a price system [Wellman & Wurman 97]. In other words, the market is treated as a multi-agent system in which every decision is really about *resource allocation*; resources are usually limited and of various types.

In this line of research, the term "Market-Oriented Programming" has been widely used by Michael P. Wellman, following Yoav Shoham's "agent-oriented programming" name and concept of layering constraints on top of object-oriented methodology. Wellman's market-oriented programming environment provides a set of generic constructs for specifying the elements of a computational economy, and implement some of the facilities to manage the interaction of these elements, according to defined protocols [Wellman 96].

In addition, an automated trading system called "Michigan Internet AuctionBot" [Wurman *et al.* 98] has also been developed, which is a configurable auction server with a WWW front-end as well as a software API, so it supports human as well as artificial trading agents (more information can be found at `http://auction.eecs.umich.edu`).

There are a number of systems described in the literature that have been put into practice in financial markets. Due to secrecy and certain success (if ever found), these descriptions tend to hide a great deal of detail. Nonetheless, it is important to devote a chapter to mention the scope of real applications of AI in Finance. This chapter also describes examples of systems that are not being developed commercially, but are relevant to the work presented in this thesis.

6.2 Neural Networks

Financial Institutions are replacing conventional systems by neural networks for pattern recognition, classification and curve fitting in various domains such as classification of credit applications, credit scoring, fraud prediction, for modelling and forecasting bankruptcy, securities trading, and portfolio evaluation. The following are a few examples of such applications:

6.2.1 Mellon Equity Associates and TSB

Mellon Equity Associates is an active equity and balanced strategic asset allocation manager. Although the firm has not made public how exactly they use NNs, it has been reported in [Enrado 94] that they found that certain nonlinear relationships between various valuation measures which were not being captured precisely through their own standard linear regression system in asset allocation process and specific stock selection. The use of NNs has improved the performance of their models. Currently the firm has over $38.8 billion assets under management, so finding better systems is always a priority.

In particular, they reported in [Enrado 94] that one of the problems of using linear techniques is that they fail to account for the difference in curvatures in the relationships between the price-to-earnings (PE ratio) of the equity markets and short-term interest rates. It has been observed that when extremes in interest rates exist, linear systems may show equities as being over-valued (too high PE ratios), whereas nonlinear systems may indicate equities as being undervalued. Thus, the information they were obtaining through their conventional methods was imprecise and prone to forecasting errors.

Finally, to summarise the success of their NNs, in the same article, a specialist from the firm pointed out: "our basic incentive to go with neural networks was to improve the performance of our models. So far, the gains have been fairly significant." No additional information has been found regarding the use of NN systems at present.

6.2.2 LBS Capital Management

Dean Barr and Ganesh Mani reported in [Barr & Mani 94] their results on experiments predicting S&P 500 Index values. Their first experiment uses 21 indicators: 7 inter-market relationships derived from prices of other markets such as the CRB Index and Dollar Index, and 14 indicators based on price, volume, and put-call ratio. A total of 126 inputs are derived from the 21 distinct inputs via telescoping and using data from 5, 10, 15, 20, and 25 periods back, in addition to the data from the current period. A network consisting of 126 input nodes, 14 hidden nodes, and 1 output node was selected as the most accurate during the testing phase, after experimenting with 6 different configurations (varying hidden nodes from 10 to 20 in steps of 2) and 20 learning parameters (learning rate from 0.5 to 1.0 in steps of 0.25). The training set was 164 trading days long and testing was done with the following 18 days. The accuracy is calculated from the correct predictions of the direction of change in the S&P 500, which is reported to be from 15-17 over 18 testing days during 3 different trials varying on the number of epochs. The authors indicated that the training time for such network was only 3-5 minutes in a 486 PC, and that once the net is trained, obtaining the prediction can be made almost instantly. However, this measure of time does not include the time it takes to chose the correct inputs and the *selection process* in which they performed a total of 120 experiments to obtain the best configuration (6 X 20).

A second experiment was made in which 4 themes (cyclical, defensive, growth and financial) were created with 4 large stocks fitting every theme as inputs, creating a type of *synthetic index*. They fed the net with the change of prices of the stocks of every theme along with other 2 indicators per stock (total of 3 inputs per stock), increasing to 60 inputs after telescoping them into 4 different values back (4, 8, 12 and 16 weeks back). The best net found had 8 hidden nodes (a 60:8:1) and the accuracy in direction of change of the S&P 500 was 69% on the test set.

Thirdly, as a rule extraction procedure, they performed *sensitivity measures* (i.e. with a trained net, vary the value of an input by a percent and observe the change in output) to infer which are the most relevant variables, which they used to generate approximate rules. They found that 2 indicators (P/E ratio and bull/bear sentiment) strongly affect the S&P 500 Index two weeks forward. With this information, they created rules which were used to enhance the probability of decisions within an expert system, along with other hand-crafted rules bringing a total of 268 rules. This expert system is called OMNI and is used by LBS Capital Management. Compounded annual returns from 1987 to 1992 of 26.80% over 6 years were reported using this system, against a 14.04% generated by *buy-and-hold*.

In addition to this system, the use of a Modular System of NNs called AXON [Ganesh & Barr 94] reported increases of more than 6% in invested capital within a few months of the system started in practice.

6.2.3 Hitachi Ltd, Fujitsu and Nikko Securities

Hitachi Ltd. researchers are working on the Neuro Portfolio securities selection system based on a massively parallel neural network. The neural network software, if run on Hitachi's supercomputer, "might take up to 10 minutes to select the optimum five stocks out of the 1,500 on the Tokyo Stock Exchange, as compared with four years or more using normal von Neumann-type computers. They created a model based on a 100-issue stock exchange; each neuron in the Neuro Portfolio model stands for one stock, connected by synapses. Hitachi has established a long-range, corporate-wide strategy to pursue neural computer-related development in a systematic effort" [Sakurai 88].

Takashi Kimoto, with the Fujitsu team and Morio Yoda of Nikko Securities [Kimoto & Yoda 93] jointly developed a Model of Multi-layer Perceptron, which gives 62% of accuracy in stock price predictions. The model predicts the Tokyo Stock Price Index (TOPIX) using economic and technical data. The inputs include price, volume, interest rate, foreign exchange rate, and the New York Dow-Jones Industrial Average. Every input is converted into three variables: a moving average, a regression coefficient and the difference between the current index value and the moving average.

They use economic and technical data as inputs, including price (regression coefficient), volume, interest rate, foreign exchange rate, and the New York Dow-Jones Industrial Average. Every input is converted into three variables: a moving average, a regression coefficient and the difference between the current index value and the moving average.

The system reports a 62% accuracy in stock price predictions and 1.6% annual return as opposed to -20.63% of the TOPIX *buy-and-hold* strategy.

They concluded that the system is fairly useful when "predicting stock price changes associated with factors that are readily adaptable to numeric representation, but not so useful for changes caused by other factors" [Kimoto & Yoda 93].

6.3 Genetic Algorithms

GAs are widely used to evolve portfolio management strategies. They are powerful optimisation techniques that can aid in the traders' decision-making process. While their usefulness has been acknowledged in the investment community, very few people describe how they use them.

One of the Institutions that has acknowledged a high GA activity is the The World Bank (Washington, D.C.), which has been using GAs for developing portfolio strategies. Guido Deboeck presents a good introduction to strategy optimisation through GAs in his book "Trading on the Edge", where he has edited a good collection of AI applications in finance. Another good introduction to GAs and investing strategies can be found in [Bauer 94].

Olsen & Associates Research Group has been working on GAs for a long time. In their library they have a good collection of papers addressing the use of GAs when building trader models [Dacorogna 93]. In [Pictet *et al.* 96] the authors specifically address the problem of designing trading models with relevant indicators and propose the use of GAs to optimise simple trading models.

6.4 Learning Classifier Systems

GAs and NNs have definitely found their way as accepted AI technologies in financial applications. Unfortunately, this is not the case for LCS. So far, no commercial system for financial predictions or decision making, based on LCS, has been found in the literature. It seems surprising that no banks, nor world leader research groups and research boutiques such as Olsen Group, which develops a great variety of AI commercial applications for predictions, nor any other successful companies in forecasting such as Prediction Company have been found to use LCS in their systems. However, only two economic models using LCS were found in literature In this section, these two important economic models using classifier systems, both of SCS type, will be briefly described. The first model was developed by an Italian group of researchers, Luca Beltrametti, Riccardo Fiorentini, Luigi Marengo and Roberto Tamborini, who reported in [Beltrametti *et al.* 97] a series of experiments with real world data from the *FOREX market*. They use two sets of exchange rates, with almost 20 years of data each. The first one corresponds to the the US dollar – D mark rate, and the US dollar – yen rate, which the authors believe is representative of two possibly different market environments.

This model is composed of a fictitious *artificial agent*, called AA, who observes a set of variables of the exchange rate, and which has three possible decisions at each time step: do nothing, switch one unit of wealth to dollars (i.e. buy 1 dollar), or switch one unit of wealth to marks (yen). The experiment does not consider commission payments when performing these transactions.

The reinforcement scheme rewards or punishes the decisions made by the agent, immediately after they have been made, trough a comparison of the exchange rate between the two currencies. Gains and losses drive the AA's search for better decision rules.

They train the agent with monthly data from 1973 to 1990, and then tested it out-of-sample, with the remaining 2 years, from 1990 to 1992. Training is performed in a special way. A number of runs are performed with the idea of keeping the rules evolved for the next training runs. i.e. good evolved rules accumulate, but the wealth is set back to zero after every run. Their reported results refer to the last training runs, once the agent had found a stable configuration of rules.

The forecasting performance of the AA is compared against the performance of some theoretical fixed rules (TFR), the random walk model of exchange rate (RW) and an estimation of the exchange-rate data generation process based on vector auto-regression (VAR).

Results reported appear somewhat mixed. However, in the yen market, AA performed better than VAR on the 8th run of the training set, but the opposite happened with the mark, where VAR outperformed AA. The RW rule behaves better than any other theoretical model, consistent with empirical findings in the FOREX market. Forecasting with the validation set was performed by not allowing genetics of rules, i.e. it is a test of whether the rules learnt during training could be carried over to a newly observed data set. Here only VAR was measured against AA. Results show that VAR performed better over the validation set.

At this point it is important to mention that this model represents two important similarities with the model proposed in this thesis: (1) real world information is considered and (2) the agent's learning process is represented by LCS. Therefore, it will be addressed throughout this thesis, specially in Chapter 7, where a number of differences between these two approaches will be discussed.

The second model found to use LCS, is a model of heterogeneous and imperfect rational agents, described in [Marengo & Tordjman 96]. The authors show that speculation of exchange rates is intrinsically, a disequilibrium phenomenon. This is caused by the heterogeneity of the "models of the world," or beliefs, developed by the adaptive agents. The dynamics of the asset prices are, therefore, affected by the changing learning properties of the agents and the non-equilibrium state of the environment.

They used SCS to simulate the behaviour of an artificial FOREX market consisting of heterogeneous, adaptive agents. As in ASMs, the goal is to study the aggregate results, in terms of prices and volume of transaction, of the interaction of such agents. Similarly to the SFI model (except that this refers to FOREX market), preliminary results show that series of simulated exchange rates exhibit some of the distinctive properties characterising real series [Marengo & Tordjman 96].

6.5 Genetic programming

GP is a growing area of research and applications. This section describes a number of selected GP applications, specially those dealing with decision-making and forecasting. One of the advantages GP offers is the ability to represent a candidate solution in the form of a decision tree, which, in turn, can be treated as a string using evolutionary operators. In addition, there is no restriction of having a fixed-size length in GP trees and the rules evolved can easily be evaluated and understood by humans.

Probably one of the most influential pieces of work dealing with applications – using real data – of GP systems in financial decision-making is the work of Edward P. K. Tsang, Jin Li and James M. Butler, who have developed an Evolutionary Dynamic data Investment Evaluator system (EDDIE), described in [Tsang *et al.* 98]. The system is intended to serve as a forecasting tool, designed to help improve the quality of investment decisions, not as a system to replace the role of experts. The idea is that the user leads the search by suggesting factors that are considered important. With historical data, the system creates hypothesis that the user can approve or reject in an interactive manner, or it can also suggest alternative rules.

EDDIE-3 is the system that was specially designed for the financial forecasting application, using S&P 500 data. Reported in [Tsang *et al.* 98], results of 10 runs were compared against the random strategy, showing an annual rate of return (ARR) of 53.59% in average over the test set, compared to 49.47% of the random strategy. In order to test the quality of the decisions, they used a trading behaviour that *invests one unit of money reflecting the index when the system predicts a positive position.* If the index does increase by 4% within 63 days, the system sells at a profit of 4%. Otherwise, it sells on day 63, no matter what the price is. There are no commission charges made when performing these transactions. When testing with this criteria, EDDIE-3 achieved 42.71% return in average over the test set, compared to 38.03% of the random strategy. These results clearly suggest that the GP approach is better than the random strategy.

Further research of the system was reported in [Tsang *et al.* 00], here called FGP (Financial GP). It is important to mention that the system not only searches the interactions between the variables under consideration, but it also searches the space of thresholds. For instance, if the price of the stock is above a certain limit, an investment decision could change considerably. The system searches for the value of such limit, generating rules that users can interpret. The goal to predict is whether the index is going to rise by *r*% within *n* trading days. The system has been tested with the S&P 500, the DJIA and FTSE indexes, as well as with 10 US shares and 10 Australian shares.

Again, it was compared against the random decisions.

Results suggested that patterns in indexes were indeed found, yielding over 65% accuracy in predictions. The rules predicted whether the price was rising by 4% in 22 days or not. Regarding stocks, patterns were found in some of them.

Further testing of FGP has been reported in [Li & Tsang 99a], where the DJIA index was analysed under the following goal: *the index will rise by 4% or more within the next 63 trading days (3 months).* The search space lies within the variables, operators and thresholds of technical indicators only, of the following type:

1. $P_t - MA12$
2. $P_t - MA50$
3. $P_t - P5_{lowest}$
4. $P_t - P63_{lowest}$,
5. $P_t - P5_{highest}$
6. $P_t - P50_{highest}$,

where P_t is the price of today, MAn is the moving average over the number of days denoted by n, and finally, the last four rules, which are of the type Pn, refer to the highest or lowest price over the number of days denoted by n.

Experiments were made with 3,700 trading days of the DJIA Index, from which the first 1,900 were part of the training set, the following 900 of the test set I and finally, the last 900 days were part of the test set II. The trading behaviour used to measure profitability is the same as the one described above, reported in [Tsang *et al.* 98].

They compared performances of 20 generated trees using FGP, versus the six technical rules, i.e., using rule 3, buy if $P_t > 1.001\%(P5_{lowest})$. The mean of FGP's accuracy and annualised rate of return over test set I and II were superior than those obtained using the six technical rules.

In addition to these experiments, [Li & Tsang 99b] report results compared against C4.5 [Quinlan 93], a machine learning type of classifier system which also generates decision trees. In this paper, the rules predicted whether the price was rising by 2.2%

in 21 days or not. The fitness used is the Rate of Correctness (RC) function, a measure of prediction accuracy, defined as:

$$RC = \frac{P+N}{N}, \tag{6.1}$$

where P and N are the number of *Positive* or *Negative* correct predictions respectively, and T the total number of predictions made.

Results over the test set show that the mean of RC of FGP is 54.78%, followed by C4.5 with 53.40% and finally, the random runs reported only 49.53%. Overall, results are encouraging and suggest that FGP and C4.5 were indeed, able to find patterns, contributing to higher performances that could not be explained by the random strategy.

Further improvements of this system are reported in [Li & Tsang 00]. In this paper, the authors propose a way to tune FGP's rate of failure according to the user's preference. This is achieved by adjusting RC, the fitness function used in earlier versions, according to a weighted sum of three parameters indicating performance rates. One of these parameters is RF, the *Rate of Failure*, and it is defined as follows:

$$RF = \frac{FalsePositives}{FalsePositives + TotalPositives} \tag{6.2}$$

Experiments were performed over DJIA index again and showed that with the appropriate parameters, performances using the new fitness formula were better than those using RC, showing that the former had a lower RF and was better in AARR, the *Average Annualised Rate of Return* and RPR, the *Ratio of Positive Returns*.

Note that while trying to reduce the number of failures, the number of opportunities are also reduced; using these constraints results in fewer positive recommendations. However, this property should not be considered a drawback as suggested in the paper, many investors want to perform less transactions as long as their AARR and RPR is not negatively affected. In these experiments, both of these increased dramatically when reducing RF.

In addition to the index, the system was run with ten shares and results were compared against various NNs and a linear classifier. Again, the best FGP rules proved to have found better RF than those obtained through the NNs. Linear classifier was the worse amongst the three methods.

To summarise all these, the combined results of a GP-based system were better than other three NN methods widely used: the time delay, recurrent and probabilistic networks. In addition, it was also better than the linear classifier and also better than C4.5, a decision-tree classifier. FGP also outperformed random decisions.

Chapter 7

The Model: ATOF World

"How poor would be the human mind without vanity! Thus, however, it resembles a well-stocked and constantly replenished bazaar which attracts buyers of every kind. There they can find almost everything, obtain almost everything, provided that they bring the right sort of coin, namely admiration." Friedrich Nietzsche.

7.1 Modelling the Financial Environment

Modelling financial markets is not an easy task. It was once said that "anybody doing applications of artificial intelligence (AI) in financial trading and investment management must possess considerable determination and vast reserves of optimism" [Essinger 90]. Why?

The generalised growth in complexity of financial markets has been among one of the most significant changes seen in the past decades. As a result, the difficulty of the decision making process has increased dramatically, becoming more dependent upon the analysis of information gathered and available data coming from many different sources and in various forms. The results of various quantitative forecasting techniques, as well as human subjective skills and judgement, serve as an important tool for strategic decisions to most financial institutions.

The globalisation and internationalisation of today's financial market requires an extensive use of engineering technologies for the development of computerised automation, and new software contemplating the needs of our society is oriented to a more efficient and intelligent way of handling information. Without these advances, a great deal of the financial technology developed would not have had such a significant impact.

However, the high levels of activity seen today, along with the vast amounts of information available produce multidimensional financial time series where the general mechanism that generates them is poorly understood. Much research has been devoted to try to find some kind of model with a good explanatory power of the variables involved, being of relatively low complexity. Therefore. with the hope of finding the key to increased wealth, a common approach that has been followed is to try to look for previously undetected non-linear regularities (such as daily or intra-daily fluctuations) from economic time series, which can be done in several ways.

In some models, data derived from only one time series is used, such as a specific stock's price, a given index or the one-day rate of return to holding a stock. Readers interested in a neural net (NN) model of this type are referred to [White 88], where inputs are defined as the *one-day return of a stock* r_t[1], and for a genetic programming model where inputs are defined as the difference between the current price of a single stock and its moving averages, highest and lowest prices, refer to [Tsang *et al.* 98, Tsang *et al.* 00, Li & Tsang 99a, Li & Tsang 00]. These GP models were explained in more detail in Chapter 5.

In addition to using a single price series, other approaches use more variables and with different levels of sophistication (see [Barr & Mani 94, Ganesh & Barr 94, Kimoto & Yoda 93], also explained in Chapter 5). These models can include over 120 inputs coming from other stock prices, volume of transactions, macro economic data, market indicators such as indexes, profit and earnings ratios, etc. Other models even include people's personal opinions, which are extracted from electronic news boards, such as [Wuthrich *et al.* 98, Thomas 00, Thomas & Sycara 00].

It then seems reasonable to suppose that in financial markets, investment decisions usually involve taking into account a significant number of elements and their relationships, which due to the complex nature of economic systems, are usually difficult to

[1]The one-day return of a stock is defined as $r_t = \frac{p_t - p_{t-1} + d_t}{p_{t-1}}$, where p_1 is the closing price of the stock in question on day t and d_t is the dividend it paid on day t. This value should be adjusted to stock splits and dividends if appropriate.

understand and highly non-linear. Even though some investment managers and traders can factor in such dependencies in the analysis, they often cannot explain their decisions and point out the significant factors and relationships contributing to their current positions; other times, due to secrecy policies, they are not willing to do so. This makes the problem of economic forecasting very hard, as there is no actual knowledge easily available and controversies and debates as to what is the right investment approach are still in question, as described in the Economics and Finance Chapters, 2 and 5 respectively.

Therefore it is probably not sensible to devote any more efforts to try to acquire such knowledge from the experts and embed it in a learning system. In addition, obtaining such market knowledge or expertise would be very difficult and costly. For this reason the present financial model starts with no previously acquired knowledge from the experts.

The problem here is to create a number of *adaptive agent-types*, capable of analysing and classifying various historical and non-historical market data to incorporate it in an investment decision-making process where the agent will trade upon in order to survive and hopefully grow in assets. The main concept is that the agent-type is born with no prior knowledge about the market and it is forced to start trading its funds (buying, selling or holding) since day one. The simplistic on-line nature of this learning process is one of the main features of the model. The decision-making process can be performed daily or intra-daily, depending on the availability of the data. All experiments shown in this thesis use daily data, so that the trader's decision is performed on a daily basis, at the end of each day, and using the *close price* of the stock and other pieces of information that will be explained later.

Similarly to the format given in section 4.4.2, a segment of data, freely available from `http://quote.yahoo.com`, looks as follows:

```
Date,Open,High,Low,Close,Volume
30-Apr-01,46.75,46.75,45.85,46.19,3837000
27-Apr-01,47.50,47.50,46,47,3480900
26-Apr-01,47.60,47.98,46.90,46.90,4296400
25-Apr-01,47.50,48.40,47.40,48.20,3247700
```

```
24-Apr-01,47.45,48.45,47.17,47.39,4343400
23-Apr-01,47.50,47.90,47.03,47.45,3330500
20-Apr-01,47.05,47.25,46.28,47,4583600
19-Apr-01,46.05,47.77,46.04,47.50,5418200
18-Apr-01,46.30,47.08,45.71,46.75,7383700
17-Apr-01,45,45.70,44.81,45.70,4407900
16-Apr-01,44.57,45.50,44.15,45.41,4143500
12-Apr-01,43.50,44.90,42.65,44.57,5214600
11-Apr-01,43,44.75,42.37,44.13,6874200
10-Apr-01,43.91,44.30,43.21,43.65,5948200
9-Apr-01,45,45.59,43.90,43.90,4115800
6-Apr-01,44.43,45.35,43.85,45,3953900
5-Apr-01,46.25,46.60,44.62,45.30,4621100
4-Apr-01,44.52,46.05,44.17,45.25,4399300
3-Apr-01,45.60,45.70,44.18,44.66,4672700
2-Apr-01,45.40,46.92,44.86,45.85,5212400
```

As it can be seen, the data is given *daily*. It starts with the Date. Date formats vary between sources. The following four columns describe the various prices that are recorded daily, starting from the opening price of the day, then the highest and lowest of the day, followed by the closing price. Finally, the last column displays the volume of transactions relative to that stock. Note that dates are given in descending order, they usually need to be re-ordered in ascending order.

Now that the environment surrounding the model has been described, two important considerations will be discussed before describing the problem and the proposed model. The first one gives a brief summary on how to build artificial stock markets, followed by describing two examples where artificial markets that use real data have been built.

7.2 Building the Artificial Market

Building an Artificial Stock Markets is not a trivial task. Before doing so, for those interested, it is highly recommended to review the literature regarding these issues, which due to space limitations, cannot be covered in detail here. Blake LeBaron has done extensive research in the new and emerging area of building artificial markets. For a good review, refer to [LeBaron 00, Chan *et al.* 99].

It was explained previously that there are still a large number of questions about market and agent design that remain unanswered. Perhaps one of the most difficult issues to tackle is: how to represent and implement ideas of evolution and learning in populations of traders, trying, at the same time, to keep the economic model as simple as possible? Usually, in an ASM, active traders, through their consumption and portfolio decisions, endogenously determine the price sequence in the market being designed. However, it has been explained that the motive here is different: design an artificial market with real data, where agents are not affecting the price. Has this previously been done?

Yes, but only partially. Martin Lettau uses external data in his model [Lettau 97]. His model, for simple agent benchmarks, can be regarded as "probably being the best place to start for anyone thinking about constructing artificial agents" Furthermore, LeBaron adds when describing Lettau's model, "the price is given exogenously. This assumption may appear to go against the entire concept of building artificial markets, but it allows Lettau to concentrate on the agent's behavior. Since the actual behavior of evolutionary agents is not well understood, this seems like a very good idea" [LeBaron 00]. The agents of this model decide how much of a risky stock to buy, against risk-less a bond paying 0% interest. In this model the price of the stock is not *a real price*, it is artificially adjusted according to a random dividend from a Gaussian distribution. The goal then is for the agents to find the optimum portfolio, which they do achieve. Learning is modelled with GAs, which seems like a good idea for various reasons.

One is that there is extensive theoretical and empirical evidence that GAs have robust adaptive ability in ill-conditioned search spaces [Goldberg 89, Holland 75, Smith *et al.* 99]. In addition, such techniques allow the system to improve and discover new strategies through experience, as the ability to perform well depends on the environmental feedback. The economic notion that agents making systematic mistakes have an incentive to modify their behaviour parallels the natural selection of good genes from a diverse gene pool [Routledge 94]. Let's now move into some modelling issues.

An economic model is considered, in general terms, "a set of decision-making mechanisms, organisational arrangements, and rules for allocating society's scarce resources" [de laMaza *et al.* 98]. However, real economic processes are complex in the sense that they are difficult to decompose into separate parts that can be studied in isolation and then added to get the whole picture of the system. For this reason the attempt is, through computer simulations, to integrate certain features that behave, in some respects, like human traders in this type of environment.

At this point, what other aspects need to be considered? Several. For instance, in simulations of real phenomena, a choice has to be made of what are the important components of the world to be simulated. As a consequence, there are no unique solutions offered and quantitative conclusions are difficult to obtain. However, some general properties, behaviours and trends can be inferred, making simulations useful and manageable techniques. The idea is to develop a computational approach which could offer new insights to the hard problem of experimenting with real socio-economic systems under controlled conditions. In addition, such model must be adaptive.

Evolutionary Computation (EC) methods are one of the most natural machine learning techniques to transfer general-purpose adaptive capabilities to agent-based systems [Smith & Taylor 98]. For instance, agents can benefit from adaptive methods in strategy acquisition through reinforcement learning. In Reinforcement Learning (RL) problems [Sutton & Barto 98], reactive and adaptive agents are given a description of the current state and have to choose the next action from a set of possible actions so as to maximise a scalar *reinforcement* or *feedback* received after each action [Sen & Sekaran 95]. Holland's Bucket Brigade algorithm [Holland 75] has been a widely used feedback distribution scheme.

All these considerations seem relevant to the problem in hand, however, can all these be taken into account in the model? If this is so, how? The following section addresses these issues.

7.3 The Approach

Many techniques have been used to model financial time series data such as stock prices, including linear and nonlinear statistical methods and more inscrutable methods such as neural nets. The hope is that if the time series can be predicted fairly well, the predictions can be used to make profitable buy/sell/hold decisions. Time series modelling may work well for physical systems, such as blood flow in veins, because the underlying physical laws do not change with time even though their effects (such as the rhythmic expansion and contraction of blood vessels) may do so. However, time series modelling in the financial world often seems to fail to track any of the causes of the data, and the causes themselves change over time. All that happens is that the model ceases to fit the real data well, and then it has become time to re-train or re-fit.

So what can we do? Continual re-training may not be the answer; in the case of (e.g.) neural networks it may be expensive, and certainly performance cannot always improve, so there is a potential problem about when to switch from the old model to a new one. Another drawback of such models is that they do not tend to identify or explain anything about the relevant phenomena that is being modelled; this is not their main goal. As a result, there is no added knowledge or hypothesis about the underlying causes of the changes in the environment being modelled.

The approach here is different. Because financial markets are complex systems where the players are human traders, the aim is to address these issues by focusing on the agent rather than on the data. The idea is to model traders' decision-making processes, as they decide repeatedly whether to buy, sell or hold a particular stock. Modelling individual traders may not be ideal – as described in Chapter 5, individuals are extremely complex, their actions are affected by their private knowledge, their emotional make-up, even their breakfast! Why then, not model the behaviour of whole groups of traders instead? The underlying hypothesis is that the aggregate behaviour of a group will be simpler to model than that of individuals and that it may still produce very good results.

Modelling the behaviour of a *group of traders* seems a sensible idea. Why? In this work each group of traders uses a certain set of real market information before making decisions. To make a simple start, the groups receive small and reasonable-looking sets, consisting of some binary indicators about how today's prices and volumes relate to various moving averages and historical extremes. The behaviour of the group is expressed as a set of rules that relate to a single stock, such as: if A and B and not C and D then buy 45% of what cash-in-hand allows; if not A and B and not C then sell 15% of current holdings; etc. These elements will be explained in the following sections.

It is important to emphasise that *trader types* are evolved rather than models of individual traders. The assumption is that the aggregate behaviour of a set of traders, each of whom is (say) basing buy/sell/hold decisions on certain sorts of data such as information about various moving averages and recent highs and lows, will be simpler to describe than the complex behaviour of only one trader. Any one trader might include in his decision-making factors such as personal knowledge of the people controlling the company, private knowledge of the local economy, etc. However, within a group who use the same basic methodology, the effects of such extra factors may tend to balance out so that the group's behaviour may at least be approximated in terms of only a modest set of easily available indicators.

For example, a *trader type* may pay attention to such factors as whether today's price is more than 20% higher than the 5-day-moving average, whether today's volume is larger than yesterday's and so on. At present, a *trader type* is characterised by the particular set of such factors f_i – each of them binary, a simple yes/no – that it pays attention to. Its behaviour is expressed in terms of rules whose conditions are a combination of terms such as $f_1 =$ yes, $f_2 =$ no, $f_3 =$ don't care (the *don't care* are wild-cards, they match either yes or no). Rule actions can be to do nothing (hold), or to buy more stock using a certain proportion of current cash invested in the risk-free bond, or to sell a certain proportion of holdings in the stock and invest that in the risk-free bond.

7.4 The Goal

The goal of each of these *trader types* is simple. They are all trying to learn behaviours (rules, market strategies which they create and improve) that will lead them to increased profits under the current market situation, i.e. they are not looking for the *optimal* behaviour, only an appropriate one. Such behaviour might even include things like "do nothing" apart from buying or selling the stock in question. The system does not directly predict price values (e.g., it does not tell you that the price tomorrow will be £123.86 like some NN models do) or the direction of change (price will rise or fall), but rather it delivers a specific action for the agent to follow, which will be rewarded only if it turns out to be profitable. Such decisions represent a measure of profitability. But what tends to drive such criteria?

A common approach to quantitative decision making in finance time-series is to train a model using a prediction criterion such as minimising the Mean Squared Error (MSE) to make predictions of the next value of the security. The measure in this model is financial profitability, not accuracy in the prediction.

A similar approach regarding the *system's goal* was reported in a NN based system developed by [Bengio 96] that shows better results when the model is trained to optimise a financial criterion rather than an MSE criterion. Results with the *buy-and-hold* benchmark yielded 6.8% yearly return on a 6 year testing period. Optimising a financial criterion yields 14.2% as opposed to 9.7% for the MSE trained network, both on the test periods.

7.5 LCS Model Advantages. Why LCS?

Having described the problem in hand, now the question to analyse is the following: is it feasible the use of LCS to represent the agents' learning process, specifically in the stock market? If so, what sort of performance could these agents achieve?

This research involves modelling sets of behaviours of adaptive agents in a complex environment: the stock market. As stated in [Holland *et al.* 86], the original idea behind the development of learning classifier systems was to model inductive processes in human-like cognitive systems. The goal is to search for a plausible way of representing how intelligent systems learn to improve by constructing internal models of their surroundings. In other words, the aim is to build a system that will aid in the study of how these intelligent beings co-evolve with their environment. LCS seemed like a well-suited candidate to explore this complex problem.

However, LCS, even in their most basic form, are "particularly difficult to analyze because of their complexity" [Horn *et al.* 94]. The simple LCS is a much more complex model than the simple GA, but no one has argued that modelling the human decision-making process is a simple one either. The competition-cooperation intrinsic character of LCS contributes to the greater complexity of LCS than GAs. Such added complexity is induced by several apportionment of credit mechanisms such as the bucket-brigade, specificity dependent bidding process, internal message lists, etc. However, within the boundaries of such complexities, the approach here is to design a system made of simple parts acting together, following the good old KISS acronym: "Keep It Simple and Stupid." As Goldberg suggests, add complexity later [Goldberg 89]. Some key features that make LCS a good candidate for this task are the following:

1. The beauty of LCS resides in that adaptation and learning occur in two different contexts: the first one, which I call the "micro-adaptation" level, happens at a short time interval, through reinforcement of the good strategies in a continual process of adaptation to an environment which is completely unknown initially. This, being the apportionment of credit algorithm component of the LCS, can be implemented as the bucket brigade algorithm, a profit sharing plan or Q-learning. In this scope, the agent can be seen as being in a sort of *do the best with what you've got* mode. The second level, which I call "macro-adaptation", offers a more powerful evolutionary component where discovery of new, invented strategies, while maintaining the diversity of the pre-existing mixture of strategies, is essential in order to "view" the entire process from a global perspective, and it is implemented via the GA. Here, the agent is in a perpetual ... *keep looking for different, new alternatives* mode.

2. LCS are well-suited to solve multi-objective problems because their design contemplates niching to tackle such tasks. Note that the search here is not for a

single rule that is capable of solving the task completely (optimality here can not be achieved and therefore it is not the issue), but instead the idea is to obtain and keep diverse rules in the population that together help the agent increase its capital.

In addition, the designed N LCS-agent topology of the model will allow to simultaneously explore many competing market hypotheses at two distinct levels:

- at the *trader level*, where the trader's computational elements possess the ability to map a significant number of states of the environment into a limited number of compact internal representations (usually 100).
- at the *trader-type level*, where each type is designed to be part of a different sub-environment which in turn, is part of the whole market environment being modelled. This promotes the specialisation of pre-defined sub domains (usually 3).

7.6 Market Structure

This section describes the market structure of the proposed model, also described in [Schulenburg & Ross 99, Schulenburg & Ross 00, Schulenburg & Ross 01, Schulenburg & Ross 02]. The model, called from now on Artificial Trader-Oriented Financial World (ATOF World), consists of the following elements, which will be described in more detail in the following sections:

1. **Time**, which is discrete and indexed by t, represents one cycle equivalent to one *real* trading day in the market. There are only about 253 trading days in a calendar year due to weekends and holidays, so when we refer to a ten year period of historical data in the following sections, it roughly corresponds to a total of 2,530 days.
2. Two assets traded, both with infinite supply: a **risk free bond** paying a fixed interest rate. In this case it is equivalent to a real bank's investment, and a **risky**

stock whose return will be ruled by the specific stock in question, with the restriction that the agent must be able to afford the amount of shares it wishes to buy.

3. One **buy-and-hold agent** which represents the so called *buy-and-hold* strategy – in this case, this agent simply puts all available cash into the stock at the start and then keeps it there, buying the stock at the initial price P_t, paying a commission percentage when doing this single transaction. Many people believe this is the best performing strategy in the long run. It is of great relevance because it is necessary to compare the adaptive trader's performances against a strategy which works very well for highly performing stocks.

4. The **bank agent**, which keeps all money in the bank at a good rate of interest, never buying the stock. Therefore this agent does not own any shares, all its possessions are cash, compounded in the bank at an interest rate of 8% p.a. (but the user can alter this value). When given shares, it immediately sells them, paying the appropriate commission for the transaction.

5. One **trend-following agent**, representing a strategy that varies according to price moves. This is a type of *momentum trader* that buys all its money available in stocks at the end of the day if the price increased with respect to the previous day. Otherwise, it sells all the shares owned. This agent also pays commission for every transaction performed.

6. The **information set**. This is the available raw data about the market, as well as data which has been processed in various ways. It includes basic daily information about the stock such as its current price, volume of transactions, splits and dividends, and some derived information such as price differences, moving averages, current standing of the *buy-and-hold* strategy and the *bank* investment. It will be explained in section 7.8.

7. An **accounting procedure**. A basic accounting procedure takes place for every agent. It calculates and updates the possessions of every trader and the various other strategies (bank, buy-and-hold, etc) according to the new price; it works as

follows: At the beginning of each trading cycle, the agent's holdings and cash accounts are updated, including stock splits, cash dividends and commissions for all transactions executed. If the trader decided to buy it must own all the cash incurred in the transactions, including the commission fee. When selling, the trader must own all the shares it wants to sell. Agents cannot borrow money or sell short. At the end of the day, the wealth of each trader is updated by adding the interest paid during one cycle to the cash account. The wealth $W_{i(t)}$ at time t of agent i is thus given by the equation:

$$W_{i(t)} = (1+r)M_{i(t)} + H_{i(t)}p(t), \tag{7.1}$$

where $(1+r)M_i(t)$ is the cash invested at an interest rate r and $H_{i(t)}p(t)$ are the holdings calculated at the current price of the stock $p(t)$.

8. Three **heterogeneous agents**, which will be described in the following section. Note that in this model the stock price does not change according to the supply and demand governed by the artificial trader-types, but rather by changes of real phenomena outside their scope. Therefore the agent's actions do not affect the price of the stock.

7.7 Trader Types

A *trader-type* (also referred as *agent*) of the present artificial stock market model is represented by a Michigan-style strength-based Learning Classifier System (LCS) as described in Chapter 3. It is designed to learn and adapt to a market environment that is partially understood and where the domain characteristics can change rapidly over time, but how many, and how simple can these agents be?

There seems to be many relevant factors, but there exists some work already, which uses a number of simple agents when trying to simulate financial markets [Ankenbrand & Tomassini 97], [Rajan & Slage 96]. Such work typically uses (i) various heterogeneous agents of each type (n_1 of type 1, n_2 of type 2, etc.), which might lead to considerable computational complexity, or, (ii) a number of – usually called – heterogeneous agents or market participants [de laMaza & Yuret 94], [Palmer *et al.* 94] who in reality only differ from each other in the strategies they evolve (generally through a genetic algorithm) while all receiving the same type of information about the market, which does not resemble realistic features of diverse traders.

In order to keep the model as simple as possible, only three types of traders have been introduced (that is, three agents) in the market and to make them truly heterogeneous, they all receive different sets of market information and make their decisions by using different models. In what follows, the three types (agents) are called **Tt1**, **Tt2**, and **Tt3**. One type can not evolve into another type, but it can go broke and thus halt while another one can get rich. It is also possible for an agent to learn to ignore any particular field in the daily market information it receives, but it cannot ask for extra fields beyond those given to it.

In the work reported here, the market environment consists of genuine daily price information about a specific stock, typically over a ten-year period. The agents make investment decisions of whether to buy, sell or hold different proportions of the risky stock and the risk-free bond (essentially, an immediate-access bank account paying a reasonable but fixed rate of compound interest). Having different views about the market, each type of agent creates, develops and explores a vast pool of expectational models based on the recommendations of those strategies that perform best over time.

Although it would be simpler to give each different agent type the same information, in the present model it is also possible to choose to vary what an agent sees according to its nature. The agent types are distinguished from each other by the fact that each day, they receive different sets of information-environmental messages.

Each agent buys, sells, or holds its possessions and adapts by receiving feedback from the changing environment by monitoring, updating and evolving new procedures. The idea is that agents therefore continually form individual, hypothetical expectational models of "theories of the market," test these, and trade on the ones that predict best. An agent might interpret a certain state of the market (events) according to its past performance. Thus there is a dynamical approach in the sense that future behaviour is affected by past performance importing feedback (new observations) from the environment. This inductive way of reasoning has been explained in more detail in Chapter 2.

7.8 Information Sets

The Information set simply represents the market state, which is summarised in a binary state vector of fixed length for each trader type, i.e. Tt1's state vector could be seven bits long, while for Tt3, the length might be only five, as described in the previous section. This is due to the fact that each trader type receives different information about the market. The traders are thus limited to put their position based on the state variables given to them, but they are capable of ignoring any of these state variables. This is an important feature of LCS, which makes them similar to GP in the sense that the length of the rule can vary in size. This information selection capability is controlled by the # symbol explained in Section 3.4.1, which implies that a match exists no matter what the value of the specific bit of information is. This will be explained again in Section 7.9.

Initially the agent's rules are randomly generated, although with an inbuilt bias towards generality. Thereafter they are periodically modified by the LCS, so that rules which perform better than average are more likely to produce descendants, but the forces of recombination and mutation also work to produce new and steadily better rules.

Each type of agent controls its market position by using and evolving a set of classifiers. These classifiers are used to match binary conditions in the market, which requires predefining a set of binary states or environmental messages for each agent that can be used when making decisions. All three agent types receive the first difference in price (1 if the current price of the stock is higher than the previous day's, 0 otherwise), followed by a basic set of market information which differs for each type of agent.

The first information set contains **price statistics** and **moving averages**, based on standard moving average types of trading rules, the second set basically contains **volume statistics**, and the third set incorporates apart from the price and volume, some other **news** such as the accumulated wealth of both the buy-and-hold strategy and the bank's, and a recollection of its past action. Only these simple environmental messages are reported here, because we want to study simple models first. It is possible in this model, for example, to allow an agent to make use of what other agent(s) chose to do on the previous day, allowing a poorly performing agent to switch from seeking to improve its wealth to seeking to emulate another, better-performing agent but still using its own information set for the purpose.

For present purposes Tt1's state vector is seven bits long, Tt2's is five, Tt3's is six, but these can be defined differently by the user. Table 7.1 shows the meaning of each position for Tt1 and Tt2 during the current trading period t. The actual value of each cell is 1 if the condition is satisfied. For example, bit number 1 is 1 if the price today is higher than yesterday's: P_t is the price on day t, P_{MA5} is the five-day price moving average, $P_{highest}$ and $V_{highest}$ are the highest price and volume for any day so far, etc.

Table 7.1: Environmental Message for Trader Types 1 and 2

Trader Type 1

Bit Number	Representation
1	$P_t > P_{t-1}$
2	$P_t > 1.2 * P_{MA5}$
3	$P_t > 1.1 * P_{MA10}$
4	$P_t > 1.05 * P_{MA20}$
5	$P_t > 1.025 * P_{MA30}$
6	$P_t \geq P_{highest}$
7	$P_t \leq P_{lowest}$

Trader Type 2

Bit Number	Representation
1	$P_t > P_{t-1}$
2	$V_t > V_{t-1}$
3	$V_t > V_{MA20}$
4	$V_t \geq V_{highest}$
5	$V_t \leq V_{lowest}$

Similarly, the third trader's market state vector is summarised in table 7.2. The third trader differs from the others in a couple of ways. The first one is that it receives both the difference in price and the difference in volume, with the idea that it will use this in conjunction with an additional piece of information, and that is "how" it stands in the market. For this reason it receives a comparison of its daily accumulated shares with respect to the total shares owned by the the buy-and-hold investment, represented by $Shares_t$, as well as information about the accumulated wealth of its competitors: the *buy-and-hold* and *bank*. Note that it also receives its own previous action, so there is a 'reflective' aspect to this trader. The other two types do not receive a note of their previous decision as part of the environmental message.

Table 7.2: Environmental Message for Trader Type 3

Bit Number	Representation
1	$P_t > P_{t-1}$
2	$V_t > V_{t-1}$
3	$Shares_t >$ Buy and Hold $Shares_t$
4	$W_t > W_t$ Bank
5	$W_t > W_t$ Buy and Hold
6	Tt3 action at (t-1)

7.9 The Rules

As usual, classifier rules contain two parts: $< condition >:< action >$. The condition is a bit string matching the current market state vector which has a fixed length with each position taken any of the three values 0, 1 or #. The 1 and 0 match corresponding bits in the state vector, while # is a wild-card symbol which matches both 0 and 1 of the conditions and allows traders to dynamically decide which pieces of information are relevant to them by ignoring some state variables if they wish to do so. This symbol is responsible for providing the trader with an information selection property.

The action is a bit string of length 4 indicating the decision of whether to buy, sell or hold possessions on the current day. The first bit represents a buy (1) or sell (0) signal, followed by three more bits which represent the percentage of available cash to use if buying, or the percentage of shares to sell if selling. Table 7.3 shows the mapping between bits 1 to 4 and the percentage to trade. Buying or selling 0% corresponds to simply holding.

Table 7.3: Transaction Percentages

Bit 1	Bit 2	Bit 3	Bit 4	Transaction	Trade
0	0	0	0	H	0%
0	0	0	1	S	15%
0	0	1	0	S	30%
0	0	1	1	S	45%
0	1	0	0	S	60%
0	1	0	1	S	75%
0	1	1	0	S	90%
0	1	1	1	S	100%
1	0	0	0	H	0%
1	0	0	1	B	15%
1	0	1	0	B	30%
1	0	1	1	B	45%
1	1	0	0	B	60%
1	1	0	1	B	75%
1	1	1	0	B	90%
1	1	1	1	B	100%

7.10 Decision Making Process

Although in general terms the goal of the artificial agents is to maximise their profits by making investment decisions regarding their current portfolio, they form their expectations in very different ways. Each day, they all have a choice of (i) leaving their money in the Bank, in which case there is a fixed interest rate paid on a daily basis equivalent to 8% annually, or (ii) they can buy or sell a stock, represented by a given **real** stock, such as Merck & Co.

7.10.1 LCS Representation

The framework for representing the adaptive agents relies in the following LCS components, more throughly described in [Holland 92, Holland 95]:

1. The performance system, consisting of (i) detectors for extracting information from the environment, (ii) classifiers to represent the agent's capabilities for processing the environment and (iii) effectors, through which the agent acts on its environment.

2. Credit Assignment Algorithm. For situations in which many rules fire at the same time, it is responsible for strengthening rules that set the stage for later more rewarding activities.

3. Rule Discovery Algorithm, through which plausible hypotheses are generated and past experience is incorporated.

7.10.2 Credit Assignment and Rule Discovery

Two levels of learning can be identified in the model:

The **first level** happens rapidly as the agent learns which of his strategies are accurate and worth acting upon, and which ones should be ignored. This happens during the auction among currently matched classifiers in the apportionment of credit algorithm.

The **second level** occurs after several trading periods have gone by, when the structure of these market strategies is modified by the use of a Genetic Algorithm which replaces some of the worst performing rules by new ones created with portions of the better rules.

Associated with each classifier rule, there are other parameters necessary for (i) the apportionment of credit algorithm, which is encharged of picking the winner rule among all the current matching rules and for (ii) the reinforcement algorithm, encharged of giving a reward to those rules that matched the market state and also satisfied a given criterion. These parameters are the classifier's strength, bid, ebid, match flag and specificity. More detailed definitions of these parameters can be found in [Goldberg 89].

On each trading cycle zero, one or more rules can match the current state of the market, but only the action of the best one, amongst the ones relating to a specific market descriptor, will be used in that period. The credit-distribution system used in the present model is a simplified bucket brigade algorithm in which a payment equal to the bid of the current winner is transferred **only** to the old winner. For each agent there exists an option of setting this payment distribution on or off. No payment chains are implemented because we are not interested in rewarding sequences of classifiers for more than one step. The assumption here is that the current decision to buy, sell or hold the stock should not depend on previous decisions. This is a single-step environment where, for instance, the decision of Friday is not dependent on the decision taken on the previous Wednesday, whether it was good or bad.

In order to describe these two levels of adaptation, consider the following bit string, which reveals the meaning of each position for one type of trader during the current trading period t. The actual value of each cell is 1 if the condition is satisfied. For example, the value of bit number 1 could be 1 if the price today is higher than yesterday's, and so on. For the given current market state **1 1 1 1 1 1 0**, there is a subset of matching classifier rules which might be advocating the same or different actions:

```
1 1 1 1 1 1 0 : 1
1 1 1 1 1 # 0 : 1
# 1 1 # # # # : 0
1 # 1 1 1 # 0 : #
```

Each rule of this subset of rules makes a bid proportional to its strength and specificity, and only the one with the highest combined worth is selected. The winner then pays its bid, which according to Holland's Bucket Brigade Algorithm, is distributed among classifiers which contributed to its activation in the previous steps. As explained earlier, this is not the case in this model since the activation or success of a particular classifier rule is independent of previously active classifiers, therefore there is no need for bucket brigade chains of winning classifiers systems in this particular problem. The corresponding action of the winning classifier rule is sent to the environment, which then provides with a reinforcement indicating its goodness. It is in this sense that the first level of learning occurs.

In general terms, the rule discovery system makes use of a genetic algorithm that is activated every so often to allow the agent to improve the quality of its rule-set. This corresponds to the second level of learning, and it works as follows:

A proportion of the population is selected for reproduction during a given genetic algorithm invocation. In the examples given in this thesis, this figure is usually set to 20% (as suggested in [Goldberg 89] when solving the six-multiplexer problem starting from a randomly generated set of 100 rules). Using roulette wheel selection, a pair of mates is selected, where crossover and mutation take place in the normal way and the replacement candidates are chosen from a *low-performance subpopulation*, rather than being chosen from the full population at random, as it happens in the conventional crowding mechanism. Since the search is for a well adapted set of rules, crowding replacement is used to choose the classifiers that die, inserting new offspring in their place on the basis of similarity. For more details on how the genetic algorithm works, refer to [Goldberg 89] or to Chapter 3 of this thesis.

Both levels of learning allow the agent to achieve the following goals:

1. **Creation.** Primarily, new rules are derived from existing ones. This occurs every so many generations, when the GA is applied and rules are inserted in the set, replacing other bad performers. Creation of rules also happens when an incoming message does not match any rule. If no existing rule matches the input, a new rule is created to match it. In this model, because more specific rules tend to be more valuable to a trader, an additional pressure has been introduced here. The new rule's condition part is an exact replica of the market condition encountered – no extra wild card symbols are added. This is explained later in Section 7.11, under *automatch* flag. However, In XCS, some random wild-cards (usually $p_{\#} \sim 0.33$) can replace some bit values in the condition part of the rule to allow the rule to be tested later on in a number of different situations – this *covering* mechanism was described in Section 3.4.2.5. It is also given a random action and low initial prediction, depending on the classifier used. For example:

   - If Input is: `11000101` and no action matches,
   - New Rule is: `1##0010# : 011` $\longrightarrow$ 10.

2. **Generalisation.** Any LCS looks for maximally general rules. When one classifier is a generalization of another one with the same strength, the more general one will have more reproductive opportunities than the more specific one because it will match more states, and if it is correct more times, it will receive more environmental payoff. Consider the following two rules in an XCS: $R1$ and $R2$ with the given *prediction* `P`, *error* `e` and *fitness* `F` values:

```
                      p      e       F
R1: 10#001:0   ==>  1000   0.001    920
R2: 10##0#:0   ==>  1000   0.001    920
```

   The more general occurs in more action sets and as the GA acts on the action set [A], it will reproduce more. Therefore R2 will gradually drive R1 out of the population.

3. **Specialisation.** Transformation of general rules into more specific rules. While an LCS searches for accurate general rules as explained in the previous paragraph, the formation of specific rules is also encouraged in such systems. The population of classifier rules is initialised randomly, with low-specificity rules; that is, there is a higher proportion of wild-cards in the condition part in order to cover more market states initially. This model presents three types of pressures to encourage the formation of specific rules over general rules: the first one, described in [Goldberg 89], is implemented in the *apportionment of credit component*, in the bidding mechanism, by making the rule's bid proportional to its strength and specificity; the second and third ones, which are a novel property of this system, are implemented in the *reinforcement component* and in the *covering* mechanism, known in this model as *automatch*. Reinforcement is applied according to specificity as well, and *covering* inserts only specific rules. These will be described in more detail in Sections 9.1.2 and 7.11.

4. **Diversification.** Introducing more heterogeneity in the rule set. For instance, good diverse sets of rules can be kept in LCS due to the use of a niching mechanism, such as the ones described in Section 3.4.1.2 – this model uses a modified

crowding scheme – and by apportioning the bid payment among the matching classifiers to ensure the formation of diverse rules, which together, will cover different types of behavioural requirements that provide ample payoff to the system, without undue interspecies competition.

7.11 Description of Parameters

The simulations start by randomly initialising the strategies of each agent, all with strength values equal to 10. There are 100 classifier rules per agent, although they are designed to have any number of classifiers, completely independent of each other. The probability of having # symbols, $P_{\#}$, in classifier conditions, is usually 0.5. The rest is filled equiprobably with 0s and 1s. The reason to set the probabilities this way is because the goal is to start with more general rules to match the given environmental states. The action part has no # symbols; bits are 0 or 1 with equal probability. The measure of performance of any given agent is the amount of money it accumulates through its actions and it is calculated at the end of each day.

An example of **arbitrary parameters** for the three agents are shown in table 7.4. Note that this set of parameters does not necessarily correspond to a true run. It only intends to illustrate which sort of values can be assigned. The possibilities are ON/OFF for the binary ones (flags), and almost any number for the rest. These values will be explained in detail in the following paragraphs.

GA period refers to the number of generations or cycles that have to pass before the genetic action takes place. No GA invocation is denoted by *GA period* = -1, and the maximum GA invocation possible would be equal to the number of trading cycles. For instance, if there are 900 days to trade in the simulation and the GA is invoked every day, then the *GA period* = 900, which corresponds to the maximum number of possible calls to the GA. Although this parameter has been experimented from values of [-1, max], a good range of this parameter normally goes from [50, 200]. Fifty corresponds to a very fast learning rate (GA is acting very often, every 50 days), while 200 corresponds to a slower learning rate (every 200 cycles the GA is called). In all cases presented in this thesis, when the GA is ON, 20% of the current population undergoes

Table 7.4: Possible Values of Parameters for a Given Simulation

Parameter	**Tt1**	**Tt2**	**Tt3**
GA period	200	-1	1000
$P_{\#}$	0.5	0.3	0.9
Automatch Flag	ON	ON	OFF
No. of classifiers	100	100	100
Specificity Flag	ON	OFF	ON
Noise Flag	ON	OFF	OFF
Bucket Brigade Flag	OFF	ON	ON
Reinforcement	5.5	0.2	1.0
Number of shares	0	100	2,000
Initial Cash given	10,000	10,000	10,000
Commission per trade	0.1%	0.1%	0.1%
Annual Interest Rate	8%	8%	8%
Seed	0.19	0.95	0.52

reproduction, crossover, mutation and replacement by crowding, the remaining 80% stays unchanged. This means, for example, that the population of rules with a GA invoked every 100 cycles is turned over 4.55 times in 2,275 iterations.[2]

Automatch flag refers to a covering algorithm similar to the one suggested in [Wilson & Goldberg 89], that creates a matching classifier when there is no current match and inserts it in the population replacing the classifier with the lowest performance. There are differences between covering and *automatch*. In *automatch* the action is selected according to price behaviour (i.e. if the price today is higher than yesterday, the action is set to 1) and the proportion is selected randomly. Also, the classifier's strength is chosen to be the larger of 10 and the average value of the population's strength. Covering selects an action randomly, but experiments suggested that applying some type of heuristics when inserting the action is beneficial.

[2] $(0.2)(100)(2{,}275)/100 = 455$ new offspring

Additional experiments showed that using the initial default value of $strength = 10$ in all cases, did not improve performance; on several occasions, a non-matching situation was encountered when the average of strengths was well above the default value of 10 and the inserted classifier became one with very low strength and ended up disappearing, either because other non-matching classifier was inserted, or because the genetic algorithm would select it for replacement, forcing the system to make unnecessary re-insertions of the same classifier later on, when the same environmental message appeared again. In the majority of experimental runs, the use of this automatch algorithm showed improvements.

Note that the *number of classifier rules* per agent is a fixed quantity that the user chooses (e.g. 100), not a dynamic list because we want to develop the simplest possible type of agent first, one able to generalise good enough, as well as to specialise by recognising more narrow market states. If we simply keep adding to the population of rules new market states, we would end up with a very large set of specific rules and the complexity of the exploration cold be very large. Different numbers of classifier rules, such as 200 and 50, were also experimented. 50 showed the lowest performance, but 100 and 200 were very similar in performances. 200 classifiers did not show any improvements, which suggests that the number of possible market states can be summarised in 100. Therefore, most runs reported here, were performed with 100 classifiers per agent.

The flags *specificity*, *noise* and *bucket brigade* refer to the apportionment of credit subsystem, where an auction takes place to select a winning classifier from the set of matched classifiers, as described in [Goldberg 89]. When the bucket brigade flag is on, a payment equal to the current winner's *bid* (amount proportional to the classifier's strength and specificity) is given to the previous winner. This corresponds to the *implicit bucket brigade* algorithm first developed by S. Wilson in [Wilson 85]. In general, results show higher trader's performance when this flag is on. This is due to the fact that the payment received by the old winner balances the bid payments (and others such as taxes) it made; otherwise it could lose strength when its action might have actually contributed to more wealth.

And finally, *reinforcement* represents the amount of payment awarded to those classifiers that satisfy specific market criteria. The reinforcement scheme corresponds to a modification of Holland and Reitman's epochal credit allocation plan scheme, the *profit sharing plan* (PSP) [Holland & Reitman 78, Grefenstette 88], where a constant fraction of the current reward is paid to each classifier that becomes active since the last receipt of reward. Instead of paying a constant fraction, classifiers are rewarded according to their *specificity*: a classifier's parameter which determines the number of non # symbols in its condition part. The more specific the classifier is, the higher the percentage of the reinforcement reward it gets. Rationing the reward payments this way allows the creation of a good distribution of low specificity rules to cover more general market scenarios with respect to more specific rules to cover the more specific cases.

In addition to this, classifiers are paid if and only if they satisfy certain conditions. There are a number of market characteristics to consider when designing the reward scheme. Many different criteria have been tried, i.e. the simplest criterion corresponds to paying a reward to all matched classifiers whose action (yesterday) was a decision to buy when today's price is 2% higher than yesterday's price or 2% higher than the last 6-week moving average. Another criterion rewards classifiers that suggested holding when the actual price fluctuates 5% above or below the last 6-week moving average. Note that the reward payment is made once the current price of the stock is known, i.e. rewarding the classifiers that matched on the previous day AND satisfied the given criterion.

The following parameters will be explained in more detail in the following sections due to their relevance in the trading process. These parameters stipulate the status at the beginning, on day one, of each agent. For example, in the example of parameters shown in table 7.4 Tt1 was given 0 *shares* at the beginning of the simulation, £10,000 of *initial cash*, it is charged a *commission* of 0.1% on every transaction it performs and was paid 8% of *annual interest rate* for all its money invested in the bank. There are many alternatives to these values. A trader can start with any number of shares, cash, or both. In the same way, the other parameters such as interest can vary according to what the user defines. Finally, *seed* refers to the value of the internal seed, which is used during the first random generator number routine call to initialise values.

To summarise, in experiments, a *trader type* starts with a certain capital and gradually learns how to trade a certain stock, keeping cash not invested in that stock in a risk-free fixed-interest bond, essentially a high-interest immediate-access bank account paying 8% annual interest compounded daily. Every trade involves paying a commission too, normally set to 0.1% of the value of the transaction. The initial rules are created at random, except that there is a high proportion of wild-cards involved. Rules that, with hindsight, contributed in some way to good trading decisions get rewarded and receive more attention whenever rules compete to make the next decision. Every so often a Genetic Algorithm (GA) is run in order to evolve the rules further; rules with a good track record are favoured in this process.

7.12 The Ecology of Rules

The agents' evolved strategies are perhaps the most interesting outcome of any complex adaptive system model, yet they have traditionally been the least important focus of attention of researchers, probably because of their intrinsic complexity. Such ecologies represent various interactions and non-linear dependencies of rules which are involved in competitive and cooperative relations with each other. In [Holland *et al.* 00], Rick L. Riolo points out some important challenges that will face those involved either in the design and applications of LCS, or those using them for modelling complex adaptive systems: "For instance, we shall again be faced with the emergence of parasites and free riders, which are ubiquitous in natural ecologies and other similar complex adaptive systems ... these and other unanticipated side effects will provide great challenges."

When modelling these types of systems, not all the rules are meant to be useful, nor they can all be equally useful at all times. This state of perfection would be not only unrealistic, but also impossible to achieve in any complex adaptive system. At the end of the day, there are always a number of rules that seem to serve no purpose. This can be due to a number of reasons: for example, when analysing rule sets, one can observe that some rules have not had the chance to be tested either, because they could be part of new offspring added to the population, or because they could have been added by the auto-match algorithm (similar to Wilson's covering) as a result of having no current match to the environmental state, but then that particular state may never had occurred again. Also, in this pool of strategies, there are always some which are not commonly used or others which are not good that have not had the chance to be removed by the GA and are still part of the agent's current beliefs.

The analysis of such evolved sets of rules seems easier to study and explain than those derived from other paradigms such as NN or GP. In NN models it is difficult to know what the neural network is actually encoding during training, and when performance starts to drop and it is time to retrain, it is difficult to know what the net is recalling and what is forgetting and therefore we might not even know if it is encoding the original structure at all. These issues were addressed in Section 3.2.3. In GP models the rules can become extremely large and complex, making it very difficult to understand them (see Section 10.1). This section is devoted to provide the reader with several ways of gaining understanding of the relevance of the behaviours learned. As these come in many flavours, an example to illustrate where they apply will also be given in the following sections.

7.12.1 Analysis at the Rule Level

As explained in chapter 3, classifier rules contain two parts: $< condition >:< action >$. The condition is a bit string matching the current market state vector which has a fixed length with each position taken any of the three values 0, 1 or #. The 1 and 0 match corresponding bits in the state vector, while # is a wild-card symbol which matches either of the conditions and allows traders to dynamically decide which pieces of information are relevant to them by ignoring some state variables if they wish to do so. This symbol is responsible for providing the trader with an information selection property. Now, let's consider rule number 26 of a set derived from an experiment: **1001#01:1101.** This rule, of fixed length 7, evolved from trader-type 1 and indicates to buy 75% of all cash available whenever the price of today is:

1. higher than yesterdays's, and
2. not higher than 120% of the moving average of the previous week and

3. not higher than 110% of the moving average of the previous two weeks and

4. higher than 105% of the moving average of the previous four weeks and

5. it does not matter if it is higher than 102.5% of the moving average of the previous six weeks and

6. is not the highest price ever and

7. is the lowest price ever

The relevance of each rule of the evolved set can easily be studied by the following procedure, see figure 7.1: In chronological order, divide the total period analysed in two parts, a *training* set, composed of about 90% of the cases (year 1 to 9) and a *testing* set, with the remaining 10% (the 10th year). Initialise the training set with random strategies and run the simulation as normal, with both, the RL and the GA enabled in order to allow the agent to evolve its own new set of rules. After passing through the training period once, i.e. going from year 1 to year 9 of the market environment, stop and take **one** rule out along with its exact duplicates and run the test scenario where there should be no macro-adaptation involved (GA off), but only the micro-adaptation level is on (improvement of present rules, i.e. RL on). Measure and compare the performance of the trader during this testing period in two distinct cases, one with and one without the rule present in the set. Run the simulation without one of the rules each time, for a total of *N* times, the number of rules in the original set.

What can we learn by doing this? On one hand, in the case of decreasing performance by removing the rule, one can easily infer that the rule in question was indeed a relevant one (note that this assertion is only valid for the testing period analysed). But on the other hand, if performance improves without the use of that rule, the rule had no importance during that period. Let's emphasise what *during that period* means. Situations in the stock market are never guaranteed to reappear, and the successful use of a learned strategy in the near future only means that conditions have not changed dramatically. This would not be the case if the testing set was chosen to be in the year 2000 when most stocks collapsed. However, even in this case, if would be a good idea

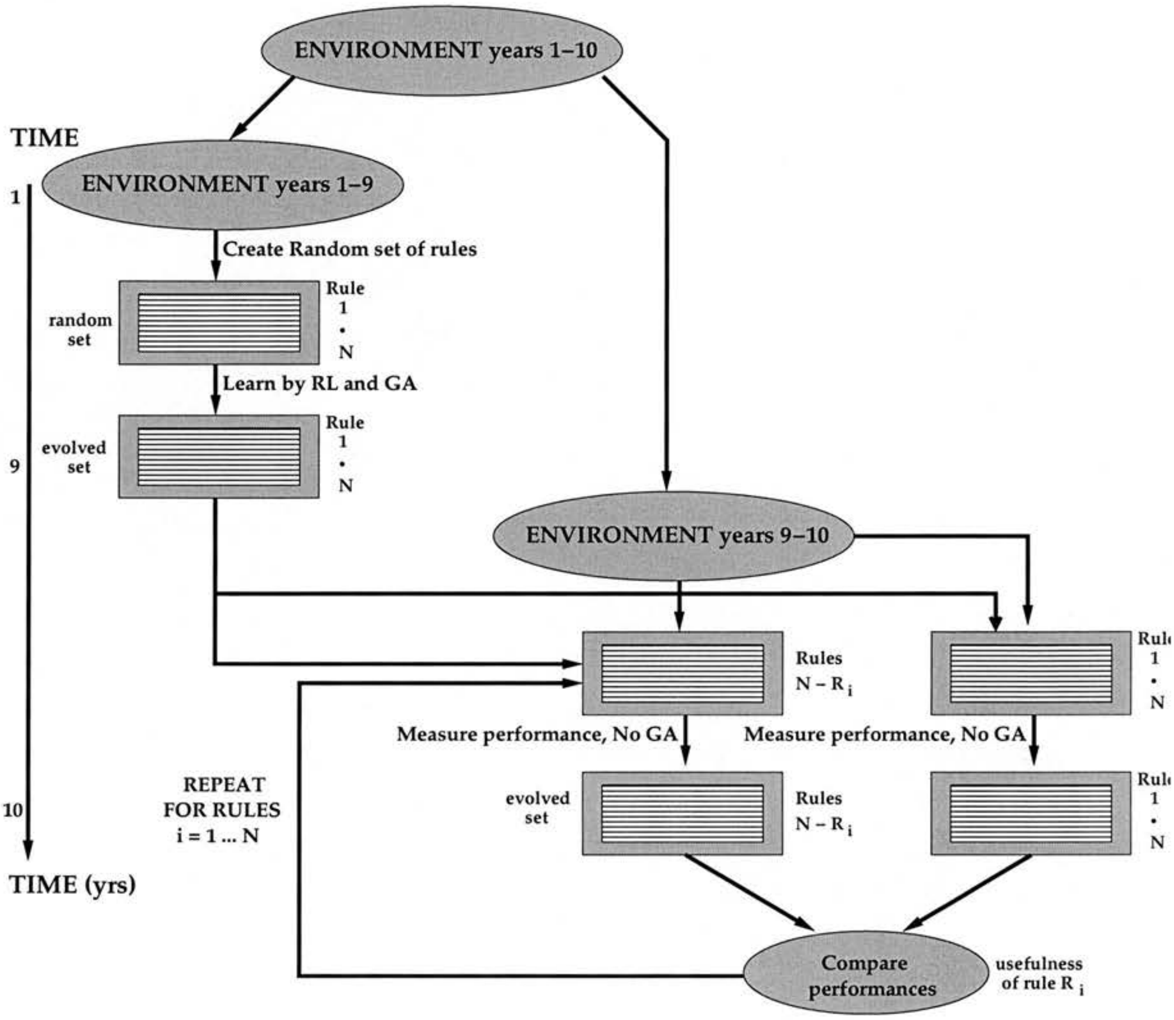

Figure 7.1: Testing the Relevance of Each Rule of the Evolved Set

to find out if the rule appeared in the continuous learning case, because perhaps it is a rule that was going to be out of the set anyway.

7.12.2 Analysis at the Bit-of-Information Level

One of the most relevant problems forecasting models face is that many indicators can be potentially relevant. "Determining a good set of indicators is crucial in generating good predictions as well as reducing training times to reasonable levels" [Barr & Mani 94]. This is perhaps one of the most crucial problems affecting NN-based systems reliability, and it is undoubtedly, a problem which is present in any system.

The aim of this section is to analyse this potential problem and how would a system of this type handle such constraints. There are a number of ways in which a sensitivity analysis might be performed. Typically, in NNs, the values of a certain variable are altered by a small percent (usually 5-10%) while keeping all the others fixed as they were, and the outputs after this change are then compared with the output using the original values. This is called a *weak sensitivity measure* because the whole idea is that there are non-linear relationships between the variables and therefore the change in one should affect the others as well. For instance, if one tests the relevance of variable 4, obtaining a negative change in performance of 25% and variable 8 showing a decrease of 10%, then if both variables are changed, the drop in performance cannot be added i.e, expecting a decrease of 35% in performance, it is not a linear relationship. The true result of altering variables is not easy to infer because of the non-linear interactions of such variables. On top of this difficulty, sensitivity analysis also vary over time, i.e., changing one variable over one period of time might not have the same effect of changing it over another period of time.

However, a *sensitivity analysis* of this type can give us an idea of which are the most influential indicators in a model, providing an alternate way of understanding the evolved sets of strategies (in a NN it would be a way of understanding a trained net and which inputs are emphasised, as reported in [Barr & Mani 94]).

Let's now consider bit number 6 of rule number 26 from the previous section: **1001#01:1101.**

This bit of information refers to the current price *not* being the highest ever. But is this piece of information important at all? How do we know if it is so? One might want to discard it once and for all if it does not look too important via a quick inspection (the price can not be the highest ever too many times, right?), but there exists the anxiety that getting rid of it might be the wrong decision to make. How can we test the relevance of such bits of information? Two ways of doing so will mentioned here:

First, for every trader-type, there is a different ecology of strategies that is dependent upon the stock being analysed. This is perhaps one of the most valuable lessons learned from using LCS: one can learn about the complexities of the environment being modelled. One can easily analyse, for example, taking the case of a bad performer such as Forest Oil, whose *buy-and-hold* performs worse than the *bank* investment. This stock has certainly very few cases matching the 6th bit: today's price is the highest price ever. At any given point of the period analysed, the rules evolved can be observed to assess their relevance.

To illustrate this point, table 7.5 shows an arbitrary segment of Tt1's rule set, evolved from a run of Forest Oil. As it can be seen, bit number 6 in most cases indicates a 0, which means that P_t is not higher than $P_{highest}$. The price of this stock is dropping most of the times, so Tt1 very quickly learns it. This is an important property of the stock being analysed and it might be important to keep this bit of information, but if another type of information appears to be more important for the designer, this bit could be a good candidate to be replaced. This set of rules also reflects important information about the stock. It is clear that most values are set to 0. This is revealing that the current price of the stock is, most times, *not* higher than the previous price, than the moving averages of various days, etc.

The second way to asses whether one bit of information is relevant or not is to discard that bit completely. For example, after analysing the rule set evolved by Tt1 with Forest Oil, one might just create an alternative Tt1 that does not have bit 6 in the condition part of its rules and evolve it in the absence of that piece of information. Then, both performances can be measured in order to assess the importance of that particular piece of information. This is similar to the *sensitivity analysis* mentioned at the beginning of this section, except that instead of changing the value by a small percent, this approach deletes the entire variable, in this case, the one describing whether the price of today is the highest ever. This will be exemplified in Chapter 8.

7.13 Model Description as a Truly Artificial Market

So far, for the reasons given throughout this thesis, the proposed model (fully described in this chapter) uses real data, i.e. the agents do not affect the stock's price. However, it is important to describe the structure of the model if it was a *truly artificial stock*

Table 7.5: An Arbitrary Segment of Forest Oil Set of Evolved Rules

Condition			Action
Bits 1-5	Bit 6	Bit 7	
00#00	0	0	1001
0#0#0	0	0	1010
#111#	0	0	1101
#111#	0	1	1101
#0001	0	0	1101
#1110	0	0	1101
101##	1	1	1101
#1100	0	0	1101
10001	0	0	1101
10#00	#	0	1111
10#00	0	1	1111
00101	0	0	0100
#0##0	0	#	0101
00##0	0	0	0111
00001	0	0	0111
10000	0	1	0010
#000#	1	#	0011
10#0#	0	1	0101
1000#	0	0	0101
1000#	1	1	0101
10001	#	0	0101
10011	1	1	0101

market. How would the stock's price be generated? How would the agents be described in this scenario? This section explores the inclusion of a different type of agent in ATOF World, one which has been extensively described and explored in the literature, and a highly inspirational piece of work that has motivated the development of this model in the first place: the *SFI stock market agent.*

This type of agent is economic in nature and it has been proven to behave in the theoretically predicted rational expectations behaviour (under slow learning), as well as to present a more realistic market behaviour (under fast learning). The basics of this agent type will be briefly described. The following readings are highly recommended for more accurate descriptions [Arthur *et al.* 96], [LeBaron *et al.* 99], [Palmer *et al.* 94], [Palmer *et al.* 99], among others cited throughout this Chapter. For more recent experiments of the SFI stock market regarding the effect of changing the system's evolutionary learning rate, etc., see [Joshi & Bedau 98], [Joshi *et al.* 98], [Joshi *et al.* 99].

This section is devoted to the development of a truly **artificial stock market scenario** in which the stock price is created endogenously according to a demand function. Because **all the agents** will actually affect the price in this scenario, not only the agent's behaviour, but also a new type of market behaviour –derived from their actions– can be analysed. The overall design of the artificial model proposed in this thesis and the SFI artificial stock market and agent rationale will be addressed, but a detailed quantitative analysis of results of this approach is outside the scope of this thesis. This is only a description of how the structure of the model with this agent would be; the actual SFI economic agent and the endogenous price formation mechanism have not been implemented. The idea here is that the original Tt1 will be replaced by the SFI agent. The other agents, Tt2 and Tt3, will be the same as they were in ATOF World. So the market still has three agent-types, one is the SFI agent and the other two are the same as they were in ATOF World.

7.13.1 The Arbitrary Agents

The pre-defined Tt2 and Tt3 agents are called, for simplicity and because they use real data in the original model, *arbitrary types*. Recall that in these agents, the decision rule represents (i) a specific action to take on the current period: buy, sell or hold, and (ii) the proportion to trade. If either of these agents decides to buy, it must also decide "how much" of its current cash available is going to spend in shares; or vice versa, if the decision is to sell, it must specify "how many" shares from its current holdings is going to sell on that day. This is exactly as they were in the original model. However, all of these mechanisms will be different for the new Tt1. This agent is explained in the following section.

7.13.2 The Economic Agent

The Economic Agent (hereafter Tt1) is myopic one period and has a constant risk aversion (CARA), expected utility of the form:

$$U(W_{t+1}^i) = -e^{-\gamma W_{t+1}^i}, \tag{7.2}$$

where γ is the coefficient of absolute risk aversion $\gamma \in [0,1]$ and W_{t+1}^i is the wealth of agent i at time $t+1$.

The new price is calculated endogenously in this scenario; only dividends are exogenous to the system, and are calculated by following the stochastic process:

$$d_t = \bar{d} + \rho(d_{t-1} - \bar{d}) + \varepsilon_t \tag{7.3}$$

ε_t is Gaussian, independent and identically distributed noise, the autocorrelation factor $\rho = 0.95$ and $\bar{d}$ is simply the moving average of the dividends over the period.

The expected or predicted price and dividend P_{t+1} is a linear combination of the parameters a and b, which will be calculated by the classifier system to maximise the utility function:

$$P_{t+1} = a(p_t + d_t) + b, \tag{7.4}$$

where $P_{t+1} = p_{t+1} + d_{t+1}$.

The mean price $\bar{Q}_t$ is

$$\bar{Q}_t = \frac{\sum_{i=1}^{t}(p_i + d_i)}{t} \tag{7.5}$$

The variance of $(p_t + d_t)$ is calculated by:

$$\sigma_t^2 = \frac{\sum_{i=1}^{t}(p_i + d_i)^2}{t} - \bar{Q}_t^2, \tag{7.6}$$

If the variance is slightly negative, consider $\sigma_t^2 = 0$

The demand for holding shares (percentage of current wealth that the trader wants to have in stocks) H_t is given by:

$$H_t = \frac{P_{t+1} - p_t(1+r)}{\gamma\sigma_t^2}, \tag{7.7}$$

where p_t is the price of the risky asset at t, $\gamma\sigma_t^2$ is the conditional variance of $p+d$ at time t, for agent i; γ is the coefficient of risk aversion.

The number of shares S_t the trader wants to hold/have must be calculated because H_t refers to the whole trader's wealth, which is its holdings (stock valued at current price) plus its cash available:

$$S_t = \frac{H_t W_t}{p_t} \tag{7.8}$$

Recall that W_t is the total wealth of the agent, including its current holdings and cash available, so we need to compute the net amount to trade T_t, which is the difference between the number of shares the agent wishes to hold/have S_t and the number of shares owned, s_t:

$$T_t = S_t - s_t \tag{7.9}$$

The sign of T_t indicates whether the agent is bidding to buy or offering to sell the number of shares specified by its magnitude. If the magnitude of T_t is zero, the agent wishes to hold his possessions as they are.

$$T_t = \begin{cases} > 0; & b_t = T_t, \\ < 0; & o_t = T_t, \\ = 0; & b_t = 0 \text{ and } o_t = 0. \end{cases} \tag{7.10}$$

For each t, the bid to buy the stock is denoted by b_t and the offer to sell is denoted by o_t.

B_t is the total number of bids to buy the stock, and similarly, O_t are the total number of offers to sell the stock.

$$B_t = \sum_{i=1}^{N} b_i(t) \tag{7.11}$$

$$O_t = \sum_{i=1}^{N} o_i(t) \tag{7.12}$$

With the following equation we adjust the future price:

$$P_{t+1} = p_t \left[1 + \eta \left(\frac{B_t + 1}{O_t + 1} - 1\right)\right], \tag{7.13}$$

$\eta \in [0, 1]$. Use a small η for very small adjustments to prices and large, such as η = 0.95 for large oscillations.

7.14 ATOF World vs The SFI Model

Although there are other models of markets, this section describes some of the differences between this model and the SFI model, summarised as follows:

- The main difference between the two models in market structure is that the price is created endogenously in the SFI model. Therefore the dynamics of prices and the **price formation mechanisms** are completely different.

- The **goal of the artificial agents** in later versions of the SFI model [LeBaron 95, Arthur *et al.* 96, LeBaron *et al.* 99] is to build forecasts of the future price and dividend according to a utility function. This corresponds to a condition-forecast type of classifier rule. In an earlier version [Arthur 94a], the rules were of the type condition-action, where the action on each period was a binary choice to bid (1) or to offer one share (-1). If no rules were activated on that period, it was considered that the agent made neither a bid nor an offer. In ATOF world, the action is a two-part investment signal dictating on the first part, whether to buy, sell, or do nothing on the current period, and in the second part, stipulating what proportion to trade. Holding is one of the agent's valid outcomes (when the proportion is 0) and in fact, it is considered so important that agents indeed receive a positive feedback from the environment in cases in which the decision made was to hold and the fluctuations of price or volume were not very significant i.e. within +/-5% of the current value. This means that the agent will hold when he expects a small change in prices.

- The **heterogeneity of traders** of the SFI model resides only in the divergence of rule sets, which means that the agents have different operating principles, i.e. agents become heterogeneous by using different rules. Apart from that, traders are identical [LeBaron *et al.* 99]. All classifier rules are of the same length (i.e. 60 rules with the condition part being as long as 70-80 bits in earlier versions of the model. Also, agents have the same number of such rules. In ATOF World, agent heterogeneity takes a more rigorous role in the sense that there is a wider variety of psychologically plausible behaviours. Agents' heterogeneity is dictated by the fact that rule sets are different from each other in meaning, size, and number. For example, Tt2 might have 50 condition-action strategies with the condition part of the rules of being of length 7, Tt3 30 strategies of length 5 and Tt1 having 80 classifiers, with 10 bit long conditions each. The effect of this is that an agent might decide to sell because of negative economical news (although not formally represented in ATOF World due to limitations in data) while a more objective and confident trader might decide to buy because it takes into consideration other facts such as price or volume ratios. It would follow what the stronger strategy in its current set of classifier rules dictates.

- In the SFI model there is a fixed **number of agents** in the market, equal to the total number of shares in the market. The number of agents goes from about 25 to 100. In ATOF World there are only three agent types, and each one stands for the aggregate effect of all traders of its type. Therefore one type might be more successful than others, depending on the stock being analysed. It is a much simpler version in the sense that when running the simulation only once, i.e. with one seed, each one of these three types of agents is representing only one agent, but each instance can represent one or any other number of agents of its type, like 100 or 1,000. The number of types can also be modified. Here we only use three types, but other types can be easily developed and added.

- Other components not present in SFI model are the *buy-and-hold*, *bank* and *trend-following* strategies, as well as the commission, interest and reward schemes.

7.15 Analysis of Agent Behaviour due to Changes in Evolutionary Learning

The purpose of this section is to address an important issue regarding the analysis of emergent behaviours due to the effects of different types of learning. As it has been explained in sections 7.10.2 and 7.11, the evolutionary rate in the model is controlled by the *ga-period* parameter. The model uses a steady state GA which is invoked deterministically every *ga-period* generations (where one generation is equivalent to one trading day). There are also other methods that could be used, such as invoking the GA stochastically (called probabilistically with average period *ga-period*), or conditioning it to events such as lack of matches or poor performance.

In the Santa Fe artificial stock market, the prices of the stock are created endogenously according to supply and demand rules. Arthur et. al. [Arthur *et al.* 96, LeBaron 00, Palmer *et al.* 94] showed that varying the rate at which individual agents learn new investment strategies reveals two different kinds of overall market behaviour, i.e. the stock market converges to one of two attractors, one corresponding to the equilibria posited in conventional theory and the other exhibiting the more volatile qualities observed in the real stock market. In other words, if investment strategies evolve slowly (when individual agents have slow rates of learning about conditions affecting the market), the market showed behaviour generally consistent with the prediction of traditional economic theory -conventional behaviour. But if the strategies were allowed to evolve more quickly, the market showed the kind of instabilities and statistical properties typically observed in real-world markets -the artificial stock market behaves in the more volatile manner.

More recently, in the same SFI model, Joshi and Bedau [Joshi & Bedau 98] have analysed and attempted to explain the market and agent behaviour exhibited by varying the rate of evolutionary learning. Their results show that when the GA interval is moderately low, rule complexity (measured by number of non '#' symbols in the rules) is very high, the agents' wealth is low and there is significant technical trading. When increasing the GA interval, the complexity in forecasting rules decreases, allowing an increase in wealth earned by the agents and technical trading decreases as well.

Schulenburg and Ross [Schulenburg & Ross 99, Schulenburg & Ross 00] have shown that increasing the GA invocation is proportional to accumulation of wealth up to a certain point (highest returns obtained with *ga period* ranging from 50-100 depending on the agent type), after which wealth starts to decrease. Because the model uses daily real data, there is a limitation in the number of time periods that can used to approximately 2,500 generations (ten years of trading), which is much smaller compared to the SFI model, where the simulations performed can run for up to 300,000 periods.

Chapter 8

Experimental Results

"Why does man not see things? He is himself standing in the way: he conceals things." Friedrich Nietzsche, *Daybreak*, translated by R.J. Hollingdale.

8.1 A Brief Review of the Model

This chapter reports on a number of experiments using the model fully described in Chapter 7, where the three different groups of artificial agents learn, forecast and trade their holdings in a real stock market scenario given exogenously, in the form of easily-obtained stock statistics such as various price moving averages, first difference in prices, volume ratios, etc., as well as other non-technical factors.

These artificial agent-types trade while learning during –in most cases– a ten year period. They normally start at the beginning of the year 1990 with a fixed initial wealth to trade over two assets (a bond and a stock) and end in the second half of the year 2000. The agents' learning process is represented with Learning Classifier Systems (LCSs), that is, as sets of bit-encoded rules. Each condition bit expresses the truth or falsehood of a certain real market condition. The actual conditions used differ between agents. In this chapter, the forecasting performance of the adaptive agents will be compared against the performance of the *buy-and-hold* strategy, a *trend-following* strategy and finally against the *bank* investment over the same period of time at a fixed compound interest rate. To make the experiments as real as possible, agents pay commissions on every trade. The question to address here is whether the agents are capable of performing in sensible ways in a **real stock market** environment which will be given to them.

8.1.1 Reviewing the Strategies

The baseline strategies are:

1. *Buy-and-hold*: simply put all available cash into the stock at the start and then keep it there, and

2. *Bank*: merely keep all money in the bank at a good rate of interest, never buying the stock. Note that the *buy-and-hold* strategy can be a very good one if the stock performs well over the years. *Bank* strategy can be very good if the stock performs poorly over the years.

3. *Trend-following* strategy: this strategy buys into a stock when the price starts to show an upward trend, and sells the stock when there is potential evidence that the price will continue to fall.

The *trend-following* strategy is simple and yet will outperform both the *buy-and-hold* strategy and the *bank* strategy if the stock price does not show a clear long-term trend upwards or downwards. It assumes that there will be an uptrend if it sees that the price of the stock today is higher than yesterday, and therefore buys the total number of shares that its available cash allows, minus the commission. If today's price is lower than yesterday's it assumes a continuation of that downward trend and sells all shares in possession. In both cases no other transaction is made until there is a reversal in trend. It pays commission in the same way the adaptive agents do. Note that when this strategy sells its stock and puts its money in the bank, the commission it pays on the transaction is higher than three days of interest earned on the bank deposit. On any given day, its holdings are either all shares valued at the current price, or all cash in the bank and earning interest. The frequency of transactions varies from stock to stock with this strategy, and as it will be shown in the next section, its performance is quite unpredictable and unreliable; it can be the best or the worst, depending on the stock.

8.1.2 Interest and Commission

To make the model as real as possible, the model contemplates an annual fixed interest rate of 8% to all cash invested in the bank. This means that all the available cash the agents have is invested in the bank (with instant access) at this rate and it is compounded daily, meaning that at the end of each day the agents get their interest payment. Although it might seem a bit high in the year 2000, a decade ago interest rates were higher than at present times. The possibility of making the interest variable rather than leaving it fixed for such a long period of time was considered, but again, there are limitations obtaining the data needed for the experiments. It would be needed, for example, quarterly or monthly interest rates for the whole past decade for the different countries these stocks trade in.

The commission chosen to be charged is 0.1% of the total amount to trade. Commission costs are much lower than they used to be. Although it might seem low at first sight, which seems to be a realistic figure. To make the point clear, two examples will be illustrated. First, consider that on March 31st 2000 a trader decided to spend 75% of its cash available of MSFT shares (a bad decision, the price dropped from $106.25 to $90.875 in one day!), having only $36,000 of available cash, the commission charged at 0.1% would be $27.00. Nowadays most executions on the Internet charge less than $15.00 per transaction. These low fees normally include all types of orders - Market, Limit, and Stop - up to 5,000 shares traded online with no hidden costs. Total wealth and amount to trade vary from agent to agent. As a second example, imagine another agent has $1,000,000 worth of MSFT shares and decides to sell 64% of them, then the commission it pays is $640, which is extremely high! In cases such as these the commissions might average out, but another possibility is to charge a flat $15.00 (or £15.00 for UK stocks) per trade, no matter what size.

In the future the idea is to experiment with models that use flat commissions and a lower bank interest rate. No dramatic differences are expected in the results. Another factor is that the model uses the readily available **daily closing prices** rather than the actual bid and ask price (again these data was not obtainable), but the spread might be advantageous for the traders. For this reason, the commission could increase to 2%, a figure which should include more than the bid/ask difference (normally 1/16) plus commission. Note that using closing prices is less than ideal, since nobody can actually trade at those prices.

8.2 Previous Work

First, in order to explore the degree of reliability that artificially intelligent agents can have when applied to real life economic problems, an initial evaluation was conducted to test whether an LCS is able to represent competent traders in a **real market scenario** in which daily stock prices, volume of transactions, dividends and information about other simple investment strategies are given to the agents **exogenously**, allowing full concentration on the dynamics and evolution of the behaviour of these evolving traders without having to be concerned about how their actions affect the market. This initial investigation was reported in [Schulenburg & Ross 99, Schulenburg & Ross 00, Schulenburg & Ross 01]. Further explorations were reported in [Schulenburg & Ross 02].

The results in the earlier papers, using the very well performing stock Merck and Co. over a period of ten years, showed that the artificial agents, by displaying different and rich behaviours evolved over time, were indeed able to discover and refine novel and successful sets of market strategies that could outperform baseline strategies such as (i) *buy-and-hold*: simply put all available cash into the stock at the start and then keep it there, and (ii) *bank*: merely keep all money in the bank at a good rate of interest, never buying the stock. Note that the *buy-and-hold* strategy can be a very good one if the stock performs well over the years, and *bank* strategy can be very good if the stock performs poorly over the years.

Arguably, Merck and Co. is not a representative example of a stock (no stock would be), nor is it an unequivocally impressive illustration of the approach. Figure 8.1 shows how it performed over the years; in that figure, the lowest of the jagged lines shows how the value (in US dollars) of the *buy-and-hold* strategy grew over the years –the horizontal axis represents trading days. Apart from a fall over a period of about one and a half years in the middle, the stock's value grew fairly steadily, and at a very healthy rate in the most recent years. It was therefore fairly easy to make money from it. As the figure also shows, the best examples of the three *trader types* each managed to outperform the *buy-and-hold* strategy by a healthy margin. Although the fluctuations in their wealth were similar to the fluctuations in the stock's value (as shown by the wealth of the *buy-and-hold* strategy), they were able to do better by transferring money from the stock to the bank during downturns and transferring it back in upswings, as a close examination of the sequence of decisions revealed. The *bank* strategy of course did poorly; its wealth is shown by the lowest, smoothly rising line in figure 8.1.

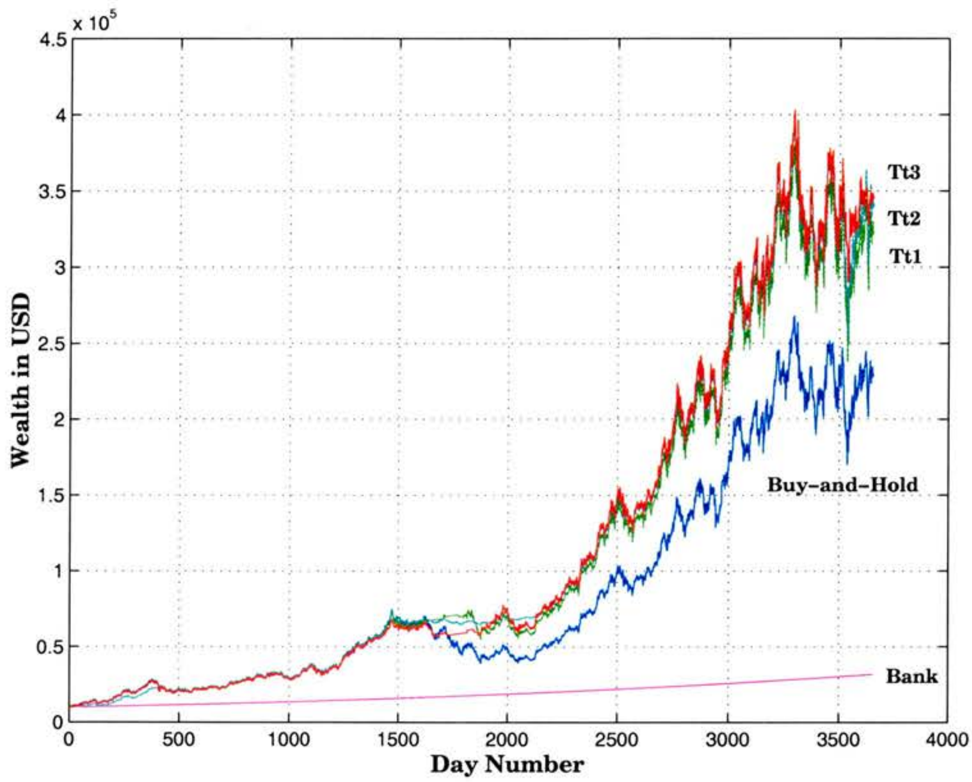

Figure 8.1: Merck and Co. Stock. Wealth (in US dollars) of *Bank* Investment, *Buy-and-Hold* Strategy and *Trader-Types* 1,2 and 3 over a period of 11 years

In order to test if the approach works, that is, whether the best of the sets of rules evolved perform more profitably than other common strategies, the experiment typically uses ten years of historical data. The simulated traders apportion their money between holdings in a reliable high-interest bond (8%pa compound) and holdings of the stock in question, and they are obliged to pay a realistic commission on every trade

(0.1% of transaction value). The best of them seriously outperform strategies such as keeping money in the high-interest bond, or buying then holding at the start of a prolonged general growth. The models are cheap to evolve, taking one run only seconds on a reasonable PC; not all models are great performers, but the cost of finding good performing models is still low.

8.3 Describing the Market Environment

Although one single stock can capture some features of real markets, more investigation is needed across different stocks to analyse agent behaviour under a full range of market phenomena of interest, such as abrupt changes in trading volume, bubbles, crashes, market psychology and moods. For this reason, results of simulations of a selection of stocks representative of two potentially different market environments (the UK and the US) will be presented in this chapter; five of which trade in the London Stock Exchange and three in the New York Stock Exchange.

Table 8.1 describes general characteristics of the stocks analysed such as the company name, the industry sector they belong to and their index membership.

Table 8.1: Profile of Stocks Analysed

COMPANY NAME	SECTOR	INDEX MEMBERSHIP
BP Amoco Plc.	Oil and Gas	FTSE 100
GKN Plc.	Automobiles	FTSE 100
Hanson Plc.	Construction	FTSE Mid 250
Lloyds TSB Group Plc.	Banking	FTSE 100
WPP Group Plc.	Media	FTSE 100
Microsoft Corp.	Tech., Software	Nasdaq 100, S&P 500, Dow Ind.
Cabletron Systems, Inc.	Tech., Comp. Networks	S&P 500
Forest Oil, Corp.	Oil and Gas	N/A

Table 8.2 is a continuation that refers to their symbol and the actual dates used in the experiments described in this chapter. The length of the series is not the same in all stocks due to the fact that availability of both, prices and volumes is needed for the agents to process. There are two stocks with shorter series: Lloyds TSB Group (series of length 1,659) and WPP Group (length 1,752), the others are longer than 2,500 trading days. Unfortunately, Lloyds' price series after December 29th, 1995 was unavailable, as well as WPP's volume series up to June 30th, 1993. Note that one calendar year (365 days) corresponds roughly to only 253 financial or trading days.

As it will be observed later on, these stocks include good performers that have generally grown over the period and others that have pretty steadily fallen. From now on, they will be addressed by their symbol, as shown in Table 8.2, and their properties and trends by day numbers rather than the actual dates when they occurred.

Table 8.2: Dates of Stocks Analysed Corresponding to Table 8.1

SYMBOL	FROM	TO	NO. OF DAYS
BP.L	06/19/89	06/07/00	2,780
GKN.L	06/19/89	06/07/00	2,780
HNS.L	06/19/89	06/07/00	2,780
LLOY.L	06/19/89	12/27/95	1,659
WPP.L	07/01/93	06/07/00	1,752
NASDAQNM: MSFT	13/03/86	25/08/00	3,654
NYSE: CS	31/05/89	28/04/00	2,757
NYSE: FST	26/03/90	13/03/00	2,513

8.4 Describing the Layout of Results

It is important to recall that the results of the following sections represent the performance of the best traders of a number of runs. This is an on-line learning model where the agents start trading in a real environment at day one with initially completely random strategies. Every new run is independent of previous ones and starts with a set of random strategies, with each agent initially having 10,000 dollars or pounds depending on whether the stock in question is a US or UK one. The random number seed is changed for each run; running the model with the same seed value would give exactly the same results. As described in Chapter 7, with probability 0.5 any bit-position in an initial rule's condition will be a #, so that the initial rule-set is reasonably general; 0s and 1s each occur with probability 0.25. However, random strategies are not uniformly bad. For a given run, the initial strategies that one agent gets can be genuinely better than what another agent gets.

Look at figure 8.2, which describes experiments with BP Amoco stock. The left-hand table shows the final wealth of various strategies: the *bank*, *buy-and-hold* and *trend-following* strategy; in this case the *trend-following* strategy is the most successful of these three, accumulating £74,943. The next table shows the wealth of the best-performing Tt1, Tt2 and Tt3 agents, for various values of how often the GA is run (every 50 days; every 75 days; …; every 200 days). For each of these learning rates, 1001 runs were performed and the figure in the Tt1 column gives the final wealth of the most successful Tt1 agent in those 1001 runs. For example, when the GA is run every 50 trading days, the most successful Tt1 agent made £69,294, the most successful Tt2 agent made £80,122 and the most successful Tt3 agent made £65,843.

The graph associated with the table shows the growth in wealth of *bank*, *buy-and-hold*, and the best instance of each of the three *trader types* Tt1, Tt2 and Tt3 –taken from the best results with the GAs shown on the left-hand table. Due to space limitations daily wealth of the *trend-following* strategy is not graphed, but as explained before, the final wealth of this strategy is shown on the left-hand table.

The horizontal axis shows trading days, the vertical axis shows wealth – in pounds for UK stocks and in dollars for US stocks. All start with an initial 10,000 pounds or dollars. The smoothly rising line represents the wealth of the *bank* strategy. The lowest of the jagged lines is the wealth of the *buy-and-hold* strategy, and therefore reflects the stock price over time. Over the last ten years the stock markets have generally shown an upward trend, so it makes sense to use the *buy-and-hold* strategy as a common baseline for assessment in judging how easy it is to produce a *trader type* that is at least reasonably effective at making money.

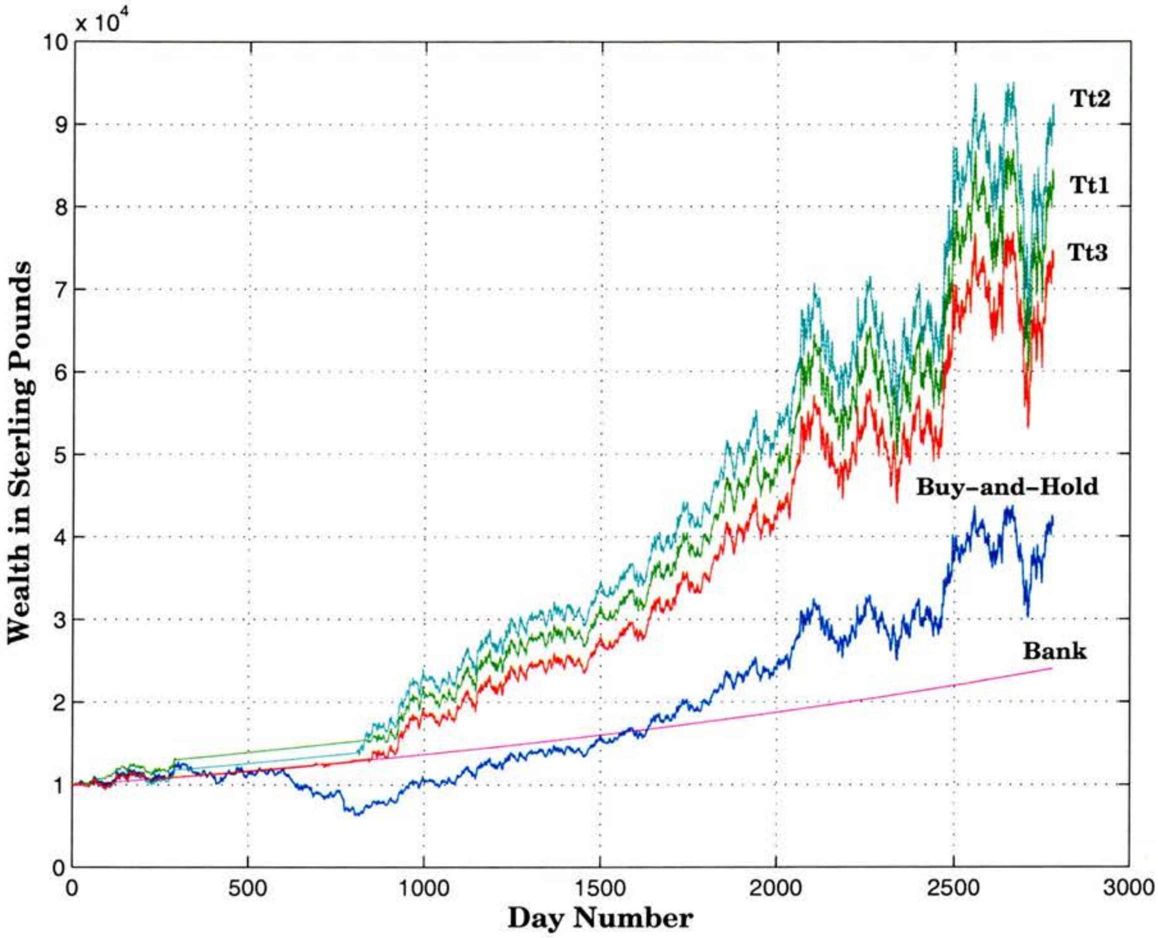

	GA	Tt1	Tt2	Tt3	Tt1	Tt2	Tt3
Bank	50	69,294	80,122	65,843	26%	18%	12%
24,107	75	80,166	79,876	68,464	32%	19%	18%
Buy-and-Hold	100	80,830	85,062	65,776	32%	19%	18%
41,248	125	81,964	83,640	64,689	32%	19%	17%
Trend-Following	150	75,942	86,959	66,430	32%	21%	20%
74,943	175	80,381	74,431	72,531	32%	20%	22%
	200	71,160	89,697	69,553	26%	21%	19%

Figure 8.2: UK: BP Amoco Stock. See text for explanation

The right-hand table shows the percentage of those 1001 runs in which each *trader type* was able to outperform the *buy-and-hold* strategy. For example, when the GA is run every 100 days, 32% of Tt1 strategies, 19% of Tt2 strategies and 18% of Tt3 strategies were able to beat *buy-and-hold* in terms of final wealth. These figures give at least some crude idea of how many times one might expect to have to run the LCS in order to arrive at an acceptable strategy by taking the best-performing *trader type*. And as the graph suggests, in general the *buy-and-hold* strategy is outclassed consistently from very early on; it is not the case that the *buy-and-hold* strategy leads for most of the time but fades near the end.

The graph conveys other interesting information too. For example, observe in this case that the price of BP Amoco stock showed a downward trend from around day 300 to around day 800. The *trader types* all managed to spot that it was best to keep their money in the bank during this period!

8.4.1 Is there a Better Trader?

It is noticeable that Tt2 sometimes outperforms the others by a large margin, for example in figure 8.2. Why is that? To understand, it is important to know what information the different *trader types* base their decisions on.

Trader type 1 uses price information: is today's price higher than yesterday's?, is it a new high?, is it a new low?, is it higher, by some small fixed percentage, than moving averages taken over periods ranging from 5 to 30 days? *Trader type* 2 pays attention to whether today's price is higher than yesterday's, and also some volume indicators such as whether today's volume is a new high or new low, or higher than yesterday's, or higher than the 20-day moving average. *Trader type* 3 pays attention to whether today's price and volume is higher than yesterday's, and also to how it compares to the bank and *buy-and-hold* strategies. It also knows what action it took yesterday, so it is crudely reflective.

Trader type 2 often benefits from knowing more about volume information than the others. Volume, when used in conjunction with other data, is a useful determinant in identifying whether a continuation of or a reversal in the prevailing trend is likely. Volume at low levels reflects uncertainty regarding the future direction of the market in question. If the volume is relatively high while the market is going up and remains relatively low during corrections, the inference is that the market is in a strong upward trend which should continue. In the other case, when the volume is high while the market is going down and relatively light during upward retracements, then the market is weak with a continuing downward trend more likely.

This does not necessarily mean that Tt2 is the best possible model for a *trader type*, of course. The whole question of what is a "good trader" is a complex one. As it was described in Chapter 5, in financial markets a good trader is not necessarily the person who makes the "right decision" most of the times. It could be someone who only makes three good decisions in a year, probably obtaining larger profits than another one who "got it right" most of the times, but with a smaller profit average. So as long as they keep earning a profit, they might stay in the market. But what about the others? Many others who report losses are also in the market. They are all survivors. Thus the market is composed of all sorts of traders (experienced and unexperienced, mathematical and gamblers, etc) each competing and hoping to get the most out of their investment at the expense of others. Some not only rely on their own tactics but might also copy what others, that they consider good ones, are doing. They imitate and possibly learn something by trial and error experimentation. In this pool of traders, some simply leave the system when they choose to do so, others are forced out (traders are kicked out of the market usually by punishment, negative criticisms, etc). Some can have a very short lifespan in the market place, others a long or even indefinite one. This approach favours the "fund manager" model in which the aim is to learn to make a steady stream of reasonably good decisions, rather than (say) a single major killing followed by exit.

8.5 UK Stocks

8.5.1 Case 1: BP Amoco Plc.

This has been partly discussed in section 8.4. Note again in figure 8.2 that Tt1 and Tt2 beat even the *trend-following* strategy, for most choices of GA-interval. It is also easy to find strategies that beat *buy-and-hold*, even though the stock has trended upward for most of the time. As mentioned earlier, the best strategies responded clearly and appropriately to the significant period when there was a pronounced slump in price.

8.5.2 Case 2: GKN Plc.

In the graph in figure 8.3, the *buy-and-hold* strategy is the lowest of the jagged lines. As the tables show, the *trend-following* strategy is hugely better than either *buy-and-hold* or *bank* strategies, because of the early gentle decline of prices and because the price fluctuated a lot towards the end of the period.

The best *trader types* managed to avoid the early decline, and all managed to do very much better than *buy-and-hold* but not as well as *trend-following*. It was reasonably easy to generate strategies that beat *buy-and-hold* by a healthy margin.

In highly volatile stocks such as GKN, the *trend-following* strategy performs much better than the other methods because there are a large number of well-defined trends during the period analysed, specially from day 2,000 onwards. Figure 8.4 shows – in addition to the information plotted in figure 8.3– the wealth of the *trend-following* strategy and it focuses only in the last part of the series, i.e. days 2,000 to 2,780. The arrow at around day 2,550 indicates a 100-day drop of the *trend-following*'s wealth: this strategy can have huge losses (far bigger than the losses of trader-types 1 & 3 and *buy-and-hold*) due to the highly risky operations of buying and selling so frequently. From the start of the arrow until the last day of the series, only Tt1 and Tt2 managed to make a profit.

During unstable periods such as this one, the traders usually manage to reduce their losses by keeping more money in the bank, by holding a small number of shares or by selling early in the downward trend. Figure 8.5 shows the shares owned by *buy-and-hold* (constant number of shares throughout), *trend-following* strategy (the lighter with more shares) and Tt2, which is the trader that can have a similar behaviour under highly volatile environments. To make the graph clear, it only covers the last 230 days of the series, i.e. from day 2,500 onwards. In this graph it can be seen that Tt2's frequency of transactions is smaller than the *trend-following* strategy's and that it usually holds a smaller number of shares, thus minimising its losses when bad times arise.

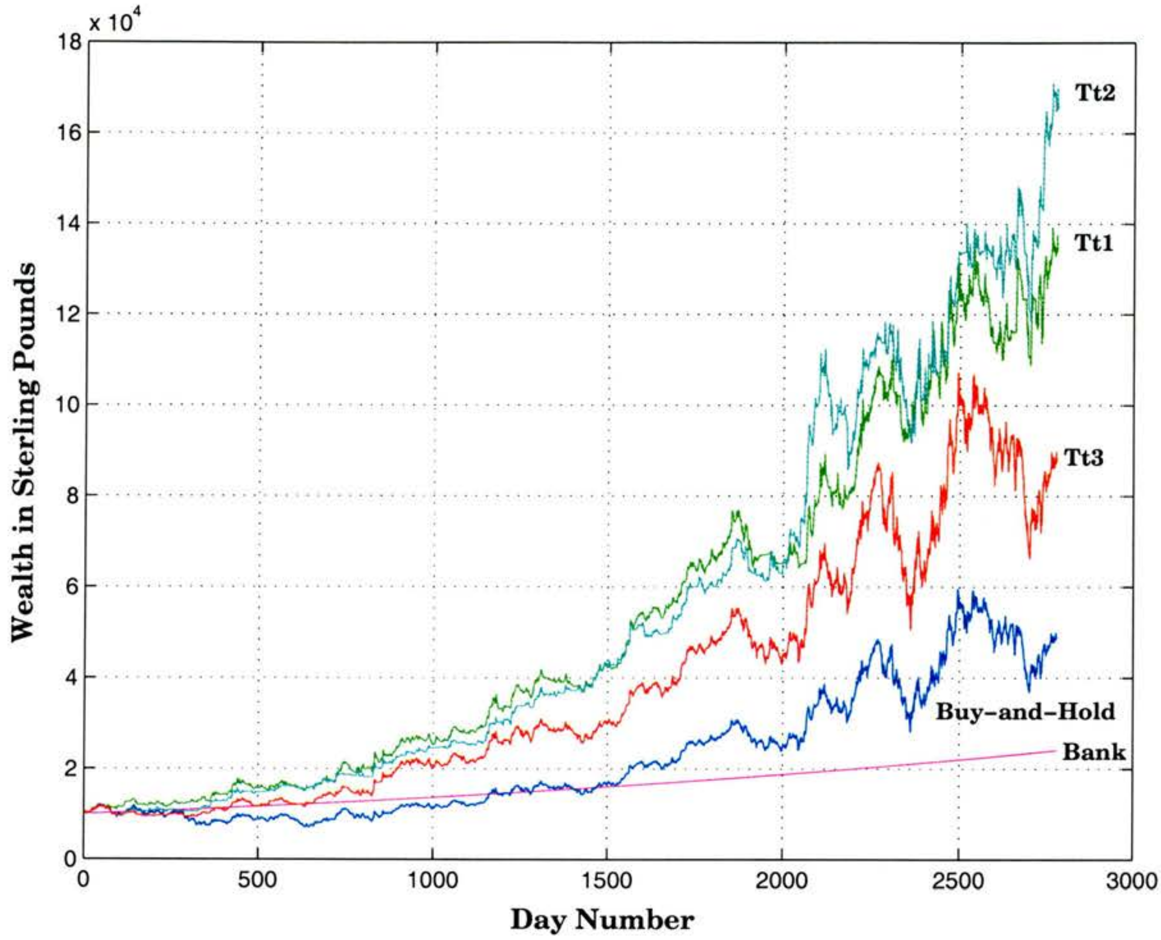

	GA	Tt1	Tt2	Tt3	Tt1	Tt2	Tt3
Bank	50	114,956	100,835	84,856	36%	26%	18%
24,107	75	91,232	106,680	81,436	38%	29%	19%
Buy-and-Hold	100	106,360	116,577	74,670	40%	28%	18%
49,072	125	108,628	105,056	83,179	35%	31%	18%
Trend-Following	150	98,655	112,793	88,491	32%	31%	19%
202,612	175	134,737	106,896	73,800	30%	29%	17%
	200	92,933	165,859	76,321	26%	30%	15%

Figure 8.3: UK: GKN Stock. See text for explanation

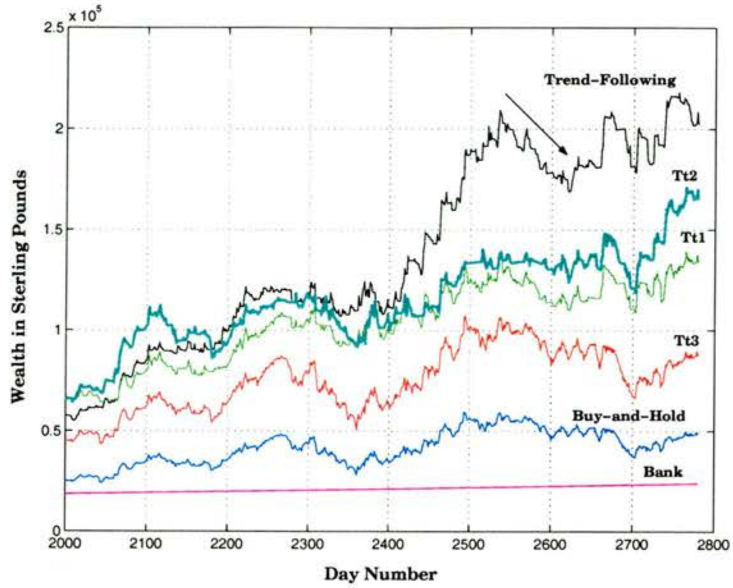

Figure 8.4: UK: GKN Stock. Wealth (in Sterling Pounds) of *Bank* Investment, *Buy-and-Hold*, *Trend-Following* Strategy and *Trader-Types* 1,2 and 3 from day 2,000

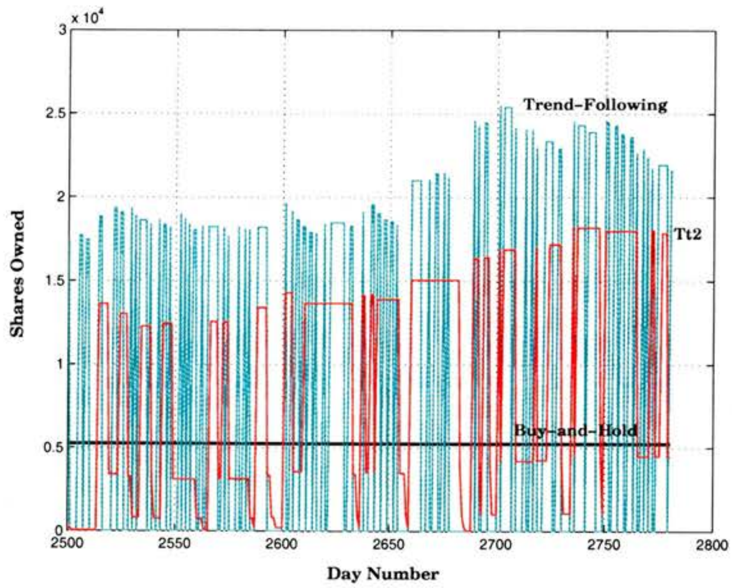

Figure 8.5: UK: GKN Stock. Shares Owned by *Buy-and-Hold*, *Trend-Following* Strategy and *Trader-Type* 2 from day 2,500

8.5.3 Case 3: Hanson Plc.

Figure 8.6 is interesting. Again, the *buy-and-hold* strategy is the poorly-performing one, the stock fluctuated a lot early on and then went into a pronounced decline before rallying and then fluctuating again. Tt3 did worse than the *bank* strategy for a while, but eventually learned to exploit the upswings to come out well. During the early decline, all the *trader types* rode out the slump by keeping money in the bank although Tt2 managed to creep ahead by successfully exploiting some of the upturns at that time. Overall the *trend-following* strategy did only slightly better than the *bank* strategy, and all the best *trader types* beat them both by a large margin. It was easy to generate strategies that beat both the *bank* and the *trend-following* strategies.

8.5.4 Case 4: Lloyds TSB Group Plc.

Lloyds stock performed well over the six and a half year period, with a fairly steady upward trend punctuated by minor falls. In the graph in figure 8.7 and 8.8, there is only one long decline in prices from around day 1,200 to day 1,500. As a result, *buy-and-hold* soundly beat its non-evolved rivals, growing the initial £10,000 to £37,321. However, the best traders all did even better, with Tt2 being the best performer. It was, however, quite hard to generate *trader types* that beat *buy-and-hold*; only 4-8% of runs did so. A possible explanation for this is that Lloyd's prices change direction too often and too quickly (i.e. trends do not last more than a day or two), making it for the traders really expensive to respond to these changes by trading, so they tend to ignore small price movements and follow more or less what *buy-and-hold* suggests after day 400 approximately.

8.5.5 Case 5: WPP Group Plc.

As it can be seen in figure 8.9, this stock performed very well over the period, with a modest slump during a short interval around day 1,300. As with the previous cases, the *buy-and-old* strategy is the lowest jagged line. The best trader types learned to gradually pull ahead of *buy-and-hold*, but only Tt2 was able to do much better than the *trend-following* strategy. Tt2 was also the easiest type when it came to generating

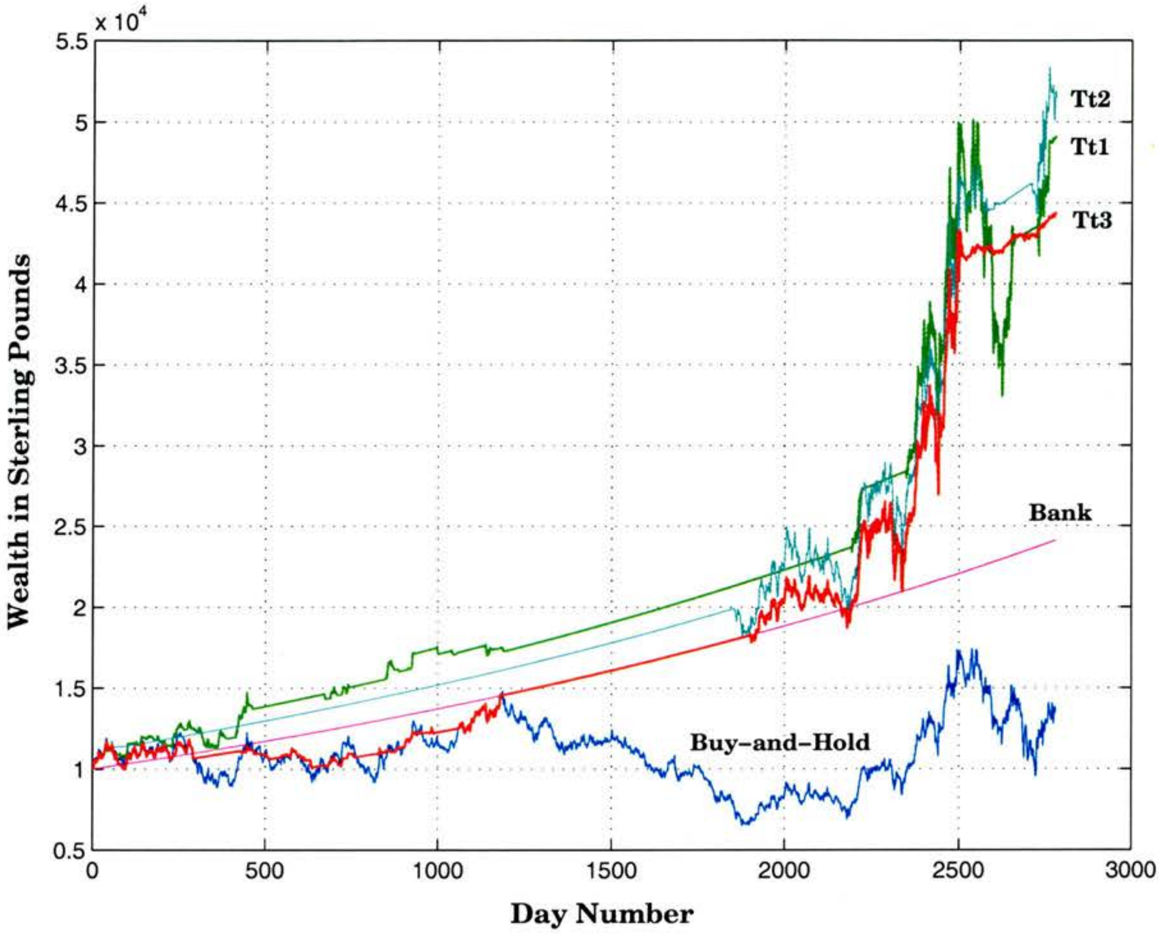

Bank	24,107
Buy-and-Hold	13,437
Trend-Following	25,873

GA	Tt1	Tt2	Tt3	Tt1	Tt2	Tt3
50	49,031	51,428	44,311	43%	79%	52%
75	38,321	48,026	35,044	43%	76%	54%
100	39,623	40,354	34,298	44%	77%	53%
125	33,996	39,046	32,981	47%	75%	51%
150	33,194	41,593	30,185	49%	73%	51%
175	33,179	37,360	31,176	46%	75%	50%
200	30,278	35,278	27,772	48%	72%	51%

Figure 8.6: UK: HNS Stock. See text for explanation

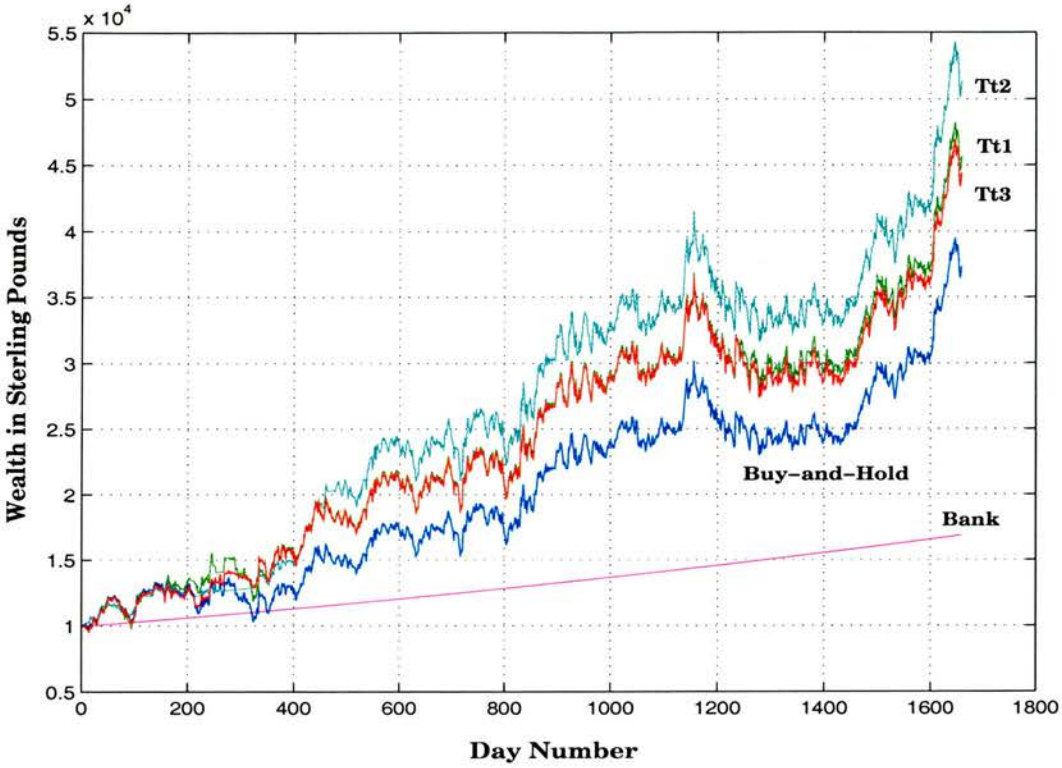

Figure 8.7: UK: LLOY Stock. See text for explanation

	GA	Tt1	Tt2	Tt3	Tt1	Tt2	Tt3
Bank	50	45,300	49,817	44,404	4%	5%	6%
16,906	75	41,453	44,724	43,961	4%	4%	8%
Buy-and-Hold	100	42,245	51,316	40,353	4%	5%	5%
37,321	125	40,651	45,664	43,082	4%	5%	6%
Trend-Following	150	40,407	43,170	44,375	4%	6%	7%
22,082	175	42,798	50,390	44,233	5%	4%	7%
	200	45,611	45,011	41,091	5%	6%	7%

Figure 8.8: UK: LLOY Stock. See text for explanation

strategies that beat *buy-and-hold*. Up to day 1,300 there was not much to be learned about this stock except that it is very stable. It steadily increases its value at a higher rate than the bank, so that all *trader types* prefer this option over the bank's investment over this period, but after it finishes, the price moves more often and we can see a quick response from the traders: it then becomes more profitable to trade the stock during the ups and downs. As a result they outperform *buy-and-hold*.

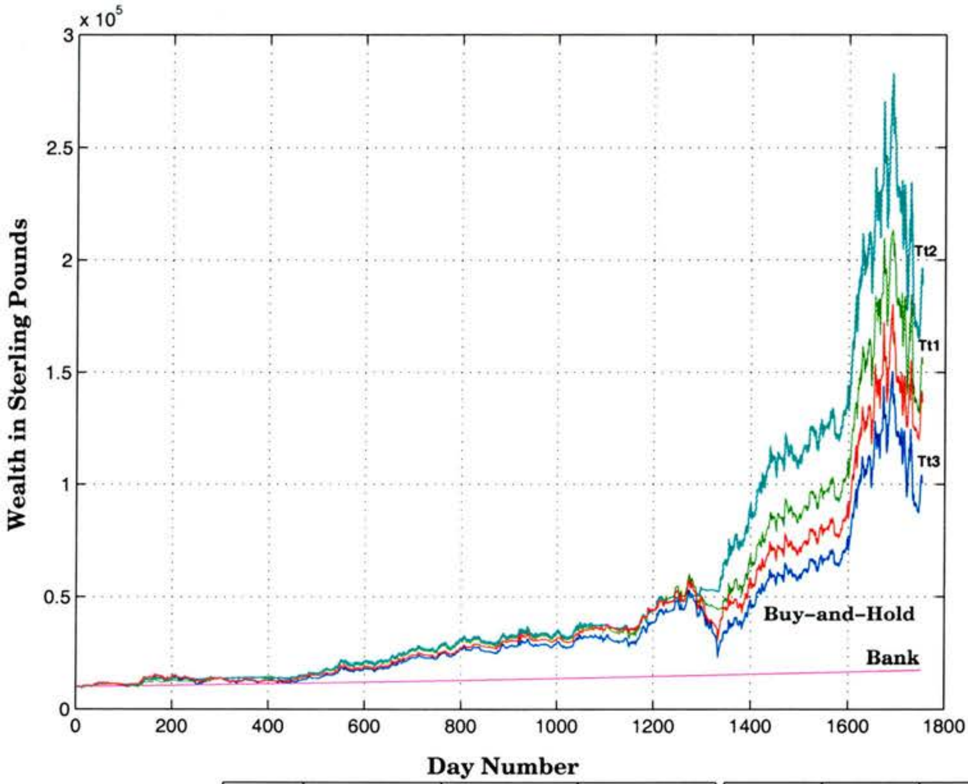

	GA	Tt1	Tt2	Tt3	Tt1	Tt2	Tt3
Bank	50	145,186	155,327	128,486	9%	16%	8%
17,411	75	153,689	175,426	134,215	6%	17%	7%
Buy-and-Hold	100	146,521	149,360	117,254	6%	18%	7%
100,465	125	130,382	146,653	127,413	7%	18%	8%
Trend-Following	150	143,507	188,925	134,263	5%	20%	7%
140,829	175	130,432	164,280	136,399	5%	20%	9%
	200	144,962	151,375	121,086	4%	20%	9%

Figure 8.9: UK: WPP Stock. See text for explanation

8.6 US Stocks

8.6.1 Case 6: Microsoft Corp.

Figure 8.10 corresponds to Microsoft's stock. This performed exceptionally well, so that the *buy-and-hold* strategy was excellent, far outperforming the *trend-following* strategy, which performs very poorly in this stock. As it can be seen from the table in the centre just below the graph, the best runs all happened with GA-period of 75.

As it is hard to recognise what is happening from this graph in such a large period, in figure 8.11 the last 654 days of the series are zoomed and the *trend-following*'s wealth was added as well. All traders outperformed by far the *bank* investment, the *buy-and-hold* and the *trend-following* strategy. It is interesting to note how exactly they managed to do so: Tt1 was more or less holding a number of shares until around day 3,500; see figure 8.12, when on day 3,484, it abruptly sold all its 50,297 shares. This trader's profit from the last days comes from keeping its money in the bank (that is why its wealth steadily increases bank-like from this day on), it ends owning no shares at all.

Tt2's earned profit after the drop in prices comes from an interesting mix of buy/sell behaviour. It did not adopt any particular fashion, instead as you can see from figure 8.12, its shares oscillate somewhat during this unstable period, while before that time, it was steadily holding around 50,000 shares (all it could afford). At the end it owns 88,994 shares. Note that with such transactions it earned almost 40,000 shares!

Tt3 almost precisely followed the *buy-and-hold* strategy until the end, when thanks to the generalised drop in technology stocks, lost about 40% of its wealth, still keeping at the end all its wealth in 83,234 shares.

8.6.2 Case 7: Cabletron Systems, Inc.

Figure 8.13 shows Cabletron Systems, a very volatile stock that performed somewhat similar to GKN, where the *trend-following* strategy was unbeatable. However, Tt2 came close to equalling the trend strategy. The table below the figure shows it was quite easy to generate strategies that beat the *buy-and-hold* strategy; in particular, Tt2 found it the easiest, as it beat it almost half of the time.

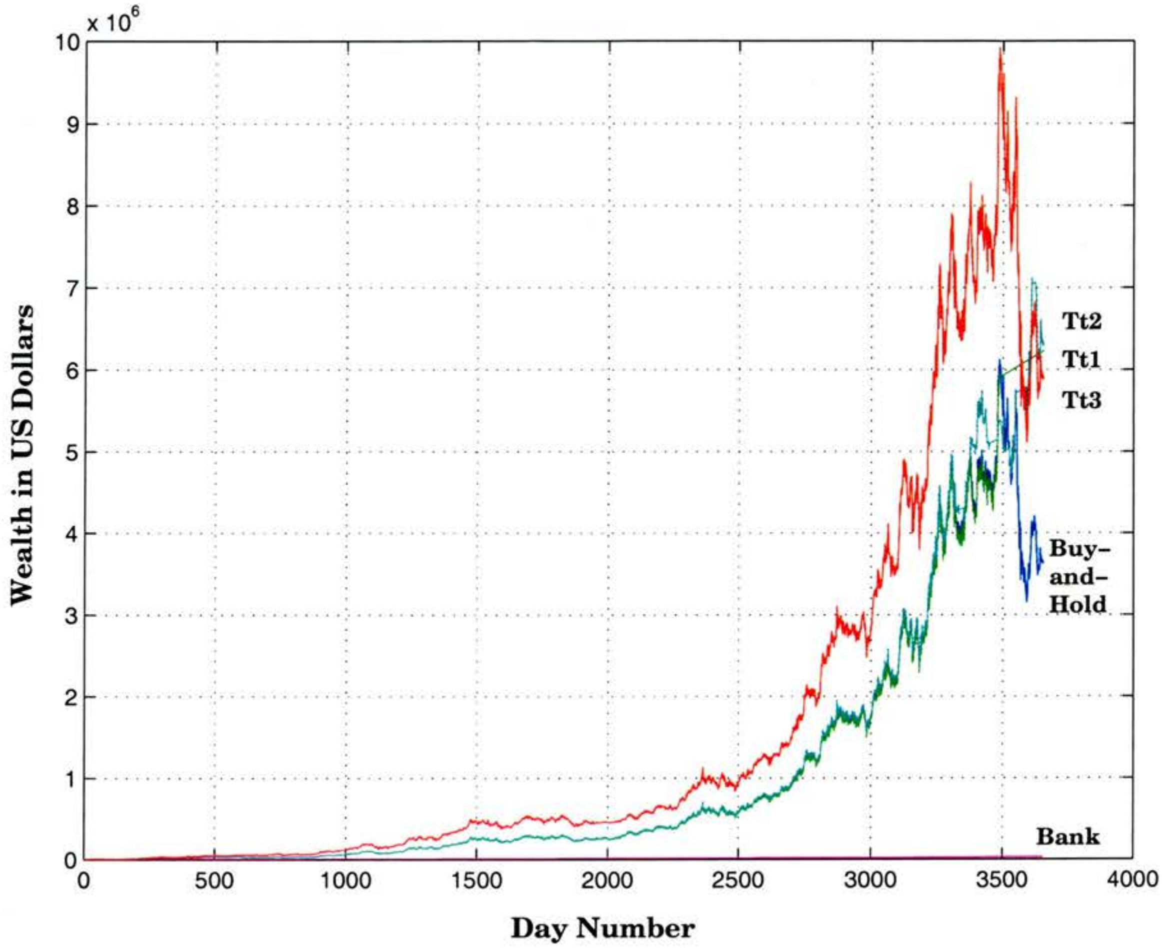

Bank	
31,791	
Buy-and-Hold	
3,629,278	
Trend-Foll.	
331,860	

GA	Tt1	Tt2	Tt3	Tt1	Tt2	Tt3
75	6,229,070	6,285,237	5,878.406	4%	2%	8%
100	5,548,541	5,647,181	5,508,190	3%	2%	8%
125	5,509,811	5,878,918	5,171,595	3%	2%	6%
150	4,733,077	7,615,733	4,598,360	3%	2%	7%

Figure 8.10: USA: Microsoft Stock. See text for explanation

8.6.3 Case 8: Forest Oil, Corp.

Figure 8.14 shows Forest Oil which, like many oil companies in the last decade, lost value reasonably steadily. The *bank* strategy therefore beat *buy-and-hold* easily, but

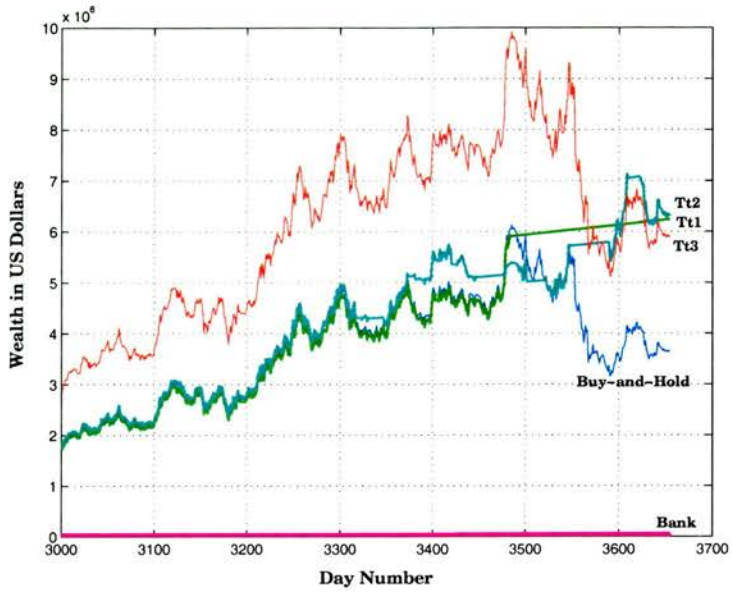

Figure 8.11: USA: Microsoft Stock. Wealth of Bank, Buy_and_Hold, Trend_Strategy and Traders Tt1, Tt2 and Tt3 from day 3,000

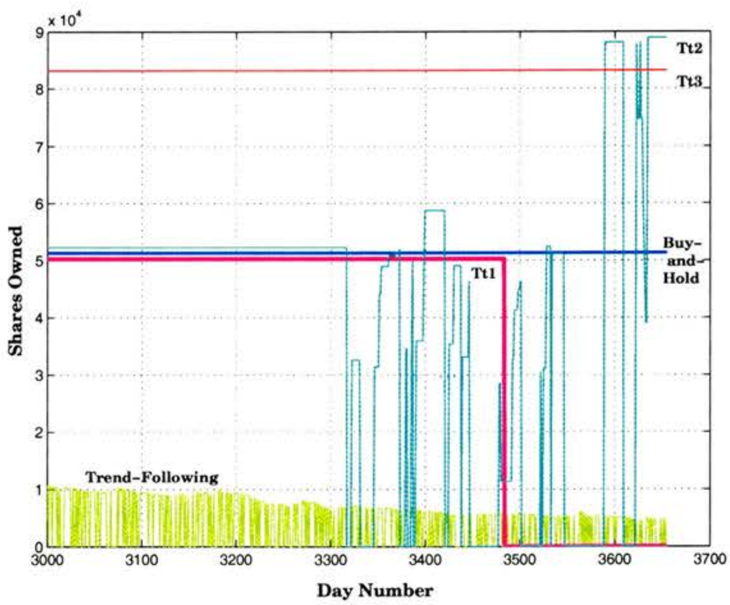

Figure 8.12: USA: Microsoft Stock. Shares Owned by Buy_and_Hold, Trend_Strategy and Traders Tt1, Tt2 and Tt3 from day 3,000

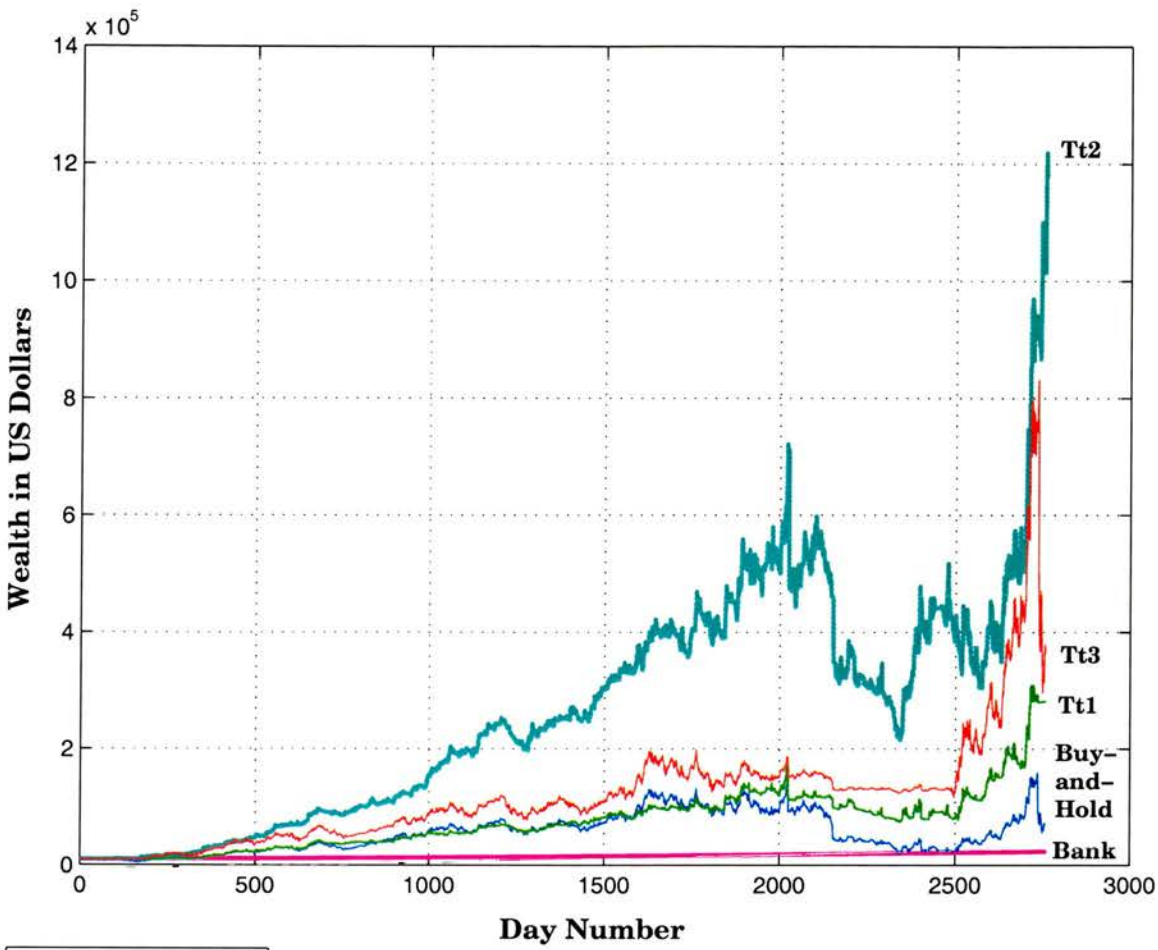

Bank	23,933
Buy-and-Hold	69,772
Trend-Following	1,312,833

GA	Tt1	Tt2	Tt3	Tt1	Tt2	Tt3
75	189,432	306,233	365,622	30%	45%	25%
100	175,472	1'175,583	194,739	30%	45%	26%
125	199,400	399,772	260,258	28%	47%	23%

Figure 8.13: USA: Cabletron Systems Stock. See text for explanation

all three *trader types* beat the *bank* spectacularly and Tt1 and Tt3 did very well to exploit the brief upswing around day 2,250. *Trend-following* strategy was the worst performer of all. As it has been observed in the previous cases, it seems that this strategy is particularly good for highly volatile stocks. This stock shows the opposite.

Low volatility and lower prices persist throughout. So how can the agents make a profit under these conditions? They have clearly learned to trade more during periods of high movement, keeping their money in the bank only when price constantly drops and shows no signs of improvement. It was extremely easy to beat the *buy-and-hold* strategy. In fact, in all cases it has been beaten by far margins.

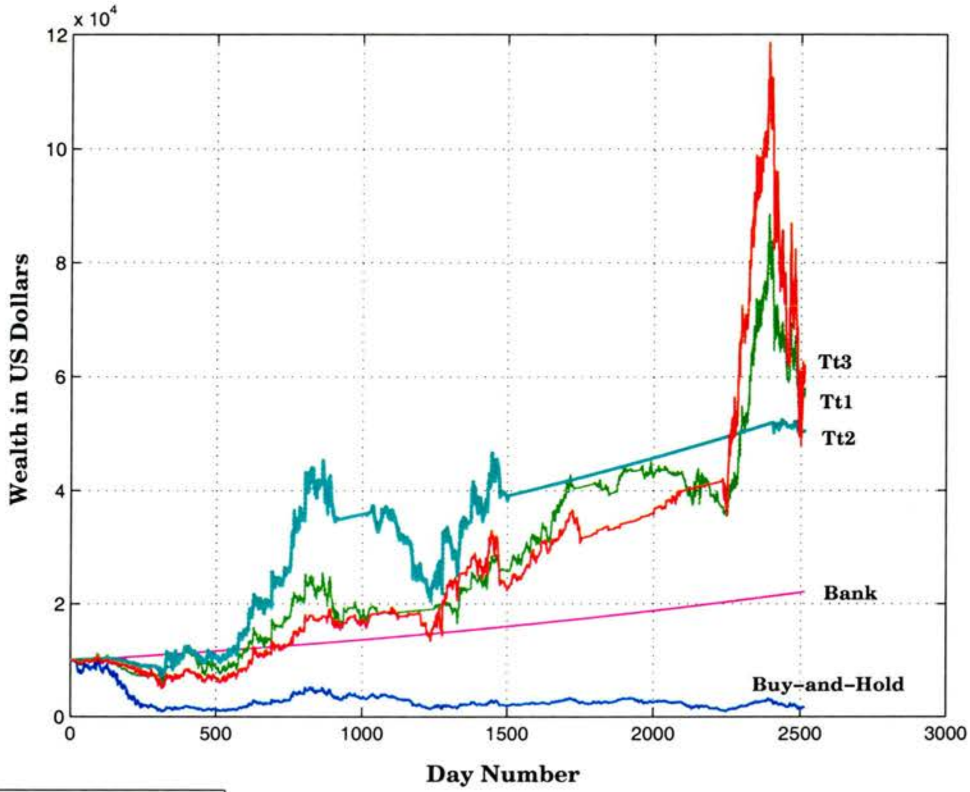

Bank	22,154
Buy-and-Hold	1,716
Trend-Following	19

GA	Tt1	Tt2	Tt3	Tt1	Tt2	Tt3
150	58,010	30,252	23,459	100%	100%	100%
300	44,842	50,421	20,120	100%	100%	100%
700	13,972	21,975	62,207	100%	100%	100%

Figure 8.14: USA: Forest Oil Stock. See text for explanation

Chapter 9

Agent Dynamics, Continuous Learning and Adaptability

> 'The greatest weight. What, if some day or night a demon were to steal after you into your loneliest loneliness and say to you: 'This life as you now live it and have lived it, you will have to live once more and innumerable times more; and there will be nothing new in it, but every pain and every joy and every thought and sigh and everything unutterably small or great in your life will have to return to you, all in the same succession and sequence - even this spider and this moonlight between the trees, and even this moment and I myself. The eternal hourglass of existence is turned upside down again and again, and you with it, speck of dust!' Would you not throw yourself down and gnash your teeth and curse the demon who spoke thus?... Or how well disposed would you have to become to yourself and to life to crave nothing more fervently than this ultimate eternal confirmation and seal?" Friedrich Nietzsche, *The Gay Science*, s.341, translated by Walter Kaufmann.

Three important aspects that deal with the process of evolution of the agents are presented on this Chapter. First, Section 9.1 describes ways of examining the *agent's dynamics* more closely; then Section 9.2 presents a discussion of why the learning process proposed in this model is *continuous* and finally, Section 9.3 addresses whether these agents *truly adapt* to the market environment. Recall that the agents start with no prior knowledge and therefore are forced to make random decisions from the start of their learning path. Over time, they evolve successful sets of rules and continually adapt to different market environments that they have not encountered before.

9.1 Analysing Agent Dynamics

Throughout this thesis, it has been pointed out that one of the advantages of this model is that it is particularly well suited to the analysis of *agent dynamics*. Section 7.12 describes a number of methods to analyse the relevance of both, certain bits of information, and particular rules from the evolved ecology. Chapter 8 described specific examples where the agents were analysed from a broad perspective; for instance, by looking at how successful the agents can be at learning useful strategies, how they react to certain changes in the real market environment they live in, etc. These properties are important to describe, many important issues were be addressed when viewing the agent from this perspective. This type of analysis of agent behaviour will be referred as the *macroscopic* approach to agent dynamics.

However, one would also like to know more about another aspect of the agent's learning process: its internal composition. For instance, it would be important to address how the agent's structure of beliefs is formed, how it changes over time and whether it would be possible to control some aspects of the learning process. This model provides ways to learn more about the agent internally, in addition to the external behaviour already described. This refers to the second way in which one can learn about the agent dynamics: the *microscopic* approach, which will be illustrated in the following sections. This approach concentrates in studying the agent's evolved rule sets under four different reinforcement criteria and various market environments, which will be explained in the following sections, after describing the market environment.

9.1.1 Describing the Market Environment

The US stock that will be considered for the purpose of analysing the agent from a *microscopic* perspective is shown in table 9.1, where the profile and dates of the series analysed are described. From now on, the stock will be addressed by its symbol, as shown in table 9.1, and its properties and trends by day numbers rather than the actual dates when they occurred.

Table 9.1: Profile of Stock Analysed

COMPANY NAME	SECTOR	INDEX MEMBERSHIP
IBM Corp.	Tech./Comp.Hardware	Dow Industrials and S&P 500

SYMBOL	FROM	TO	NO. OF DAYS
IBM	02/01/85	13/01/00	3,800

9.1.2 Reinforcement Scheme and System Parameters

In section 7.11 the mechanics of the reinforcement scheme were explained in detail. Specifically, it was described that reward is given to rules according to their **specificity**. To illustrate this, consider the following two rules:

Rule No.	Condition	Action	Specificity	Condition Size	Reward Received
1	011100#	1 101	6	7	$R\frac{6}{7}$
2	###1#0#	1 010	2	7	$R\frac{2}{7}$

According to the reinforcement algorithm developed for this model and therefore adopted throughout this thesis, if both rules are correct, rule No. 1 will receive more reward than Rule No. 2 due to the fact that it is more specific. The argument that motivated the development of this new scheme is that under normal rewards (as described in [Goldberg 89]), Rule No. 2 is subject to receive more reward because due to its high number of # symbols, it matches more environmental states than the more specific rule. However, in financial markets, finding specific trading strategies is more valuable. The search here is for specific rules, but a general rule is also desirable, as long as it has proven to work better than chance, i.e., in cases where it is a genuine generalisation of a specific rule, supporting hierarchy formation. However, these good general rules are very difficult to obtain due to the presence of non-stationary in financial markets.

The reinforcement scheme developed here can be regarded as a novel property of this system. So far, it has not been found in literature a system that supports the formation of highly specific rules through the *reinforcement component* of a LCS. David Goldberg, in [Goldberg 89], describes a way to encourage specific rules over general rules from the *apportionment of credit component*; that is, directly through the auction algorithm, making bids proportional to the classifier's specificity, strength and other input parameters. A more specific classifier will make a higher bid than a general one, thus winning rewards more often –if it is correct of course. This bidding structure according to specificity has also been adopted in this model. However, it was observed that it was not enough to create pressure towards the formation of specific classifiers through the bidding mechanism alone. The sets of evolved rules were extremely general and provided no useful information. As a result, pressure was added from the *reinforcement component* as well.

So far, the reinforcement approach proposed in this thesis works well. However, there is still some concern regarding this matter. For instance, what if Rule No. 1 (the specific one) is not as good as the general one? Rationing the reward according to specificity puts pressure on the creation of more specific rules, which as explained earlier, is a desirable property. All experiments given so far use this reinforcement scheme. However, rule ###1#0# could genuinely be a better rule than its more specific rival 011100#! This can happen when the bits of information indicated by # symbols of the more general rule are not relevant, so the agent only finds that bits 4 and 6 of the general rule are the most important ones. In this case, it seems important to define a reinforcement scheme that rewards the rules according to their **accuracy**. Rules that do not apply too often, but that are very accurate, will receive more reward and therefore increase their chances to remain in the evolving population of rules. As described in section 3.4.2, this property motivated the development of accuracy based classifier systems. This issue will be explored in this section.

In addition to reward according to specificity and accuracy, another important factor that has to do with the *type of decision* needs to be explored in more detail. So far, matching *buy-rules* are rewarded if the stock's price increases more than 2% in the following period. In the same way, *sell-rules* are rewarded if the price drops more than 2%. This percentage was chosen to make the agent sensitive to bigger price changes and to account for transaction costs. However, in occasions of uncertainty or when the market is not trending in any one direction, it might also desirable to hold. Therefore, rules that hold also deserve to be rewarded by the environment. The criteria developed is that a *hold-rule* is rewarded when the price fluctuates within 5% of the price moving average of the last six weeks. In earlier experiments, *hold-rules*, fluctuations as high as 20% were rewarded, but inspection of the rule sets showed a clear dominance of such rules, allowing the agent to buy and sell very rarely. Experiments showed that rewarding 2% and 5% produced better performances than 20%; however, sometimes it might be desirable to be "less reactive" to price changes. In this case, it is suggested to use higher reward percentages of *hold-rules*.

The following sections provide with experiments that have been designed to propose ways of analysing the agent under the following two different reinforcement schemes:

1. According to *specificity*. The $Strength$ of a given classifier i is increased according to the following formula:

$$Strength_i = Strength_i + Reward\frac{Specificity_i}{ConditionSize_i}, \tag{9.1}$$

   where $ConditionSize$ is defined as the length of the condition part. For instance, a classifier with $ConditionSize = 7$ and $Specificity = 5$ (5 of its 7 bits are non# symbols) will receive 5/7 of the total $Reward$ amount (a typical $Reward = 0.5$). One whose $Specificity = 2$ will be paid only 2/7 of the total $Reward$ amount.

2. According to *accuracy*. When paid by the environment, the classifier's strength is increased according to the following formula:

$$Strength_i = Strength_i + Reward\frac{CorrectDecisions}{MatchingTimes}, \tag{9.2}$$

   where $CorrectDecisions$, as indicated by its name, refers to the number of times the rule was correct (i.e., decided to buy when the price increased more than 2%) and $MatchingTimes$ is the total number of times the classifier has matched. For instance, when a classifier matched 50 times in a given period of time, but only 5 of these has been correct (a very inaccurate one), it will only receive 5/50 of the reward amount, which will most probably be smaller than the taxes it has to pay for matching and its bid payment in case it was a winner, so this classifier will tend to disappear.

Every one of these criteria will be reported in two different reinforcement scenarios, where the reward of *hold rules* is given depending on whether the price fluctuates within 2% or 5% of the six-week moving average. This makes a total of four reinforcement schemes analysed.

Table 9.2 describes the system parameters used in the following experiments. A detailed description of these parameters is given in section 7.11.

Table 9.2: Parameters Used

Parameter	**Tt1**	**Tt2**	**Tt3**
GA period	[50-150]	[50-150]	[50-150]
$P_{\#}$	0.5	0.5	0.5
Automatch Flag	ON	ON	ON
No. of classifiers	100	100	100
Specificity Flag	ON	ON	ON
Noise Flag	ON	ON	ON
Bucket Brigade Flag	ON	ON	ON
Reinforcement	0.5	0.5	0.5
Number of shares	0	0	0
Initial Cash given	$10,000	$10,000	$10,000
Commission per trade	0.01%	0.01%	0.01%
Annual Interest Rate	8%	8%	8%
Seed in steps of 0.001	[0-1]	[0-1]	[0-1]
Number of Runs	1001	1001	1001

9.1.3 Case IBM

Figure 9.1 displays the price of IBM stock from Jan 03, 1985 to Jan 13, 2000. Note that the price tends to decrease during the first half of the period (bull market) and the opposite happens during the second half (bear market), when the price starts to increase quite dramatically. This is indeed a very difficult stock to analyse because initial indications would seem to suggest that this stock shows a very erratic behaviour. It would be interesting to find out if the artificial traders are able to learn some regularities when none seem to exist. At this point the question to address is: How do the agents react to a price behaviour such as this one under the following reinforcement schemes?

1. According to specificity and holding at 5%
2. According to specificity and holding at 2%
3. According to accuracy and holding at 5%
4. According to accuracy and holding at 2%

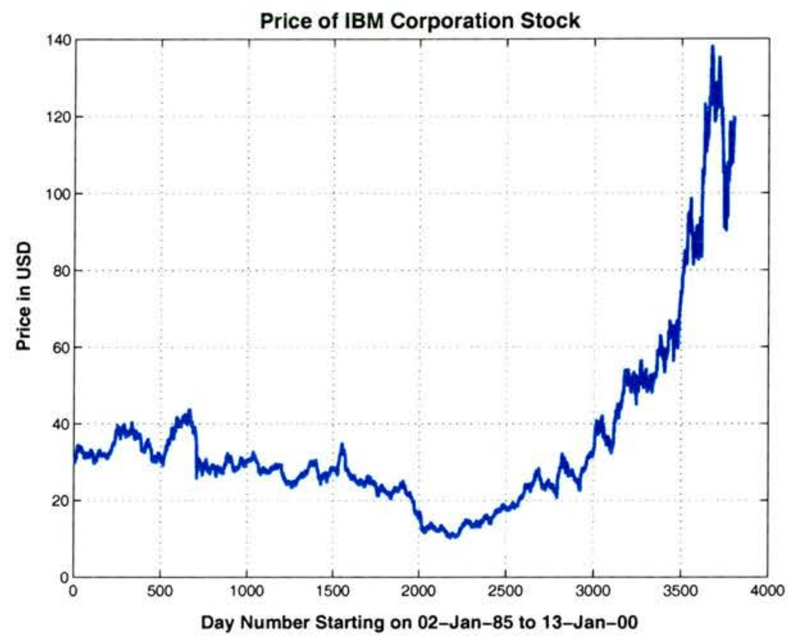

Figure 9.1: IBM Stock Price, from Jan 03, 1985 to Jan 13, 2000

The following tables of results provide with the final wealth of the *bank*, *buy-and-hold* and *trend-following* strategies, along with the average, best performer and percentage of runs beating *buy-and-hold* of IBM stock, corresponding to three different values of GA: 50, 100 and 150. Note that out of the three strategies, the best investment return was produced by the *buy-and-hold* strategy with $39,833, followed by the *bank* investment, which produced $33,294 and finally the *trend-following* strategy with only $12,557. The wealth of these three strategies will be attached at the top of each graph to facilitate making comparisons. The column *Average* provides the average return per agent, from 1001 runs using a *Ga_period* specified in the first column. Recall that in all experiments, the agents start with no prior knowledge, only with random strategies that they start using since day one. In the same manner, for every agent, the returns of the best performers over the same number of runs are also provided, along with the number of runs that beat the *buy-and-hold* strategy.

The upper half of table 9.3 provides the summary of results of runs performed with reinforcement according to specificity and rewarding the rules that recommend to hold when the price fluctuates +/-5% from the moving average of the past 6 weeks. Tt2 can be identified as the best of the three traders. With *Ga_period=50*, and an average wealth of $47,525, this trader outperforms the wealth of the *buy-and-hold* strategy by far. This is also shown in the number of runs that outperform *buy-and-hold*, which is more than 50% of the time. The best performer of all three *Ga_periods* is also Tt2, with a final wealth of $216,144. Although the best average was achieved by Tt1 with *Ga_period=150* and beating *buy-and-hold* by 52%, this type of trader does not have the best performer, which in all cases was Tt2.

The same information is provided in the bottom part of table 9.3, except that here the reinforcement is given to rules that recommend to hold at 2%. Tt2 was easily identified as the best of the three traders under specificity and hold at 5%. However, decreasing the hold percentage to 2 results in a big drop in performance. Under this scheme, this type of trader drops from being the best to the worst on average. Tt1 is a clear winner on average under all *Ga_periods*. Tt2 was more sensitive to this change; Tt1 and Tt3 seem to be less disrupted when changing the percentage of reward from 5 to 2, their performances are very similar in both cases.

Table 9.3: IBM Stock. Reward According to Specificity, Holding at 5% (top half) and 2% (bottom half)

Bank $33,294	Buy-and-Hold $39,833	Trend-Following $12,557

GA 050	Average	Best	Beat B&H
Tt1	$37,941	$182,475	31.8%
Tt2	$47,525	$216,144	50.6%
Tt3	$33,590	$163,885	17.9%

GA 100	Average	Best	Beat B&H
Tt1	$43,314	$195,859	43.4%
Tt2	$49,347	$209,519	47.0%
Tt3	$35,444	$208,200	21.0%

GA 150	Average	Best	Beat B&H
Tt1	$49,357	$174,481	52.0%
Tt2	$47,811	$211,729	42.7%
Tt3	$36,669	$173,713	22.1%

GA 050	Average	Best	Beat B&H
Tt1	$37,074	$188,712	31.3%
Tt2	$33,553	$133,978	27.9%
Tt3	$33,494	$170,742	20.1%

GA 100	Average	Best	Beat B&H
Tt1	$44,135	$172,385	43.6%
Tt2	$32,728	$183,873	25.5%
Tt3	$35,313	$189,511	19.7%

GA 150	Average	Best	Beat B&H
Tt1	$48,593	$164,851	49.7%
Tt2	$32,793	$131,769	25.3%
Tt3	$36,394	$172,657	22.0%

The next reinforcement schemes are based on accuracy. In the same way as in the specificity scenarios, two holding rewards are given: 5% and 2%, both shown in Table 9.4. In general, results for Tt1 and Tt2 are consistent throughout the average column: they are not better than those obtained with rewards according to specificity, but Tt3 clearly improved. Under 5% of hold rewards, Tt2 continues to be the best trader in terms of all criteria: average, best and beat *buy-and-hold*, except for one case, under *GA_period=50*, where the best performer was Tt1. The percentages that beat *buy-and-hold* decreased considerably for Tt1 in all cases and for Tt2 decreased for *GA_period=50* and *GA_period=100*, but not with *GA_period=150*, where there was only a slight increase from 42.7% to 43.7%.

Finally, reward holding at 2% is shown at the bottom of the table. The most remarkable change is indeed seen in Tt2. It was previously noted that this trader performs much better at 5% under specificity. However, going from 5% to 2% in accuracy seems to even be more detrimental: all three criteria decreased by large amounts. Take, for example, the drop from 46.9% of beat *buy-and-hold* under accuracy at 5%, to only 12.4% at 2% with *GA_period=100*!

There seems to be a clear relationship here: Tt2's performance drops proportional to decreases in hold reward. The reason for this is that because this trader receives volume information, it is very susceptible to price changes and it tends to buy and sell more often rather than holding. For this trader, better performances are obtained with higher rewards to hold-rules because this inhibits it, to a certain extent, from performing such high frequency transactions. To illustrate this, an additional run was made where no rewards were given to hold-rules. If this hypothesis is correct, then Tt2's performance will be lower at 0% than 2% and even lower than it was at 5%. Table 9.5 shows the trader's performance of a run with *GA_period=100* and reward according to specificity. As it can be seen in this table, the drop previously mentioned of 46.9% of beat *buy-and-hold* under accuracy at 5%, to only 12.4% at 2% went even lower at 0%, to 3.8%. The hypothesis about this trader seems correct.

Tt3 was not too sensitive to the change in accuracy from 5% to 2%, but it showed large improvements when changing rewards from specificity to accuracy. This can be explained because this trader receives, apart from the first difference in price and

Table 9.4: IBM Stock. Reward According to Accuracy, Holding at 5% (top half) and 2% (bottom half)

Bank $33,294	Buy-and-Hold $39,833	Trend-Following $12,557

GA 050	Average	Best	Beat B&H
Tt1	$34,849	$193,126	26.4%
Tt2	$40,638	$183,141	41.3%
Tt3	$39,456	$172,981	36.2%

GA 100	Average	Best	Beat B&H
Tt1	$38,646	$181,932	26.9%
Tt2	$47,158	$242,123	46.9%
Tt3	$37,775	$193,595	26.1%

GA 150	Average	Best	Beat B&H
Tt1	$39,279	$169,518	27.2%
Tt2	$44,985	$204,628	43.7%
Tt3	$37,650	$191,406	21.3%

GA 050	Average	Best	Beat B&H
Tt1	$35,616	$163,163	27.1%
Tt2	$24,745	$99,232	10.3%
Tt3	$39,691	$211,496	35.5%

GA 100	Average	Best	Beat B&H
Tt1	$39,279	$188,189	29.1%
Tt2	$25,568	$134,851	12.4%
Tt3	$38,237	$200,626	26.6%

GA 150	Average	Best	Beat B&H
Tt1	$39,911	$176,300	26.9%
Tt2	$27,513	$109,718	16.0%
Tt3	$37,377	$184,247	21.7%

Table 9.5: IBM Stock. Reward According to Specificity, Holding at 0%

Bank $33,294	Buy-and-Hold $39,833	Trend-Following $12,557

GA 100	Average	Best	Beat B&H
Tt1	$39,548	$132,245	40.7%
Tt2	$20,068	$59,249	3.8%
Tt3	$22,305	$80,192	5.1%

volume, additional information about the *buy-and-hold* and *bank* investments, i.e., whether it has more shares and wealth than *buy-and-hold* and more cash than *bank*. It also has a reflective property: it knows what it did on the previous day. This trader is very different than the others in the sense that is is not as *technical*. It receives no moving averages of any type, so it is only limited to learn (without knowing their mechanics) the behaviour of *buy-and-hold* and *bank*. This is the reason why it mostly behaves like one or the other, because it learns to mimic these strategies by checking how good or bad it is doing next to them.

For instance, if Tt3 knows that it constantly has less shares and wealth than the *buy-and-hold*, then it learns that buying will increase its wealth and starts adopting the *buy-and-hold* strategy. Then, if the trend changes and it has less than the bank, it changes strategy by adopting the *bank* one. The behaviour of this trader is extremely interesting and it could have a great impact in financial investment. Why? Many times a real trader does not know exactly what another trader is doing, but if he/she knows that he/she has less money or shares than the other, it seems like a good idea to adopt the behaviour of the competitor, whatever that is. Furthermore, if it receives relevant information that the competitor analyses in his/her decision, it could potentially find the strategies that the other is doing. (This behaviour will be illustrated later on, when showing graphs of wealth and distribution of shares of this trader in figures 9.4 and 9.7.)

The following figures show the wealth and distribution of shares of the best of every trader type under specificity rewards. These graphs display interesting behaviours. For example, look at Figures 9.2, 9.3 and 9.4, which refer to Tt1 and Tt2 and Tt3 under specificity at 5%. These traders trade at the beginning of the period (days 1-500 approx.) more or less randomly, and without accumulating excess profits (no large deviations are observed from the *buy-and-hold*). After these transactions, in all traders a large period of tranquility is observed from day 500 to day 2,200 approximately. Looking at the stock price during the period (see the *buy-and-hold* or Figure 9.1), one can see that the market was heading down and not much could be done there to achieve profits apart from leaving the money in the bank. Figures 9.5 and 9.6 refer to Tt1 and Tt2 under accuracy at 5%, and figure 9.7 under 2%; note that similar behaviours are observed with these traders.

However, another novel property of this model is that the agent keeps learning situations even though it is not *experiencing* them. This property was implemented by assigning a special type reward which is a modification of Holland and Reitman's epochal credit allocation plan scheme, the *profit sharing plan* (PSP) [Holland & Reitman 78, Grefenstette 88], where a constant fraction of the current reward (here, according to specificity or accuracy) is paid to all matching classifiers; that is, *every* classifier that becomes active since the last receipt of reward gets paid if the market criteria described above is met, i.e. rewarding buy rules when there was a 2% increase in price, etc.

This means that, if at a certain point, there were reasonable increases in price (>2%), the agent's matching strategies suggesting to buy (and the ones that are correct), are getting stronger, allowing the agent to act appropriately if a similar situation arrives later on. That is why it is observed that all three agents react extremely quickly to the sudden market changes seen in this stock. This feature allows the model to be well suited for environments that change rapidly.

As seen in the graphs, just before the market starts crimping up, the agents decide to buy IBM stock and spend all their available cash. They do that in different ways. Although trader behaviours are never exactly the same, they exhibit certain similarities, such as the ones just described. After day 2,200 they seem to be quite comfortable holding large possessions of the stock, and because they bought it cheaper than *buy-and-hold*, they end up owning about 2,000 shares, as opposed to just over 300 from *buy-and-hold*.

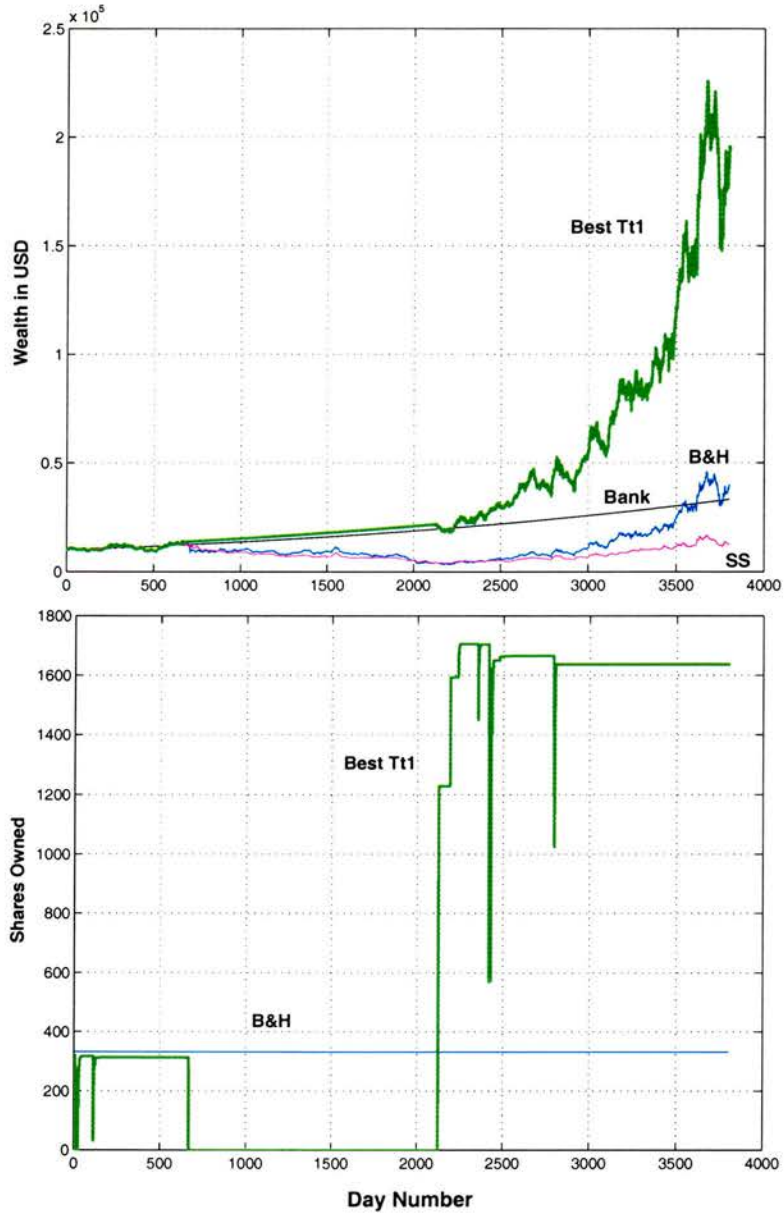

Figure 9.2: IBM Stock. Wealth and Shares of Tt1, Rewarded According to Specificity, Corresponding to Best Tt1 of Table 9.3. Final wealth = $195,859.

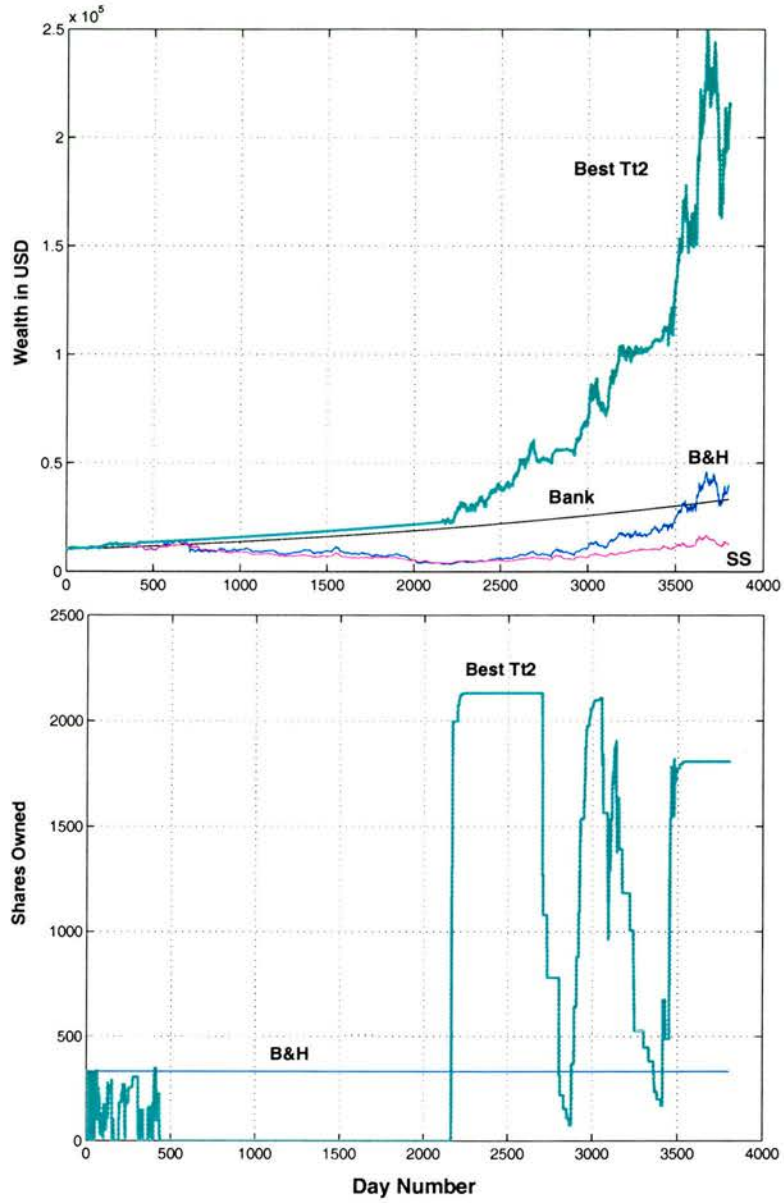

Figure 9.3: IBM Stock. Wealth and Shares of Tt2, Rewarded According to Specificity, Corresponding to Best Tt2 of Table 9.3. Final wealth = $216,144.

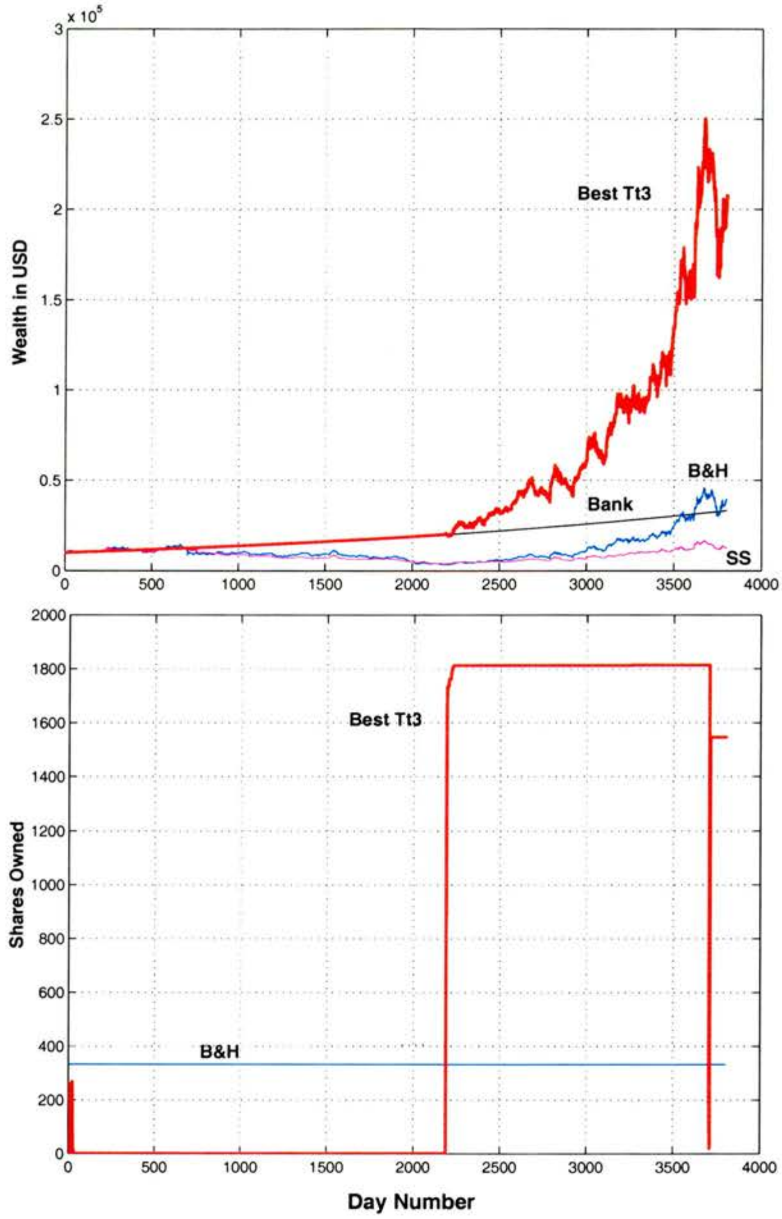

Figure 9.4: IBM Stock. Wealth and Shares of Best Tt3, Rewarded According to Specificity, Corresponding to Best Tt3 of Table 9.3. Final wealth = $208,200.

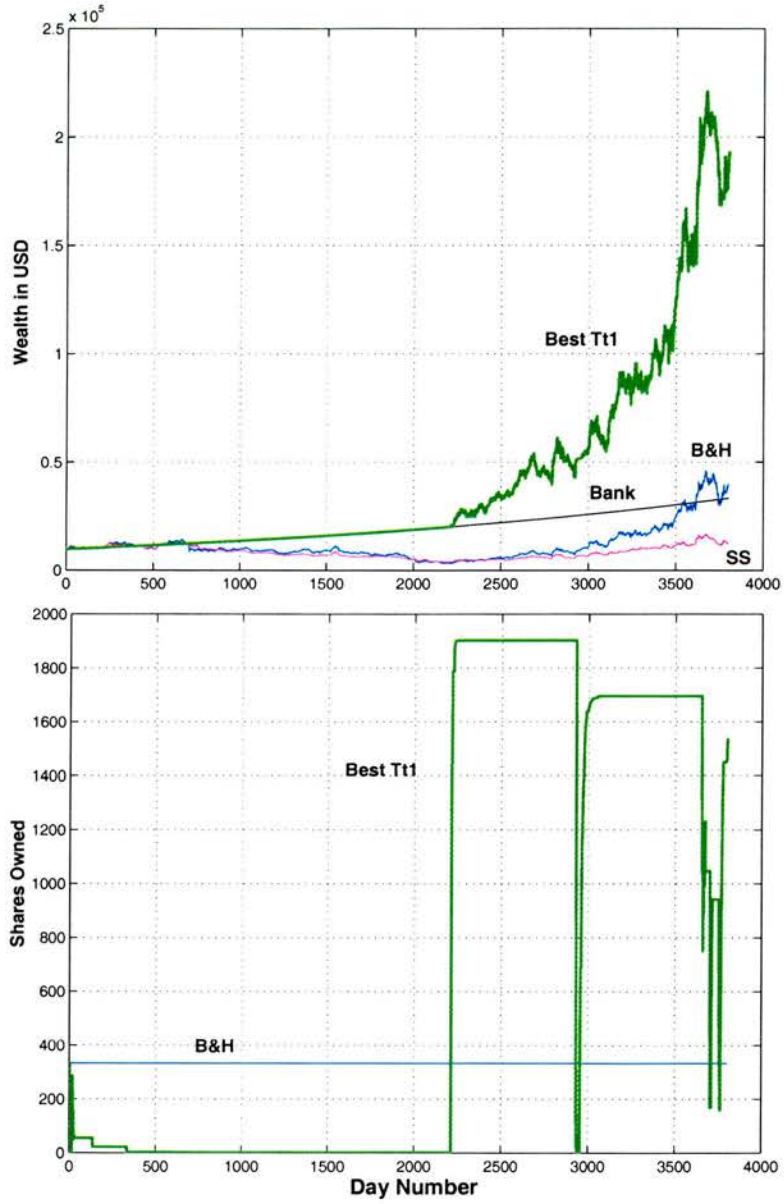

Figure 9.5: IBM Stock. Wealth and Shares of Tt1, Rewarded According to Accuracy, Corresponding to Best Tt1 of Table 9.4. Final wealth = $193,126.

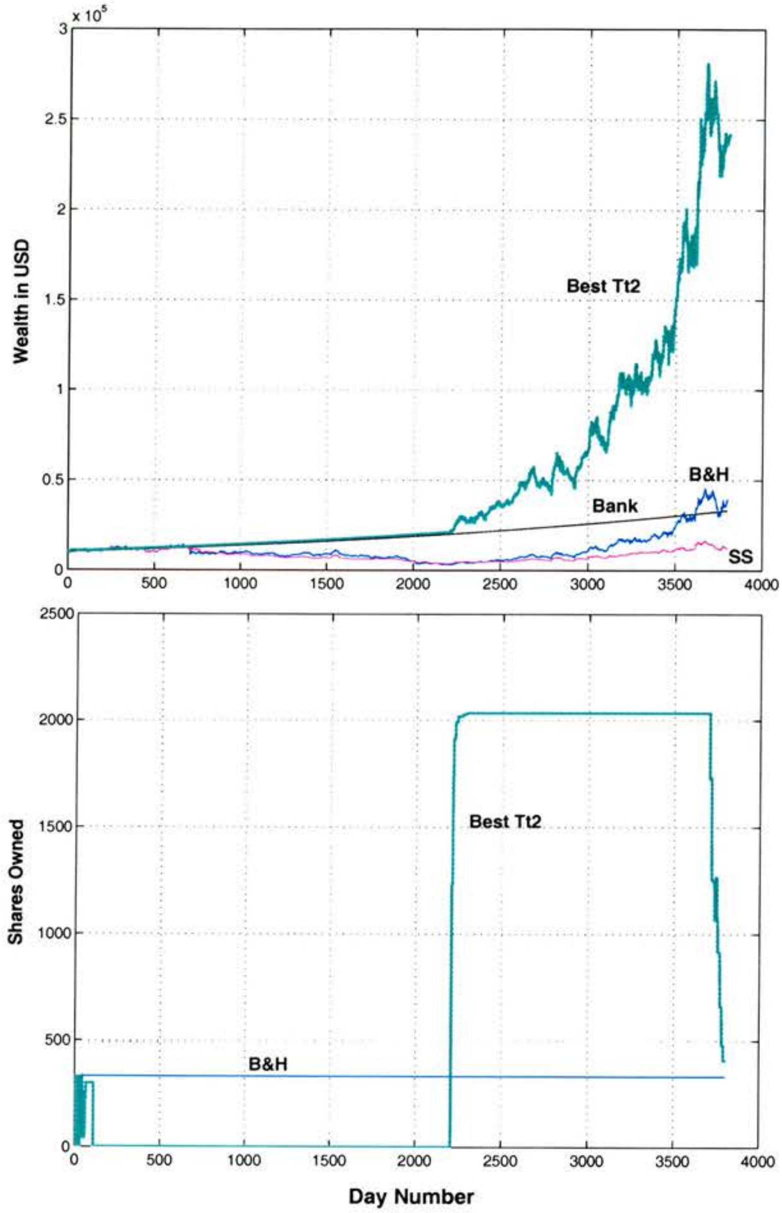

Figure 9.6: IBM Stock. Wealth and Shares of Tt2, Rewarded According to Accuracy, Corresponding to Best Tt2 of Table 9.4. Final wealth = $242,123.

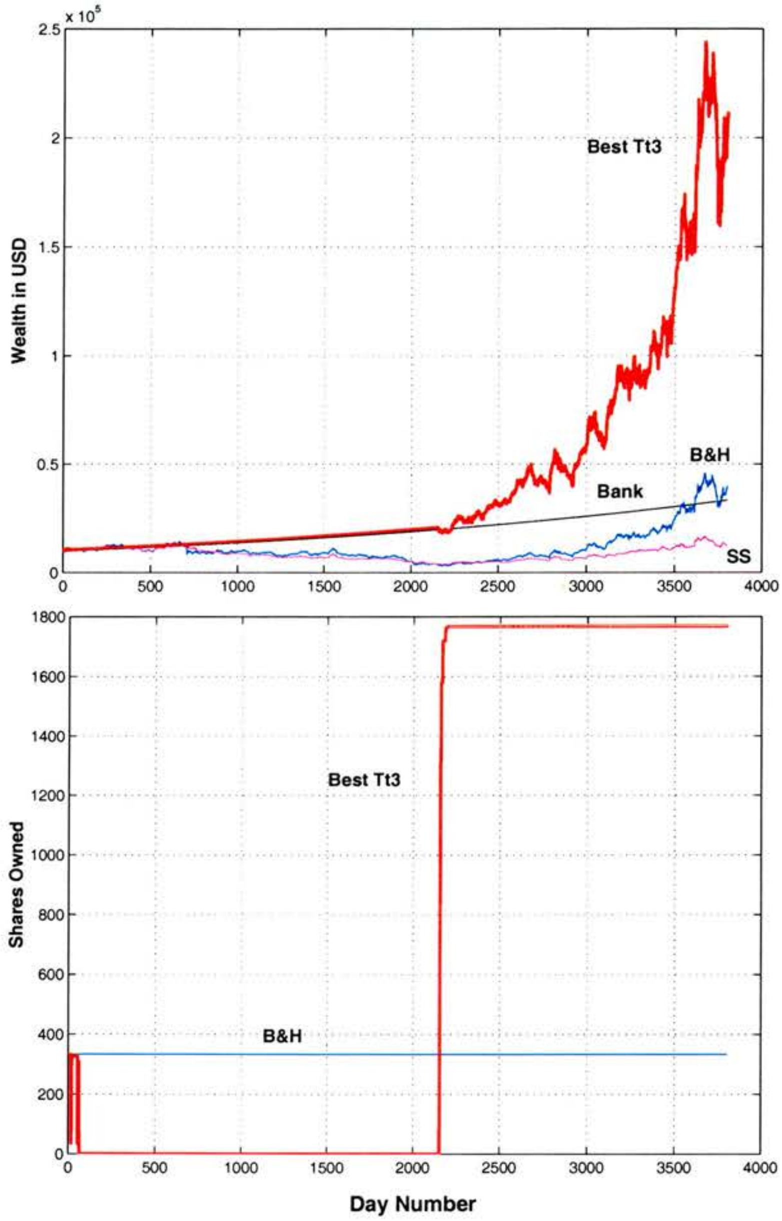

Figure 9.7: IBM Stock. Wealth and Shares of Best Tt3, Rewarded According to Accuracy, Corresponding to Best Tt3 of Table 9.4. Final wealth = $211,496.

9.1.4 Developing New Traders

How easy is it to develop a new type of trader? This section addresses ways in which trader's internal composition can be analysed and further improved. First, in section 7.8 the environmental message of Tt1 was defined. No additional explanations were given regarding the percentages that compare the variables, these thresholds were defined arbitrarily. See table 9.6. For instance, bit No. 2 of Tt1 considers whether the price today is 20% higher than the moving average of the last week. Why 20%? Again, the idea was to make the agents less susceptible to small price changes and to account for transaction costs. Suppose the average of the past week is \$10. If that bit is ON, it means that the price today is higher than \$12. This is a large price increase and the agent must be aware of it. But some stocks do not show too many increases as high as this; certainly IBM does not look like a very volatile stock, so perhaps this high threshold is not a good idea for this stock. On the other hand, many stock prices tend to increase in prices more than 50% of the times, which was the main motive to add a threshold. So, now it is the time to test whether adding the thresholds was a good design decision or not.

To do that, let's create another Tt1, one that does not exhibit those arbitrary thresholds. Table 9.6 shows the environmental message of both Tt1s, the old (called v1.0) and the new (v1.1). However, a second test will be performed in this section. Still regarding Tt1, one might wonder if the information it receives is useful or not; perhaps it could benefit with one bit less. To test this hypothesis, a new type of Tt1 will be created. The former Tt1 that has been used in this thesis will be called Tt1v1.0. A new version is implemented, v1.1, which is very similar to Tt1v1.0, except that it has no thresholds. Then, a third version is implemented, which does not have the 7th bit of information (whether the price today is the lowest ever).

The first hypothesis is that thresholds will help the trader get more wealth, that is, that performance of Tt1v1.0 > Tt1v1.1. The second one is that Tt1v1.2 (the one without P_{lowest}) will *not perform as well* as Tt1v1.1 (the one with P_{lowest} because by inspection of the stock, one can see that it presents many lows during the first half of the period. However, it is difficult to say because not many lows appear during the second half.

Trader Type 1 V1.0

Bit Number	Representation
1	$P_t > P_{t-1}$
2	$P_t > 1.2 * P_{MA5}$
3	$P_t > 1.1 * P_{MA10}$
4	$P_t > 1.05 * P_{MA20}$
5	$P_t > 1.025 * P_{MA30}$
6	$P_t \geq P_{highest}$
7	$P_t \leq P_{lowest}$

Trader Type 1 V1.1

Bit Number	Representation
1	$P_t > P_{t-1}$
2	$P_t > P_{MA5}$
3	$P_t > P_{MA10}$
4	$P_t > P_{MA20}$
5	$P_t > P_{MA30}$
6	$P_t \geq P_{highest}$
7	$P_t \leq P_{lowest}$

Trader Type 1 V1.2

Bit Number	Representation
1	$P_t > P_{t-1}$
2	$P_t > P_{MA5}$
3	$P_t > P_{MA10}$
4	$P_t > P_{MA20}$
5	$P_t > P_{MA30}$
6	$P_t \geq P_{highest}$

Table 9.6: IBM Stock. Environmental Message for Tt1 versions 1.0, 1.1 and 1.2

Table 9.7 reports the results obtained with the three versions of Tt1 described above. Note that Tt1v1.0 is the same Tt1 that has been reported throughout this thesis. The new traders are Tt1v1.1 and Tt1v1.2. With *Ga_period=50* (this corresponds to very fast learning, every 50 days 20% of the population is modified by the GA), Tt1v1.0 (the old one) is a better performer than Tt1v1.1, which means that the thresholds introduced originally were indeed favourable for the trader because the average performance of Tt1v1.0 dropped from \$37,941 to \$36,016 and the best performer dropped from \$182,475 to \$121,274, which is a considerable amount. Then, with Tt1v1.2, which has one less bit than Tt1v1.1, the average performance dropped even more in this case, from \$36,016 to \$33,949; however, the best performer increased from \$121,274 to \$138,867.

Table 9.7: IBM Stock. Reward According to Specificity and Holding at 5% for the Old(v1.0) vs two New Tt1(v1.1 and v1.2)

Bank $33,294	Buy-and-Hold $39,833	Trend-Following $12,557

GA 050	Average	Best	Beat B&H
Tt1v1.0	$37,941	$182,475	31.8%
Tt1v1.1	$36,016	$121,274	31.9%
Tt1v1.2	$33,949	$138,687	27.8%

GA 100	Average	Best	Beat B&H
Tt1v1.0	$43,314	$195,859	43.4%
Tt1v1.1	$45,766	$174,830	52.2%
Tt1v1.2	$44,280	$153,154	48.9%

GA 150	Average	Best	Beat B&H
Tt1v1.0	$49,357	$174,481	52.0%
Tt1v1.1	$49,790	$160,386	60.5%
Tt1v1.2	$48,360	$177,565	58.7%

Results for best performers with *Ga_period=100* and *Ga_period=150* suggest that removing the thresholds helps, increasing the average wealth from $43,314 to $45,766 and from $49,357 to $49,790 respectively, but Tt1.2 is still lower than Tt1v1.1 in both cases. Best performer is Tt1v1.0 under *Ga_period=100*, but it is Tt1v1.2 under *Ga_period=150*.

As results appear somewhat mixed, the averages of the values of the three *Ga_periods* were taken, revealing that the best performer of the three was Tt1v1.0 (with thresholds), with an average of wealth of $184,272, followed by Tt1v1.2 with $156,469 and finally, Tt1v1.1 with average of $152,163. The first hypothesis that thresholds help, holds when considering the best performers. However, the second hypothesis does not hold for this criteria of best performers, because Tt1v1.2 shows higher average of best performers than Tt1v1.1.

Again taking the average, but this time of the average wealth, in the three *Ga_periods*, Tt1v1.1 shows the highest, with average \$43,857, followed by Tt1v1.0 with \$43,537 and Tt1v1.2 with \$42,196. Now the hypothesis holds that Tt1v1.2 is worse than Tt1v1.1 on average performance, however, in this case no thresholds looks like a better option.

While the first hypothesis holds for the best performers and not for the averages should not be discouraging, the main point addressed in this section was to show that it is indeed possible to analyse the traders even deeper; and secondly, it is important to point out that many times we are concerned with finding best performers rather than better averages. For the purpose of this thesis, the latter one has been emphasised more. A lesson to learn from these experiments is that results should no be generalised; stocks are very different from each other. What works well in one might not work well in another. Also, an interesting piece of information revealed from table 9.7 is that all the highest percentages that beat *buy-and-hold* come from Tt1v1.1, which means they have no thresholds. This shows that after analysing trader's behaviours, one can easily change their make up and improve them. For example, if one is interested in getting the best performers, then the thresholds showed to be a better alternative. But, in the other hand, if one is interested in more traders that beat the *buy-and-hold* strategy, then it would be best not to have thresholds. As of the P_{lowest} bit, it seems more reasonable to leave it because the averages were consistently lower during its absence.

A similar experiment was performed with Tt2. A new version of Tt2 was created, which has the same information in bits 1-3 as before, except that bit 4 and 5 (the ones about volume highs and lows) were replaced by bits 3 and 4 of Tt1. This new trader is denoted as Tt2v1.1 and is described in table 9.8 and analysed in table 9.9. These results show a better Tt2v1.0 than Tt2v1.1 in both, averages and best performers for *Ga_period = 50* and *Ga_period = 100*, but the opposite for *Ga_period = 150*. A plausible explanation for the difference between the *Ga_periods* is that under fast learning the agents are able to adjust more quickly to the changes in the market environment, thus making better use of the information about highs and lows of volume. A large *Ga_period* may not inject new strategies as fast as they are needed and the agent could end up using other not so accurate rules, therefore not being able to pick up the trends.

The information about the high and lows is with respect to the current day only, so it may provide information that needs to be learned and used more rapidly than the information about the price moving average of the past five and ten days.

Table 9.8: IBM Stock. Old vs New Environmental Message for Tt2

Old Trader Type 2

Bit Number	Representation
1	$P_t > P_{t-1}$
2	$V_t > V_{t-1}$
3	$V_t > V_{MA20}$
4	$V_t \geq V_{highest}$
5	$V_t \leq V_{lowest}$

New Trader Type 2

Bit Number	Representation
1	$P_t > P_{t-1}$
2	$V_t > V_{t-1}$
3	$V_t > V_{MA20}$
4	$P_t > P_{MA5}$
5	$P_t > P_{MA10}$

Table 9.9: IBM Stock. Reward According to Specificity and Holding at 5% for Old vs New Tt2

Bank $33,294	Buy-and-Hold $39,833	Trend-Following $12,557

GA 050	Average	Best	Beat B&H
Tt2v1.0	$47,525	$216,144	50.2%
Tt2v1.1	$42,474	$174,021	44.6%

GA 100	Average	Best	Beat B&H
Tt2v1.0	$49,347	$209,519	47.0%
Tt2v1.1	$48,095	$190,946	50.3%

GA 150	Average	Best	Beat B&H
Tt2v1.0	$47,811	$211,729	42.7%
Tt2v1.1	$48,610	$213,876	52.0%

9.1.5 Examining the Rules

In previous sections, different reinforcement rewards were examined. This section looks at some of the rules evolved during those experiments. Table 9.10, column 1 shows some examples of buy-rules evolved from Tt1v1.0 under accuracy rewards. The second column displays the strength or every rule, where it can be seen, for instance, how the more general rules have lower strengths. The next columns show the specificity of each rule and the number of times the rule matched and was correct (i.e. when suggested a buy/sell before the price increased/decreased by at least 2%, etc.). The total number of matches appears in the next column, followed by its age, which indicates how many days old is the rule since it was created and inserted in the population. Observe that rules appear to be fairly young. Old rules that were bad have been replaced by better ones, and those which were good have been replaced by similar or identical descendants. The age of identical or new versions of old rules inserted in the population is set to 0. This explains why, in the case in which there was a very good random rule created at the beginning of the period, it does not appear to be 3,800 days old. Notice also that the system is indeed finding and developing highly accurate rules.

Table 9.10: IBM Stock. Examples of Evolved Buy-Rules Under Accuracy

Condition[Action]	Strength	Specificity	No. Correct	Matches	Age
1###10# [1010]	50.64	3	29	29	100
1###1#0 [1010]	50.30	3	171	171	600
0#0#1## [1010]	50.01	3	182	182	1100
0###1#0 [1110]	49.62	3	254	254	1400
#0#1100 [1001]	50.07	5	31	31	150
##001## [1010]	48.80	3	30	30	150
#0##### [1111]	24.25	1	123	200	200
00#110# [1111]	51.06	5	7	7	100
0##1100 [1111]	48.83	5	46	46	400
##0#1#0 [1001]	51.28	3	46	46	100

Due to space limitations, hold and sell-rules have not been added in this table, but they exist. Hold-rules represent the majority of rules in the evolved sets. As they are also rewarded, they achieve perfect accuracy, just like these buy-rules, but their use does not make the agent richer directly, like the buy/sell rules do when their action is taken. It seems that the main duty of hold-rules is to avoid great losses, as shown in the results of the previous sections.

Table 9.11 shows some buy-rules created by specificity rewards. The first observation to point out from these tables is that specificities under reward proportional to accuracy are lower than those evolved from rewards proportional to specificity. Both types of rules are useful, some are general cases of other more specific.

Because the rule specifies in its first part the actual conditions encountered in the market, a closer examination of the sets of such rules can give a broad description of certain properties the stock analysed. For instance, the system could count the number of 0s, 1s and #s in the entire rule set and determine, according to the bits of information considered per agent, whether the market was more bullish or bearish, whether information such as $P_{highest}$ was relevant or not, etc. For instance, in the rules evolved by specificity, one could easily infer that bit No. 2, describing whether today's price is higher than 20% of the moving average of the past week, is 0 in most rules.

This might explain why, in the previous section, the average of the experiments where such threshold was removed proved to be higher than with it. Another bit which is almost always set to 0 is the last one, indicating whether the current price is the lowest ever. This again, explains why the hypothesis that it was an important bit did not hold for the best performers. The table also shows that these rules are reasonably young. The old ones that were describing states when the market was dropping have started to disappear. This system adjusts quickly to new market information.

And finally, observe that bit No. 5 is set to 1 in most buy rules. This provides important information as well. The rule suggests to buy the stock when (among others), the current price is higher than 2.5% of the moving average of the past six weeks.These rules proved to be 100% accurate, as shown in column No. correct. Overall, it can be concluded that the specificity criterion has proven to be the most useful after analysing the rule sets, which indeed, provide with useful information about the market.

Table 9.11: IBM Stock. Examples of Evolved Buy-Rules Under Specificity

Condition[Action]	Strength	Specificity	No. Correct	Matches	Age
1011110 [1010]	47.79	7	3	3	185
000#100 [1101]	49.69	6	122	122	650
10001#0 [1111]	45.13	6	57	57	550
1001#00 [1111]	43.35	6	184	184	1600
1000100 [1111]	49.73	7	75	75	750
1###1#0 [1001]	23.30	3	155	155	500
100#11# [1001]	29.02	5	14	14	200
000#1#0 [1001]	33.55	5	72	72	350
1001##0 [1001]	33.90	5	114	114	650
000#1#0 [1001]	38.01	5	122	122	650
10#11#0 [1001]	35.21	5	236	236	1400
1001##0 [1001]	34.87	5	253	253	1700
1001#00 [1001]	35.07	6	290	290	3800
#0001#0 [1010]	32.26	5	17	17	50
00#1### [1101]	26.47	3	6	6	50
1#0#1#0 [1101]	26.51	4	58	58	200
00##1#0 [1101]	28.86	4	254	254	1400
0#00100 [1101]	25.75	6	31	31	350
##0#1#0 [1001]	24.55	3	35	35	50
##0010# [1101]	28.52	4	407	407	1900
00#01#0 [1101]	34.76	5	178	178	1700
00#01## [1111]	31.55	4	70	70	800
1#0#110 [1111]	32.79	5	69	69	1100

9.2 Continual Learning

A novel property of this model is that in all results reported in this thesis, the agents never see the stock price on any given day of the simulation more than once. The agent is trained according to how successful it will be (from day one) and later in the simulation, according to how successful it has been. Agents are not allowed to borrow money if they run out of it. The process is very different from other time series analysis models such as neural networks, where the data may be presented to the net for training literally thousands of times as part of the process of trying to minimise an error measure. Typically in a neural net system, when training is complete performance is then tested on unseen data. Such systems are usually very slow, need retraining at uncertain intervals and can be unsatisfactory because they offer no convenient explanation of why a given buy/sell/hold decision was made.

Therefore, in the studies presented in this thesis there is no separate testing phase that uses separate data. All the data are unseen, and the learning process is a continual one. The test of success is whether an agent continues to trade profitably, especially when compared with plausible non-evolutionary strategies. Remember, too, that the GA which is responsible for improving the rules runs in most cases reported here only every 50-200 trading days. It is not in the interest of this work to try to examine whether it is possible for an artificial stock-market agent to be trained to learn good trading behaviour through repeated encounters with historical data. The goal is to explore whether such agents can survive in the most human-like way found so far: where opportunities are given only once!

In [Schulenburg & Ross 99] results with separate sets of training and testing were reported. In these experiments, training was performed using 9 years of data and testing using the last year, during which the GA was turned off. In that particular example (Merck & Co.), the testing phase did produce better results than the *buy-and-hold* and the *bank* investments. However, this thesis supports the idea that it is unrealistic to do such tests; in practice, traders do not use a fixed set of rules over an indefinitely long period but change them as market conditions alter.

As an example of this point, consider Luca Beltrametti's learning to forecast experiment [Beltrametti *et al.* 97] on the foreign exchange market with a LCS. In this experiment, the authors evaluated the performance of their adaptive agent against other decision rules which followed the prescription of various economic theories on exchange rate behaviour and the performance of forecasts given by Vector Auto Regression model (VAR) estimations of the exchange-rate's determinants. Although the out-of-sample forecasting ability of the adaptive agent under performed while in-sample forecasting outperformed the rival VAR model, it really doesn't matter. It is important to stress that the authors' purpose was to use the other methods as control devices to test the adaptive agent's goodness of fit by means of a formal statistical tool, i.e. whether the agent could learn to forecast the exchange rate under the conditions they specified in the experiment, not to compare if the adaptive agent was better or worse than the other models.

The position sustained here is that the adaptive learning should never be frozen in non-stable environments such as real stock markets. Brian Arthur [Arthur 92] in the SFI stock market has also tested this, by injecting into his populations of strategies some that had been very good in the past. He observed that such transplanted strategies behaved badly in their new market environment; clearly they had been adapted to specific past conditions. As market behaviour remains unstable and never settles down, learning and forecasting should be continual activities. And because market conditions are ever-changing, results do not show continuous improvements in performance throughout all time. They genuinely adapt. This issue will be further explained in the following section.

9.2.1 Survival Test – The Dinosaurs

> "We find no evidence that market behavior ever settles down; the population of predictors continually coevolves. One way to test this is to take agents out of the system and inject them again later on. If market behavior is stationary they should be able to do as well in the future as they are doing today. But we find that when we 'freeze' a successful agent's predictors early on and inject the agent into the system much later, the formerly successful agent is now a dinosaur. His predictions are unadapted and perform poorly. The system has changed" [Arthur 92].

There are two interesting arguments to discuss in this paragraph. The first one supports the argument that there should be no testing periods, for the reason of non-stationarity given above. The second has to do with the fact that in artificial markets only the behaviours are studied. For example, instead of taking a strategy and implant it in the future, there are ways in which the rules can be studied at any given moment in time. They should not be treated as black boxes. Only in the case in which a NN is representing the adaptive behaviour, it would practically be impossible to be able to successfully track the precise motives for such actions; or alternatively, in some GP cases where the evolved trees are too complex to analyse.

Chen and Yeh also formulated Arthur's survival test in their model by finding that the number of traders with successful searches starts high and then decreases steadily up to a certain point. These findings suggest that initially, traders find a secret but as they also change market dynamics by bringing the knowledge from school back into the market, then more and more find it, making it no longer a secret. They argue that new patterns are created while exploiting the current ones, creating a "*self-destruction-and-organization-process*" [Chen & Yeh 99].

A third piece of evidence suggesting this never-ending process of market evolution was found. The model [Beltrametti *et al.* 97] shows that past strategies are no longer good in the future. For these reasons, it is suggested here that training and testing phases should not be performed separately, the model must function on-line and new strategies should be evolved fast, easily and cheaply.

9.3 Adaptability

How can it be shown that this continuous-learning system *truly adapts* to new market behaviour? What happens to performance as the agents adapt? This Chapter examines these issues in more detail, concentrating on how the agent's sets of rules evolve over time.

In Section 9.2 it was stressed that using the novel continual learning technique used in this model (where the agents constantly adapt to new market behaviour) offers a number of advantages to traditional learning processes. One of such advantages is that there is no need to have separate training and testing phases; here the system saves time and all the efforts needed to assess (i) when should the training phase end (without incurring in over or under-training), (ii) when the system ceases to fit the new information (test set), and (iii) when there is a need to retrain again. This is avoided by continually learning over the unseen data. However, one must also test whether this learning process –here represented mainly by the genetic and the automatch algorithms– is really useful or not. In other words, could it be possible that if a highly successful initial rule discovered by luck be useful throughout a long period of time with the same success? the following Sections address this question by providing illustrations of a number of stocks, where the effects of changing the genetic action is examined at various periods of time.

9.3.1 Stocks Analysed

Similarly to previous chapters, the US stocks analysed are the shown in Table 9.12, along with their symbol and dates. For simplicity, every stock will be addressed by its symbol, as shown in the table, and its properties and trends by day numbers rather than the actual dates when they occurred.

Table 9.12: Period of time and Stocks Analysed

Symbol	From	To	No. of Days
fst	17/10/90	17/10/02	3,027
hit	17/10/90	17/10/02	3,026
ibm	17/10/90	17/10/02	3,028
ko	17/10/90	17/10/02	3,026
mrk	17/10/90	17/10/02	3,028
msft	17/10/90	17/10/02	3,028
pep	17/10/90	17/10/02	3,026
sne	17/10/90	17/10/02	3,028

Figure 9.8 displays the price of these stocks from Oct 17 1990 to Oct 17, 2007. Although not very clear in the figure, after a closer inspection, it can be identified that these stocks are representative of a large range of market behaviour because they

belong to different sectors (i.e. oil and gas, software, consumables, etc.). Therefore they all suffer from several market bubbles and crashes at different periods, high and low volatilities, etc. However, they also share some features due to important global changes experienced after the year 2000 (e.g. heavy price drops after day No. 2,500).

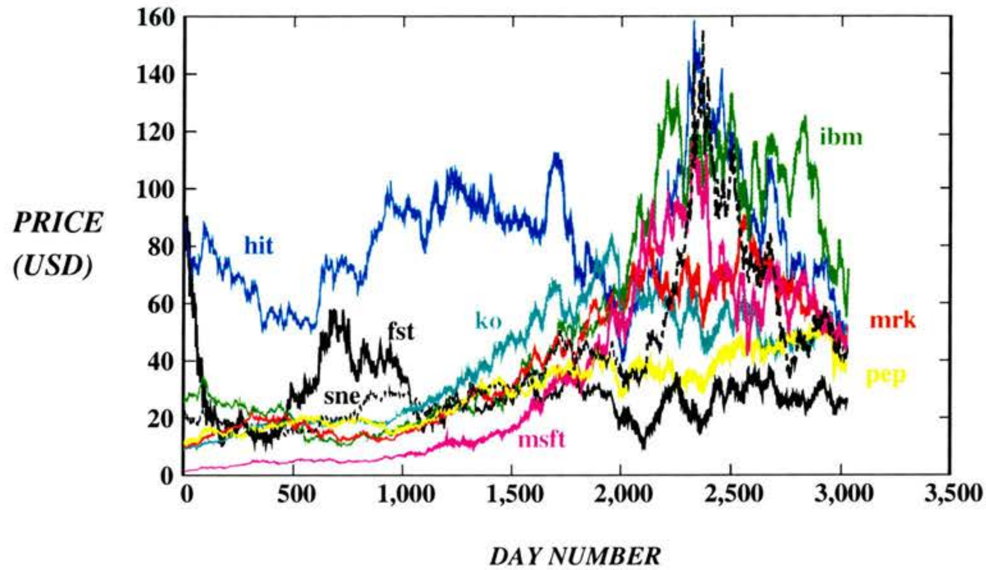

Figure 9.8: Close Price from Oct 17 1990 - Oct 17 2002 of Stocks from Table 9.12

9.3.2 Experiment: Stop Learning and Go On

The aim of this set of experiments is to find out if the artificial traders are able to adapt to new information, and indeed, if this adaptation is useful in such a fast changing environment such as the stock market. This is investigated by analysing the evolution of rules. Tests are conducted first with continual learning (GA-On) during the first n number of years of the entire data set; after this period, learning through the GA stops for the remaining period of the data $(totalPeriod - n)$. The only adaptation process happening when the GA is off is the apportionment of credit algorithm, which as explained in section Section 3.3, updates the "goodness" of existing rules, without adding any new rules. Results of these experiments are compared against a run with continual learning during 100% of the data.

First, let's look at the general trends. Table 9.13 shows the final wealth of the bank at a compounded interest rate of 8% annually, buy-and-hold and trend-following strategies when investing initially $10,000 US Dollars. As it can be seen, in the first two stocks, their *buy-and-hold* investments give a return which is less than the original investment of $10,000 US Dollars, and the first four under-perform the *bank* investment over the period, while the last four outperform both, the initial investment, and the *bank* ($msft$ is the top outperformer and fst the top under-performer of the list). The stocks belonging to the first set of four will be called *under-performers*, while the ones from the second half *over-performers*.

Table 9.13: End of Period Wealth of Bank Investment (B), Buy-and-Hold (BH) and Trend-Following (TF) Strategies with an Initial Investment of $10,000 US Dollars.

Wealth	FST	HIT	IBM	SNE	KO	MRK	MSFT	PEP
B	26,068	26,060	26,068	26,068	26,060	26,068	26,068	26,060
BH	3,277	5,213	25,800	19,049	51,459	49,552	320,759	39,543
TF	42	9,469	8,973	17,131	16,341	17,125	18,395	2,503

9.3.2.1 Under-performers

The following experiments were carried out with rewards given by *Specificity* and *Holding* at 5% (explained in Section 9.1.2). There are a total of 1,001 runs performed, all with wildcard probability $P_{\#} = 0.5$ (see Section 7.11 for descriptions of parameters, and Section 9.1.2 for parameter values). *GA period* used is 100 in all cases that follow (GA acts every 100 trading days).

Table 9.14 shows the average of the final wealth of traders-type 1, 2 & 3 over 1001 runs of fst with GA active from start to year number shown in column 2, equivalent to day number shown in column 3. For instance, the first row indicates final wealth in the absence of the GA (Ga-Off) from beginning to end. The last row indicates the opposite case, where the GA is on (every 100 days) throughout. The rows in between show the performance when the GA stops after 2, 4, 6, 8, 10, and 11 years of data. As it can be seen, it seems that the effect of the GA is slightly detrimental in Traders 1 & 2; the more GA is used, the worst these traders perform (recall that this is the worst performer of all stocks shown in Table 9.13). However, Tt3 seems to do better on average as the GA rate increases.

Table 9.14: FST Stock. Average Final Wealth of Traders 1, 2 & 3 over 1001 runs with GA-On from start until year number shown in column 2, equivalent to day number shown in column 3.

GA-On until Year	Day Number	Tt1	Tt2	Tt3
1	1	6,369	7,235	9,715
2	503	6,275	6,063	10,294
4	1,009	6,430	5,644	11,075
6	1,515	6,132	5,399	10,983
8	2,021	5,900	5,367	11,030
10	2,527	5,815	5,291	11,085
11	2,780	5,804	5,239	11,059
12	3,027	5,798	5,209	11,043

Looking at the best performers, this observation persists. Table 9.15 shows a comparison of the average performance (as shown in Table 9.14) against that of the best performers of the group. Again, the best of Tt1 and Tt2 do better with the GA-off, but Tt3 improves with GA-on. As demonstrated in Section 9.1.4, every type evolves and performs differently, it all depends on the inputs they receive. In this case, Tt3 seems to capture and learn better market strategies with the information given to it, and clearly the use of the GA helps to inject more new and successful strategies.

Table 9.15: FST Stock. End of Period Wealth of Best Trader with GA-On versus GA-Off throughout the 12 years

Final Wealth	Tt1	Tt2	Tt3
Average GA On	5,798	5,209	11,043
Average GA Off	6,369	7,235	9,715
Best GA On	34,253	67,684	61,180
Best GA Off	49,867	92,155	38,062

Now let's take a look at the other three stocks of the under-performers group. Tables 9.16, 9.17 and 9.18 show the final wealth of the average of these three types of traders. With Hitachi stock, again Tt1 performs worse with GA-on. However, in this case, Tt2 and Tt3 improve as the GA increases.

Table 9.16: Average Final Wealth of *hit* over 1001 runs with GA active until year number shown in column 2, equivalent to day number shown in column 3

GA-On until Year	Day Number	Tt1	Tt2	Tt3
1	1	15,717	12,074	12,643
2	503	14,419	12,267	14,400
4	1,009	11,817	11,965	14,512
6	1,515	10,149	12,232	14,397
8	2,021	9,115	12,376	14,377
10	2,527	9,036	12,613	14,172
11	2,780	9,061	12,308	14,137
12	3,027	8,951	12,122	14,114

The effects of the GA in the stock *ibm* are consistent with those found in *hit*. Only Tt2 and Tt3 improve as the GA-rate increases. And finally, with SNE stock only Tt3 improves performance.

Table 9.17: Average Final Wealth of *ibm* over 1001 runs with GA active until year number shown in column 2, equivalent to day number shown in column 3

GA-On until Year	Day Number	Tt1	Tt2	Tt3
1	1	37,707	26,077	26,114
2	503	38,583	28,181	25,374
4	1,009	39,118	30,641	26,221
6	1,515	38,529	30,622	26,382
8	2,021	36,963	31,195	25,797
10	2,527	36,961	32,689	26,312
11	2,780	36,728	32,924	26,287
12	3,027	36,934	32,859	26,376

Table 9.18: Average Final Wealth of *sne* over 1001 runs with GA active until year number shown in column 2, equivalent to day number shown in column 3

GA-On until Year	Day Number	Tt1	Tt2	Tt3
1	1	30,491	25,557	22,190
2	503	32,563	25,709	22,637
4	1,009	28,513	23,104	22,543
6	1,515	25,708	21,868	21,583
8	2,021	23,444	21,703	22,100
10	2,527	23,843	22,825	22,713
11	2,780	23,842	22,726	22,952
12	3,027	23,823	22,740	22,880

9.3.2.2 Over-performers

Similarly to the previous section where the results of the worst performer were presented first, in this section *msft*, as the best of the four outperformers of this group is presented first. Table 9.19 shows the average of final wealth of Tt1-Tt3 with the various GA invocations shown in each row. As it can be observed in the table, all traders improve their performance as the GA increases.

Table 9.19: MSFT Stock. Average Final Wealth of Traders 1, 2 & 3 over 1001 runs with GA-On from start until year number shown in column 2, equivalent to day number shown in column 3.

GA-On until Year	Day Number	Tt1	Tt2	Tt3
1	1	166,684	110,437	211,451
2	503	196,330	116,616	225,041
4	1,009	199,980	129,928	233,098
6	1,515	196,966	130,636	232,758
8	2,021	192,533	130,011	234,044
10	2,527	196,057	130,713	237,503
11	2,780	196,268	130,957	239,465
12	3,027	196,321	131,440	239,786

Then, results of the remaining three stocks of this group are presented in Tables 9.20, 9.21 and 9.22. Again, results are consistent with *msft*. All traders improve performance as the GA rate increases.

Table 9.20: Average Final Wealth of *ko* over 1001 runs with GA active until year number shown in column 2, equivalent to day number shown in column 3

GA-On until Year	Day Number	Tt1	Tt2	Tt3
1	1	35,697	36,243	40,787
2	503	39,954	36,550	43,262
4	1,009	40,605	36,819	43,410
6	1,515	41,129	37,986	43,491
8	2,021	41,351	39,712	44,188
10	2,527	41,325	39,814	44,884
11	2,780	41,343	39,791	44,873
12	3,027	41,292	39,841	44,942

Table 9.21: Average Final Wealth of *mrk* over 1001 runs with GA active until year number shown in column 2, equivalent to day number shown in column 3

GA-On until Year	Day Number	Tt1	Tt2	Tt3
1	1	32,497	35,767	37,288
2	503	37,109	37,026	40,281
4	1,009	38,939	38,073	40,734
6	1,515	40,357	37,131	40,897
8	2,021	40,423	37,916	40,942
10	2,527	40,554	39,303	41,318
11	2,780	40,671	39,528	41,733
12	3,027	40,666	39,552	41,863

Table 9.22: Average Final Wealth of *pep* over 1001 runs with GA active until year number shown in column 2, equivalent to day number shown in column 3

GA-On until Year	Day Number	Tt1	Tt2	Tt3
1	1	20,436	26,376	23,643
2	503	24,495	28,207	27,041
4	1,009	27,317	29,677	29,064
6	1,515	28,418	29,869	29,585
8	2,021	28,586	30,023	29,686
10	2,527	28,525	29,979	29,576
11	2,780	28,501	30,069	29,565
12	3,027	28,542	30,060	29,606

9.3.2.3 Analysing Adaptability further with FST Stock

The GA was favourable in all over-performer stocks, but not in all the under-performers, where the results were worst in the case of fst. In order to assess whether the agents *truly adapt* even when performance seems to deteriorate, we now take a closer look at the process of adaptation with the GA-on versus the GA-off. Figure 9.9 shows how the stock performs with the usual strategies.

As shown in Section 9.3.2.1, the performance with the GA-off is significantly higher than that obtained with the GA-on in fst stock. One way to analyse if the evolved rules are adapting over time or not is to examine their *accuracy* from day one to the last day. If the rule-set is not adapting, one would expect that the accuracy would remain more or less constant throughout the entire period. However, if accuracy increases, this clearly shows that the evolved rule-set is *truly adapting* to new environments.

Figure 9.10 presents the accuracy of the evolved rule-sets of the best traders in the case where there was no GA and where there was GA every 100 days from beginning to end. The results of this graph are very encouraging: the accuracy with Ga-off starts, as expected, at around 0.5 (one good decision, one bad one). It then increases to around

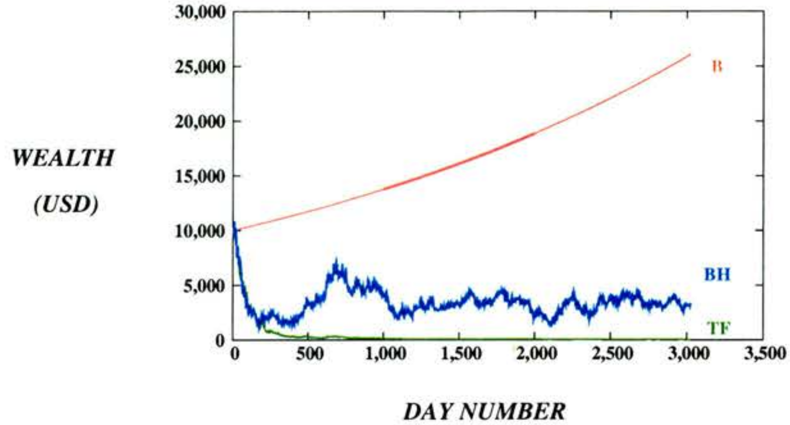

Figure 9.9: FST Stock. Wealth of Bank (B), Buy-and-Hold (BH) and Trend-Following (TF) Strategies

0.75 due to the apportionment of credit algorithm, this is the lower level of learning discussed in Section 3.4.1.1. However, after a short period of less than 500 days, its accuracy shows no further improvements. When the GA is on, again accuracy starts at 0.5, but steadily increases until it almost reaches 1. Therefore although the agent shows less wealth when the GA is off, the evolved rule-sets are far more accurate.

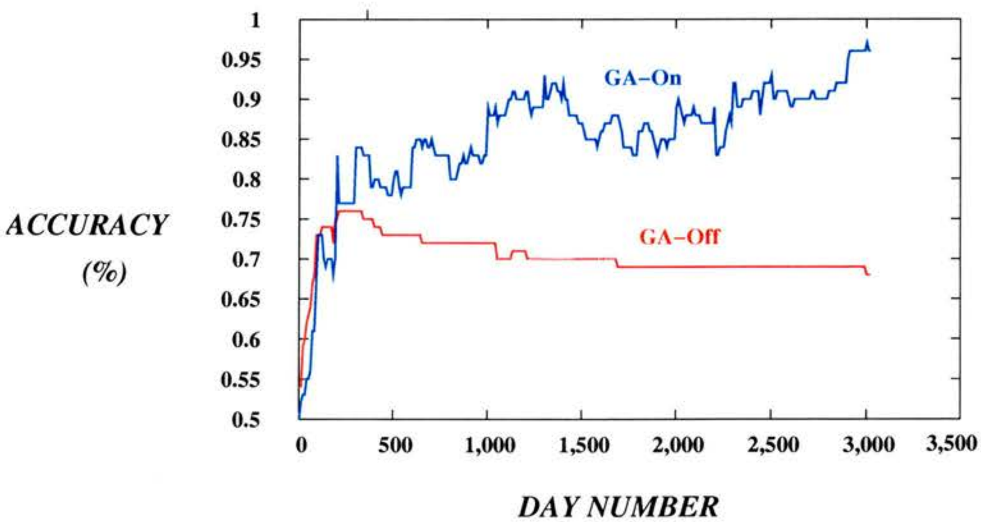

Figure 9.10: FST Stock. Accuracy of Evolved Rule Set of Best Runs with GA-On vs GA-Off

9.4 General Discussion

Finally, this Section presents a general discussion of the themes addressed throughout this thesis. It has been found that it certainly seems more defensible to try to model the behaviour of a group of traders than to model the behaviour of a single trader, or to try to predict the future of the time series that represents a stock's price, whether by local linear methods, by neural net methods or by other nonlinear techniques. Time series prediction methods generally model outcomes rather than causes, and their predictions can be invalidated if the underlying causes change. Modelling genuinely different *trader types* is at least an attempt to model certain of the causes of prices.

So what purpose could this serve? Such models could be used to actually trade, or to assist human traders by making suggestions, or even as training aids against which novice traders could compete. Further research and development of this model could help both, to gather more evidence to support it and to explore ideas such as creating a portfolio management system involving a greater number of *trader types*, each group being good at dealing in one stock, together with an arbitrator system to apportion resources dynamically between them.

Before continuing to develop this discussion, some interesting properties one can learn and investigate further from a model of this type should be stressed:

- In order for the artificial agent to survive, it must adapt to market dynamics which are not affected directly or indirectly by the agent's actions, but rather by real traders and people's expectations in the stock market. In this sense the adaptive agent can be viewed as trying to **mimic real behaviours** of which our understanding at present times is very limited. The causes of such complex behaviours are unknown to us and any attempts towards increasing our knowledge in this area are greatly valued.

- Also at the agent level, it has been shown with these results that it is possible to develop strategies for trading successfully with highly summarised and easy to process market information. It might be possible to transfer experience across different stocks in search for some universal strategies as well as testing specific rules that apply to a stock across a finite length of time

- The learning process and rule complexity are issues that must be explored more deeply. It appears that according to the agent's profitability, it might be better to have a good mix of rules, including specific ones and of higher complexity (with very few wild-card symbols). More analysis needs to be done regarding default hierarchy formation and whether it would be desirable to obtain it at all times during the simulation. In order to control hierarchy formation in the model, a special reinforcement procedure has been implemented, according to specificity.

- The model allows us to learn about three important aspects: (i) market data, what items of information are more important than others or whether they are important at all; (ii) the evolutionary process of adaptive agents of these types; and finally (iii) about the environmental reward, whether the reward function is affecting profitability in a positive way or not. A better understanding of these aspects can guide us in the design of better trader models. For example, there is considerable scope both for experimenting with the mixture of rule conditions and for seeing whether some form of hill-climbing might manage to improve the rules even further.

- Technical trading is a valid outcome in this model, but it appears to be more relevant at some times than at others. The more technical type of trader (Tt1) does not always outperform other non-technical traders. There is the possibility of creating a "super trader" which gets its investment signals from whichever trader is performing better than the rest. In this way the super trader actions can be guided by traders who are active and are performing really well. When the performance of the best trader starts to decrease, the super trader can adapt its actions by doing what a new best performer is doing.

Finally, it is important to stress that this approach is not capable of evolving agents that can spot the right moment to invest heavily in order to make a quick "killing" and then exit. The agents have a trading opportunity only once per day, rather than trading in some event-driven way. The approach as shown in this Chapter is therefore more suited to tasks such as portfolio management than the pursuit of quick profits. The use of intra-daily data could be alternative method for short-time investments. However, it is important to mention that in no cases did the agents lost all their initial wealth in any of the runs presented in this thesis. The great majority of cases show a clear return, much higher than the *bank's* investment.

Chapter 10

Conclusions

"There are no facts, only interpretations." Friedrich Nietzsche, *Nachlass*, translated by A. Danto.

The experiments presented throughout this thesis have shown that the model is capable of producing sensible results. It is understood that this economic model is limited in the sense that it captures only a small fraction of the whole repertoire of a trader's behaviour. The total number of strategies analysed by real traders in the stock markets is, indeed, far greater and much more complex than the simple market models (i.e. *trader-types*) used here. Each one of these, in turn, represents only a small subsample of all possible strategies used by real traders. However, this property should not necessarily be considered a limitation of the model, it was designed with this purpose in mind; simplicity is, perhaps, what contributes to the success of its results.

The following sections intend to assess what can a model of this type offer to the three main areas it mainly resides in: as an artificial stock market, an automated trading system and a decision-making model. Then, important issues will be brought forward, such as what sort of things can be learnt about every one of its components, as well as what are its possible applications and future work.

10.1 As an ASM

Artificial Stock Markets try to reproduce qualitative features of the real market, and some, which have already been mentioned throughout this thesis, have been very successful at doing so. But in this quest of knowledge about market dynamics, very little we have learned about the agent per se. The dynamics of *the agents* are still unknown and there is no consensus or standards established as to how exactly these agents adapt, what type of models they develop, test and improve upon. Are these market hypothesis meaningful and consistent throughout different markets? Can they keep improving or they reach a plateau and become less and less susceptible to market changes? In this model, by exploring the evolved pool of strategies even further, we can search for better ways of improving these agents. Although not all possibilities have been analysed in this work and perhaps many questions were still left unanswered, as noted in the introductory chapter, the purpose of providing with a number of alternative answers has been achieved. The intention of contributing to increase the interest and to open more ways of exploring and expanding our current knowledge of how these agents operate and interact under different market conditions, how such cognitive process can be improved, etc. has been met.

The question that still strikes me regarding artificial stock markets is the following: if the generated data from such models looks so similar to real financial markets, how useful can this be to predict *a particular* stock's future behaviour? By looking at different time series (e.g. like Forest Oil vs. Microsoft), one can infer that it would be extremely difficult to use the generated data of an ASM to try to predict the outcome of a specific stock from the real world. Here, the intention to use real market dynamics was to avoid a departure from reality, focusing only on the agent's dynamics. The hope is that better trader models will be designed which in turn, will provide us with more insights about how complex systems operate under uncertain environments.

In ASM models, trader strategies are difficult – or even impossible – to analyse. Regarding this matter, Chen and Chia-Hsuan Yeh in their paper [Chen & Yeh 99] (analysed in section 4.2.2), point out that because only about 40% of the traders benefit from the business school per day, it would be "interesting to know what kind of useful lessons traders learn form business school.... What motivates them to search and helps them to survive is in effect *brief signals*... Another way to see what traders may learn form business school is to examine the forecasting models they employ. However, this is a very large database, and is difficult to deal with directly. But, since all forecasting trees are in the format of LISP trees, we can at least ask *how complex these forecasting models are*." Then they define this complexity measure according to the number of nodes and the depth of the trees evolved. As it can be appreciated here, it can be a difficult task to try to understand trader's strategies and the tendency will be to try to avoid it by defining other types of measures that could give an approximation of what is going on inside the trader. In the same way, the SFI market analyses strategy complexity (a measure of the rule's specificity) instead of agent's strategies. In this model, traders strategies are easier to handle and, in terms of computational cost, can be evolved very cheaply.

As pointed out earlier, ASM tend to analyse the consequences of the agents' behaviours and then deduce, by the use of economic theories, an approximation of what is going on inside the artificial trader, rather than studying the dynamics of the actual strategies. The simplicity of this model makes it possible to analyse such market strategies and it is precisely this reason that motivated the idea of using real data. These issues have been addressed in the model description and results chapters.

10.2 As an Automated Trading System and a Decision-Making Tool

Even though it might seem like a very ambitious idea at first, an on-line learning system of this type shows potential to be used as a fully automated trading tool for investment decisions in financial markets. However, this potential use brings the assertion of whether human traders could eventually be replaced by modern systems, which, indeed, is a delicate one and needs to be considered with care. An interesting report entitled *Views from the Frontier: Commentary on the New World of Forecasting and Risk Management* from Olsen and Associates, cites examples where autonomous machine-based systems are already managing large amounts of capital, with the idea that a "number of people believe that human traders, with their limited information processing powers and susceptibility to emotion and unscientific ideas, are too fallible for modern markets." [Ols96]. According to this report, Andrew Lo believes that autonomous trading systems are only feasible in markets that are largely populated by human traders. Lo adds: "you will never be able to replace human interaction until you reach the point where machines become self-aware, and we might never achieve that lofty goal... the aim of the new generation of systems is to augment, rather than replace, human intelligence." This last statement brings the proposed system to another possible use: a decision-making tool.

In general, most AI systems in finance still require a great deal of heuristic knowledge. Such knowledge is provided by financial experts in the form of rules, knowledge bases, data selection and proper manipulation, which strongly suggests that they are best used to complement the decision process of existing team of experts rather than on their own. "Thus they exist more in the realm of statistical tools than 'artificially intelligent' agents. Nevertheless, they are powerful techniques and as our development of them progresses, it is likely that they will find greater and greater utilization on Wall Street" [Gilbert 95]. Such is the case of the system AXON and many others largely employed. However, even if a successful, completely automated trading system still looks far from present times, very good approximations have been deployed successfully (e.g. Prediction Company's fully automated traders), and this is the idea behind this model, to be able to act without human intervention.

So far, the model works 100% autonomously, i.e. learns and adapts without any human intervention nor any initial clues given to it, and results shown with no prior knowledge about the market are encouraging. However, one might ask if it would be worthwhile to use in this system knowledge from the *real experts?* The answer might be quite questionable.

The goal of many decision-making systems is to try to replicate the knowledge (heuristics) of a human in a particular field of expertise. The principle behind these systems is that, in an ideal situation, the human expert's wisdom can be reduced to a series of interconnected generalised rules. However, as it has been mentioned throughout this thesis, there are distinct limitations to the abilities of experts in any field to articulate the rules they follow when displaying their expertise. The expert may have no conscious access to much of his expertise. These systems are well suited for problems where there is a consensus of expertise and where there is value of retaining such expertise in a system, such as medical diagnosis, but applying them to trading, one can clearly see that there is no consensus on theories and as a result, many different views are applied.

For instance, neural networks may be more successful than expert systems when rules underlying decisions are not well understood; they may also prove better in problems where numerous examples of decisions are available and in problems where a large number of attributes describe inputs; and more advantageous in problems where uncertainty exists about the rules governing a "good" decision. However, a number of disadvantages have been addressed by NN practitioners when dealing with these types of problems. In these cases, other approaches, such as LCS and GP have been more successful in tackling the drawbacks of NN models in use. In particular, this thesis emphasises the use of LCS due to the fact that they are very powerful models of cognitive processes.

10.3 Possible Applications

Other application domains of the present model include a wide variety of areas where there is time series data available, decisions need to be considered frequently and there is some kind of feedback possible about the quality of decisions. For example, there might be applications in the insurance market, in credit scoring, in fraud detection, in marketing prospect assessment and so on. In particular, the foreign exchange market is a very good candidate to test the model. In addition, this approach could be used in various ways: to automate some dealing, to provide a benchmark for use in developing more knowledge-intensive dealing systems, as a training aid for new dealers and as part of a portfolio management system.

As a portfolio management. Each *trader type* deals with one stock and a bank, or one risky stock and one baseline stock deemed safe, that plays the role of the bank. A set of successfully evolved *trader types*, one per stock, could then be used for portfolio management if one also evolved an arbitration rule-set that decided which *trader types* would be allowed to buy at any given moment – presumably on the basis of recent performance, much as happens in other classifier systems. One would want such a portfolio to outperform a general market indicator such as FTSE 100 reliably over time and by a consistently healthy margin. For example, Donath et al [Donath *et al.* 99] have tried training neural nets to arbitrate between the decisions offered by other trained nets, with some success.

10.4 What can we learn from this model?

This section summarises brief answers to some typical questions and observations that have been addressed througouth this thesis regarding the traders, environment, reinforcement and discovery components, and finally about the LCS used in this model in general.

10.4.1 About the Traders

1. What do these agent types know? Initially, nothing, but as time progresses, they develop better sets of strategies to trade upon.

2. What do these agent types believe? One could argue that they believe in the EMH. This is a single-step environment, there is no link in rewards given for good actions in the future. In other words, the LCS does not reward chains as it is the assumption of this model that prices follow a random walk, the price of today does not depend on the price sequence prior to today. A good action today is rewarded tomorrow only if it was successful at meeting certain criteria.

3. What do they learn? The agents should be capable of adapting reasonably well to almost any type of environment. In this thesis the *stock markets case* was addressed, but there are a number of other possible applications where this model could be used. For example, it is well suited for real world applications where historical data in the form of time series is available such as in insurance, credit rating and fraud, foreign exchange markets, etc. Other applications where it can

have good potential are in the development of systems for intelligent homes and to design better transport systems, to name a few.

4. How fast do they learn? The learning parameter is controlled by the designer, and which exactly is the best value depends on the stock being analysed. For certain stocks, faster learning rates are proven to be better than slower ones. The *GA_period* parameter that best performed was about 100 in most cases. Smaller than 100 injects more strategies too quickly, without giving the agent a chance to test these models before they are exchanged by new ones again. It would be very interesting to test a self-adjusted *GA_period* parameter.

5. What exactly is the forecasting model the agent uses? Agents try to maximise their profits, but they do it in different ways: agents working on the same stock usually differ in their pool of strategies, even when they are of the same type. Exactly *how much* they differ does not seem to be relevant and was therefore not measured. This difference can be explained because as they learn through experience, their learning experiences can follow very different paths from each other; this property is common in human traders, which is exactly the target to model. However, consistency has been found (shown in Section 9.1.5), suggesting that there are patterns that are actually learnt and exploited by the adaptive traders.

6. Can agents start with prior market knowledge? Yes. Another interesting issue to explore would be to allow the agents to share their most profitable (accurate) strategies. It is strongly believed that this would increase performance dramatically, because the agents would be acting in a cooperative manner – as if they were part of a team. As noted in Section 7.1, obtaining actual market expertise is difficult and costly. However, if some expert knowledge was available, one would only need to add it in the initial rule-set rather than having a purely random start (the same applies for rules derived from other trader-types). In this model, because of the continual learning process proposed, peforming a large number of runs is very cheap and effective (see, for example, the percentage of runs outperforming the buy-and-hold strategy shown in Chapter). For these rea-

sons the model starts with no previously acquired knowledge given either by the experts or transferred by other adaptive traders. Also, other types of interactions that can be analysed are through agents dealing with different stocks.

7. How sophisticated are the traders? Again, as with the learning parameter, it is up to the designer to define how sophisticated these agents are. For example, an agent with a condition length of 70-80 bits can be made (such as the first versions of the SFI model), depending on the problem being analysed. Another type of sophistication can be measured through the specificity parameter associated to every rule, i.e. the more specific and accurate the rules are the better trading models are developed.

8. Not all runs produce any *trader types* that perform well. How much effort does it take to produce a good agent through repeated runs? Because there is no training involved, i.e. the data is only seen once, these models can be evolved very quickly, taking only a few seconds to run the system from a pure random start, with 10 years of data. It has been shown that in all cases where 1001 runs were made, a reasonable number of models were found to consistently beat the *buy-and-hold* strategy. This number of winners is, of course, stock dependent. Stocks like Forest Oil were easier to beat than other such as Microsoft.

9. Is the agent stock-dependent? Yes, agents are very much dependent upon the characteristics of the stock in question. As described throughout this thesis, the rules of an LCS contain two parts, the condition and the action. The condition matches the environmental states, therefore the sets of evolved rules have the market characteristics embedded in their condition part.

10.4.2 About the Environment

Regarding the stock analysed, a number of important issues have been raised while reporting the results of the model. In particular, the question of whether this approach would perform well when applied to a stock that tended to drop in price over the years was illustrated in section 8.5.3 with the stock HNS from the UK, and Forest Oil from the US in section 8.6.3. The aim was to investigate whether the agents would still be able to make money by buying in early in any upswing, even though those are rare in stock that perform poorly. In both cases the agents easily outperformed the *buy-and-hold*, but how would that compare with the *bank* strategy?

The *bank* in these cases is the toughest strategy to beat. That is not difficult to understand. The agents only have the option of leaving their money in the bank or buy a stock that is (almost) constantly decreasing in value. A purchase of the stock will almost certainly cause a drop in wealth. However, the best performers managed to do better than *both* strategies. Why?

By observing at their actions, these agents learned to buy previous to price increases and sell prior to price drops. They managed to profit from such transactions even though the stock is such a bad investment option during most times.

How would this approach perform when applied to a more temperamental stock, for which neither the *buy-and-hold* nor the *bank* strategy would be particularly effective? This is clearly the case of Cabletron Systems, described in section 8.6.2. Trader-types such as Tt2 are well suited for this task due to the market conditions they are built to respond to. In particular, this agent is sensitive to volume changes as well as price changes. Cabletron is a highly volatile stock and it appears that Tt2 captures and exploits the trends accurately.

Regarding information sets, it seems interesting to explore whether more sophisticated information channels such as information about a particular index, news from message boards, a broad variety of ratios, etc., could help improve performance. A trader type could benefit from more, or maybe less, information. Using more factors provides more clues, but also multiplies the size of the search space so that the evolutionary process may come to be governed more by neutral genetic drift than by genuine selective pressure. Other sorts of information, such as external political or economic conditions, might be introduced in a simplistic way by, say, adding the behaviour of the FTSE 100 (or similar gross indicators) as an extra factor. In addition to these, another interesting thing to look at, is whether a good evolved rule set can be further improved, by rule pruning, by rule aggregation or even by hill climbing.

Another aspect that could be improved is that the stock prices used in this thesis are *closing prices*, which are not necessarily the *true prices* at which one could buy it at present or in the near future. At any given instant in time, there is a best or highest "bid" price from a potential buyer and in the same way there is a best or lowest "ask" price from someone who wants to sell the stock. These facts could account in the returns reported here. Also, more frequent trading decision points could be achieved through the use of high frequency data. This would allow the system, rather than trading at most once per day, and using only closing prices, use continual information and either have several decision points per day, or make it event-driven in some way.

10.4.3 About the Reinforcement Component

The level of complexity can be tuned by the reward function used. In this thesis a number of reward schemes were proposed and tested in order to guide the learning process. The first one that was tried is the one proposed in [Goldberg 89], but the problem of over-generals arouse because it does not add pressure to more specific classifiers and the "guessers" had more chances to be reproduced, thus over taking the most specific ones in the population. To overcome this problem, a special reward function was developed here and it is the most commonly used throughout this thesis due to its higher performance over the others. Basically, it is a function of reward according to specificity. Others tested were the rewards according to accuracy, a measure that is the ratio of correct actions divided by the total number of matches. The conclusion drawn with respect to reward schemes is that the best reward function depends on the type of trader. One function might work better for one trader than for another.

The best suggestion regarding this, is to use a variety of reward schemes with the purpose of building new sets of strategies. To illustrate this, consider teaching a child to cross the street, i.e, *not* to cross it when a car is approaching and the light signalling walk is green. Unfortunately, many children learn it the hard way, by being hit by a car. Yet still others also learn it without experiencing any dangerous situations. Both strategies are correct, except that they were learnt in different ways. While one boy can have only one leg thanks to the mistake he made while learning his lesson, another might be complete. But that does not change the fact that now they both know how to cross the street, they are as good and as reliable as any others.

The fact that some traders are bad performers does not necessarily mean they have bad sets of strategies. Consider the following: Two traders, one ending the 10-year period with only £25,000, another ending with £250,000. Huge difference in profits! However, the set of strategies of one can be no worse than the other. It only means that one of them made a transaction that was influential in a more negative way than the other. The final recommendation would be to build a set of strategies from various different scenarios.

10.4.4 About the Discovery Component

The rules are evolved by running a genetic algorithm every so often, to produce new rules by selection, recombination and mutation and then replace 20% of existing rules by crowding, but more modern GAs might do better. Other niching mechanisms can be explored, including sharing, and the implicit niching of XCS.

Clearly, it matters how often the GA is invoked. If the GA is run too rarely, progress towards good rules will be slow; if it is run too often, there won't be time for rules to demonstrate their proper value. For example, a certain rule might apply only on rare occasions but be crucial when it does. That rule only gets the chance to earn its keep when that rare situation arises. If the GA runs before that time, the potentially valuable rule will be an obvious candidate for replacement because it has not yet done anything, or perhaps anything particularly useful. This is an ongoing topic of research into LCSs which raises two fundamental issues: first, how can an LCS be made to behave fairly towards rarely-applicable yet valuable rules; and second, are such rules genuinely irreplaceable? Consider a car-driving rule system. A rule which says "if about to hit a wall, slam on the brakes" is very useful but (hopefully) applies very rarely; but is it needed at all, or is it possible to produce a rule system that ensures that such a situation will never even arise?

It seems like a good idea to make the interval between GA runs be *self-adaptive*. The current work has experimented various fixed intervals. If a GA runs too often, good rules may not have the time to prove their proper worth; if too rarely, the system may fail to learn enough or may fail to detect changing market circumstances quickly enough.

10.4.5 LCS

The results presented here strongly suggest that it is feasible to model a trader's behaviour in such a simplistic way and that LCS can be a good tool to do so. However, as noted earlier, the current work uses the most simple form of classifier system and GA, which is run at fixed intervals. More modern classifier systems such as XCS (the eXtended Classifier System) or ACS (the Anticipatory Classifier System) could handle credit assignment better and could be fairer to rules which can only fire rarely but do well. There may also be as yet undiscovered improvements to be made to the classifier system, or perhaps, other models, such as XCSR, which uses 'interval predicates' $a \leq f_i \leq b$, so that real numbers a and b appear in the rule conditions rather than simple binary factors; XCSL (XCS with LISP s-expressions); XCSM (adding Memory to represent internal states) and the Corporate eXtended Classifier System (CXCS) could be explored even further. (A good summary of all these can be found in [Lanzi *et al.* 00].)

10.5 Can We Tell Why It Works?

A single reason can not be given, but here are some simple suggestions that might be acting together that could explain the success of this system. These are organised into different sections:

1. The simplistic design of the *real economic environment*, due to the fact that it deals with real data along with the use of simple Moving Averages as opposed to other derived MA indicators (in the form of ratios or comparisons) such as oscillators, MACD, adaptive moving averages, etc., which could add more complexity and perhaps the loss of some information.

2. At the *rule discovery component* of the trader-types, one concern is that the "lack of separation between global solutions can be a problem to a niched GA, because of a given population size, there is a lower limit to the distance between global solutions that the GA can discriminate. In other words, below some critical distance, two different solutions are considered to be in the same niche. In the case of a classifier system, if both of these rules are required to solve a particular

problem, over time the GA will be unable to maintain both rules in the population stably" [Goldberg *et al.* 92]. However, this problem domain might not suffer from insufficient separation between strategies in a default hierarchy, and the creation and maintenance of several niches could be indeed plausible in this model. Other factors that make a problem hard for a classifier system are those where global optima could be improperly predicted by low-order schemata and problems that have many locally optimal rules. Again, these do not appear to affect negatively the prediction task of these trader-types, as they are able to find and exploit some market inefficiencies relatively successfully. Recall that the aim here is not searching for optimal behaviour.

3. Or it might have to do with the *continual learning* approach proposed. This system offers perhaps one of the fastest ways of exploring intelligently vast amounts of information, an imperative property which adds great computational complexity to any attempt to model modern financial markets. The idea supported throughout this thesis is that there should be no need to learn (or even worse, memorise) strategies by passing the data (experiencing it) more than once. In real life this hardly happens, the world is a very dynamic system where one learns to generalise with very limited examples and by trial and error. The argument is that circumstances never repeat exactly and therefore there is no need to repeatedly be faced with previous examples to be tested on future unknown states. I propose that the generalisation that one must try to achieve is one of a higher level, in a context that is made throughout the whole life-span of the agent.

4. *Modelling a group of prototypical traders* rather than a single agent could certainly be a cause of the success of this model. If a single trader was to be modelled, a length of his environmental message in the range of the hundreds would probably not even be enough to model all the factors affecting his decision. Not having the right breakfast one morning could have a great effect, causing him to miss an important factor for his next decision, which the artificial trader could not miss under these circumstances.

5. Not allowing the adaptive agents to change the real price dynamics, obligating them this way to adapt to it by anticipating real market trends rather than by directly causing them.

10.6 Future Work

As part of future work, in this model it would be useful to do further research in the following areas:

1. Evolve *trader types* for even more stocks.

2. Explore the possibility of allowing traders learn from each other. Enable communication between agents to test whether performance improves if *trader types* can act on what other types did last, as an extra factor. Is performance in a portfolio management situation improved if there is communication between *trader types* that use the same information set but are dealing with different stocks?

3. Reinforcement scheme. So far a trader type is rewarded for managing to increase its wealth. Would it help if a trader type that was performing poorly early on were instead rewarded for mimicking the better performers?

4. As a truly ASM. In work so far, *trader types* do not affect prices; their volume of business is regarded as too small to have a detectable impact on the market. However, it is plausible that if some type of trader who bases its decisions on certain factors does well, then others will seek to do the same and emulate its behaviour. In terms of *trader types*, this suggests that a successful trader type will come to have a bigger effect on the market. If *trader types* have variable-sized effects on the market and can affect prices, then does such a simulated market 'look realistic' in some way? For example, this would generate a novel time series of prices; does such a time series have similar characteristics to genuine price time series, in terms of phase portraits, dimensions etc, or even the ability to have crashes?

5. Perform comparison with other techniques. It would be interesting to do further research in other areas to asses, for instance, how well does this approach compare with other efforts to exploit market information to make purchasing decisions, such as time series analysis, neural net methods and genetic programming?

Bibliography

[Abu99] In Yaser S. Abu-Mostafa, Blake Lebaron, Andrew W. Lo, and Andreas S. Weigend, editors, *Computational Finance 1999*. MIT Press, 1999.

[ADT01a] Forming a View. *Applied Derivatives Trading*, Consulted May, 2001. In Beginners Guide.

[ADT01b] Risk Management 1. *Applied Derivatives Trading*, Consulted May, 2001. In Beginners Guide.

[ADT01c] Risk Management 2. *Applied Derivatives Trading*, Consulted May, 2001. In Beginners Guide.

[ADT01d] Technical Analysis 1. *Applied Derivatives Trading*, Consulted May, 2001. In Beginners Guide.

[ADT01e] Technical Analysis 2. *Applied Derivatives Trading*, Consulted May, 2001. In Beginners Guide.

[ADT01f] Trading a View. *Applied Derivatives Trading*, Consulted May, 2001. In Beginners Guide.

[ADT01g] What Determines Prices? *Applied Derivatives Trading*, Consulted May, 2001. In Beginners Guide.

[Ankenbrand & Tomassini 97] T. Ankenbrand and M. Tomassini. Agent Based Simulation of Multiple Financial Markets. *Neural Network World*, 4:397–405, May 1997.

[Arthur 92] W. Brian Arthur. On Learning and Adaptation in the Economy. Working Paper 92-07-038, Santa Fe Institute, 1992.

[Arthur 94a] W. Brian Arthur. Inductive Reasoning and Bounded Rationality. *American Economic Association Papers and Proceedings*, (84):406–411, 1994.

[Arthur 94b] W. Brian Arthur. The End of Certainty in Economics. In *Proceedings of the Einstein Meets Magritte Conference*, Free University of Brussels, July 1994.

[Arthur *et al.* 96] W. B. Arthur, J. H. Holland, B. LeBaron, R. Palmer, and P. Tayler. Asset pricing under endogenous expectations in an artificial stock market. Working Paper 96-12-093, Santa Fe Instituite, December 1996.

[Arthur *et al.* 97] W. B. Arthur, J. H. Holland, B. LeBaron, R. Palmer, and P. Tayler. Asset Pricing Under Endogenous Expectations in an Artificial Stock Market. In W. B. Arthur, S. Durlauf, and D. Lane, editors, *The Economy as an Evolving Complex System II*, pages 15–44. Addison-Wesley, Reading, MA, 1997.

[Baker 87] James Edward Baker. Reducing bias and inefficiency in the selection algorithm. In John J. Grefenstette, editor, *Genetic Algorithms and their Applications: Proc. of the Second Int. Conf. on Genetic Algorithms*, pages 14–21, Hillsdale, NJ, 1987. Lawrence Erlbaum Assoc.

[Banzhaf *et al.* 98] Wolfgang Banzhaf, Peter Nordin, Robert E. Keller, and Frank D. Francone. *Genetic Programming. An Introduction. On the Automatic Evolution of Computer Programs and Its Applications.* Morgan Kaufmann, dpunkt.verlag, 1998.

[Barr & Mani 94] Dean S. Barr and Ganesh Mani. Using Neural Nets to Manage Investments. *AI EXPERT*, pages 16–22, February 9th 1994.

[Barry 01] Alwyn Barry. XCSC. C source code, Faculty of Computer Science and Mathematics, University of the West of England, Bristol, United Kingdom, Consulted May, 2001. Available at `http://www.csm.uwe.ac.uk/ambarry/LCSWEB/computer.htm`, under XCSC, file xcsc.zip.

[Bass 99] Thomas A. Bass. *The Predictors. How a Band of Maverick Physicists Used Chaos Theory to Trade Their Way to a Fortune on Wall Street.* Henry Holt and Company LLC., New York, New York, 1999.

[Bauer 94] Richard J. Bauer. *Genetic Algorithms and Investment Strategies.* John Wiley & Sons, Inc., New York, 1994.

[Beltrametti *et al.* 97] Luca Beltrametti, Riccardo Fiorentini, Luigi Marengo, and Roberto Tamborini. A Learning-to-Forecast Experiment on the Foreign Exchange Market with a Classifier System. *Journal of Economic Dynamics and Control*, 21(8-9):1543–1575, 1997.

[Beltratti *et al.* 96] Andrea Beltratti, Sergio Margarita, and Pietro Terna. *Neural Networks for Economic and Financial Modelling.* International Thomson Computer Press, 1st edition, 1996.

[Bengio 96] Yoshua Bengio. Training a Neural Network with a Financial Criterion Rather than a Prediction Criterion. In *Proceedings of the Neural Networks in the Capital Markets Conference*, November 1996.

[Bettelheim 97] Eric C. Bettelheim. The Psychology of Traders. *Applied Derivatives Trading*, May 1997. In May'97 Feature.

[Booker 82] Lashon B. Booker. *Intelligent Behaviour as an Adaptation to the Task Environment*. Unpublished PhD thesis, University of Michigan, 1982.

[Box & Jenkins 76] George E. P. Box and G. M. Jenkins. *Time Series Analysis: Forecasting and Control*. Holden-Day, San Francisco, 1976. Revised Edition.

[Bulkowsky 97] Thomas Bulkowsky. The failed trade. *Stocks and Commodities*, April 1997. in Anatomy of a Trade.

[Butz & Wilson 00] Martin V. Butz and Stewart W. Wilson. An Algorithmic Description of XCS. Technical Report 2000017, Illinois Genetic Algorithms Laboratory, 2000.

[Butz 99] Martin V. Butz. An Implementation of the XCS classifier system in C. Technical Report 99021, The Illinois Genetic Algorithms Laboratory, 1999.

[Butz 00] Martin V. Butz. XCSJava 1.0: An Implementation of the XCS classifier system in Java. Technical Report 2000027, Illinois Genetic Algorithms Laboratory, 2000.

[Campbell *et al.* 97] J. Y. Campbell, Andrew W. Lo, and A. C. MacKinlay. *The Econometrics of Financial Markets*. Princeton University Press, 1997.

[Cavicchio 70] D. J. Cavicchio. *Adaptive Search Using Simulated Evolution*. Unpublished PhD thesis, The University of Michigan, 1970.

[Chan *et al.* 99] Nicholas Chan, Blake LeBaron, Andrew W. Lo, and Tomaso Poggio. Agent-Based Models of Financial Markets: A Comparison with Experimental Markets, September 1999.

[Chen & Yeh 96a] Shu-Heng Chen and Chia-Hsuan Yeh. Genetic Programming and the Efficient Market Hypothesis. In John R. Koza, David E. Goldberg, David B. Fogel, and Rick L. Riolo, editors, *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 45–53, Cambridge, MA, 1996. The MIT Press.

[Chen & Yeh 96b] Shu-Heng Chen and Chia-Hsuan Yeh. Toward a Computable Approach to the Efficient Market Hypothesis: An Application of Genetic Programming. *Journal of Economic Dynamics and Control*, 21(6):1043–1064, 1996.

[Chen & Yeh 97] Shu-Heng Chen and Chia-Hsuan Yeh. Modeling Speculators with Genetic Programming. In Peter J. Angeline, Robert G. Reynolds, John R. McDonnell, and Russ Eberhart, editors, *Proceedings of the Sixth Conference on Evolutionary Programming*, volume 1213 of *Lecture Notes in Computer Science*, pages 137–147, Indianapolis, Indiana, USA, 1997. Springer-Verlag.

[Chen & Yeh 99] Shu-Heng Chen and Chia-Hsuan Yeh. Genetic Programming in the Agent-Based Modeling of Stock Markets. In David A. Belsley and Christopher F. Baum, editors, *Proceedings of the Fifth International Conference on Computing in Economics and Finance*, Boston College, MA, USA, 1999.

[Chen 01] Shu-Heng Chen. On the Relevance of Genetic Programming in Evolutionay Economics. In K. Aruka, editor, *Evolutionary Controversy in Economics Towards a New Method in Preference of Trans Discipline*. Springer-Verlag, Tokyo, 2001. Forthcoming.

[Chen *et al.* 01] Shu-Heng Chen, Thomas Lux, and Michele Marchesi. Testing for Non-Linear Structure in an Artificial Financial Market. *Jurnal of Economic Behavior and Organization*, 2001. Forthcoming.

[Cox 94] Earl Cox. *The Fuzzy Systems Handbook.* Academic Press, Inc., Cambridge, MA, 1994.

[Dacorogna 93] Michel M. Dacorogna. The Main Ingredients of Simple Trading Models for Use in Genetic Algorithm Optimization in Financial Applications. Internal Document MMD.1993-03-22, Olsen & Associates Research Group, May 1993.

[de laMaza & Yuret 94] Michael de la Maza and Deniz Yuret. Experimenting with a Market Simulation. *The Magazine of Artificial Intelligence in Finance*, 1(3), 1994.

[de laMaza *et al.* 98] Michael de la Maza, Ayla Ogus, and Deniz Yuret. How do firms transition between monopoly and competitive behavior? An Agent-based Economic Model. In *Proceedings of the Sixth International Conference on Artificial Life*, pages 349–357, 1998.

[Deboeck 94] Guido J. Deboeck. *Trading on the Edge. Neural, Genetic and Fuzzy Systems for Chaotic Financial Markets.* John Willey & Sons, Inc., New York, NY, 1st edition, 1994.

[DeJong 75] Kenneth A. DeJong. *An Analysis of the Behavior of a Class of Genetic Adaptive Systems*. Unpublished PhD thesis, University of Michigan, Ann Arbor, MI, 1975. Dissertation Abstracts International 36(10), 5140B, University Microfilms Number 76-9381.

[DeJong 88] Kenneth A. DeJong. Learning with Genetic Algorithm: An Overview. *Machine Learning*, 3:121–138, 1988.

[Deppisch *et al.* 91] J. Deppisch, H. U. Bauer, and T. Geisel. Hierarchical Training of Neural Networks and Prediction of Chaotic Time Series. *Physics Letters A*, 158:57–63, August 1991.

[Donath *et al.* 99] Jürgen A. Donath, Thomas Frontzek, and Rolf Eckmiller. Towards Predicting Neural Net Control of Macro-Econometric Multi-Compartment Models. In *Proceedings of the Ninth International Conference on Artificial Neural Networks*, September 1999.

[Enrado 94] Patty Enrado. Bank on Neural Networks. Application Watch. *AI Expert*, 9(6):56, June 1994.

[Essinger 90] James Essinger. *Artificial Intelligence. Applications in Financial Trading and Investment Management*. Euromoney Publications Plc., 1st edition, 1990.

[Fama 65a] Eugene F. Fama. Random Walks in Stock Market Prices. Paper No. 16 in the series of Selected Papers of the Graduate School of Business, University of Chicago, 1965. Reprinted in the Financial Analysts Journal (September-October 1965), The Analyst Journal, London (1966), The Institutional Investor, 1968.

[Fama 65b] Eugene F. Fama. The Behavior of Stock Market Prices. *Journal of Business*, 38:34–105, January 1965.

[Fama 70] Eugene F. Fama. Efficient Capital Markets: A Review of Theory and Empirical Work. *Journal of Finance*, 25:383–417, May 1970.

[Fama 91] Eugene F. Fama. Efficient Capital Markets: II. *Journal of Finance*, 46(5):1575–1617, December 1991.

[Fama 98] Eugene F. Fama. Market Efficiency, Long-Term Returns, and Behavioral Finance. *Journal of Financial Economics*, 49:283–306, September 1998. Winner of the Fama DFA Prize for the best asset pricing paper in the journal during 1998.

[Farmer & Joshi 00] J. Doyne Farmer and Shareen Joshi. The price dynamics of common trading strategies. *Journal of Economic Behavior and Organization*, October 2000. Forthcoming.

[Farmer & Lo 99] J. Doyne Farmer and Andrew W. Lo. Frontiers of Finance: Evolution and Efficient Markets. Working Paper No. 99-06-039, Santa Fe Institute, 1999.

[Farmer 98] J. Doyne Farmer. Market Force, Ecology, and Evolution. Working Paper No. 98-12-117 E, Santa Fe Institute, 1998. Version 4.1 submitted to the *Journal of Economic Behavior and Organization*.

[Farmer 99] J. Doyne Farmer. Physicists Attempt to Scale the Ivory Towers of Finance. *Computing in Science and Engineering*, 1:26–39, 1999. Also available as Santa Fe Institute Working Paper No. 99-10-073.

[Freeman & Skapura 91] J. A. Freeman and D. M. Skapura. *Neural Networks: Algorithms, Applications and Programming Techniques.* Addison-Wesley Publishing Company, Redding, Massachusetts, 1991.

[Ganesh & Barr 94] Mani Ganesh and Dean Barr. Stock-specific, non-linear neural net models: The AXON System. In *Proceedings of the Neural Networks in the Capital Markets Conference*, November 1994.

[Gilbert 95] Jeremy Gilbert. Artificial Intelligence on Wall Street. Working Paper, Brandeis University, Department of Computer Science, 1995.

[Goetzmann 01] William N. Goetzmann. An Introduction to Investment Theory. Public access in: `http://viking.som.yale.edu/will/finman540/classnotes/notes.html`, Consulted May, 2001.

[Goldberg 89] David E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning.* Addison-Wesley, Reading, MA., 1989.

[Goldberg 00] David E. Goldberg. Some Reflections on Learning Classifier Systems. Technical Report 2000009, Illinois Genetic Algorithms Laboratory, University of Illinois at Urbana-Champaign, 2000. This also appeared as part of [Holland *et al.* 00].

[Goldberg *et al.* 92] David E. Goldberg, Jeffrey Horn, and Kalyanmoy Deb. What Makes a Problem Hard for a Classifier System? Technical Report 92007, University of Illinois, Department of General Engineering, 117 Transportation Building, 104 South Mathews Avenue, Urbana, IL 61801-2996, 1992.

[Grefenstette 88] John J. Grefenstette. Credit Assignment in Rule Discovery Systems Based on Genetic Algorithms. *Machine Learning*, 3:225–245, 1988.

[Grossman & Stiglitz 80] S. J. Grossman and J. E. Stiglitz. On the Impossibility of Informationally Efficient Markets. *American Economic Review*, 70:393–408, June 1980.

[Hartle 97] Thom Hartle. The systematic trader: Richard saindenberg. *Stocks and Commodities*, August 1997.

[Hartzmark 87] M. L. Hartzmark. Returns to Individual traders of Futures: Aggregate Results. *Journal of Political Economy*, 95:1292–1306, December 1987.

[Hecht-Nielsen 90] Robert Hecht-Nielsen. *Neurocomputing*. Addison-Wesley Publishing Company, Reading, Massachusetts, 1990.

[Heitkötter 94a] Jörg Heitkötter. SCS-C: A C-language Implementation of a Simple Classifier System. C source code, Systems Analysis Research Group, LSXI. University of Dortmund, Department of Computer Science, Dortmund, Germany, October 1994. Available at `http://surf.de.uu.net/people/joke`.

[Heitkötter 94b] Jörg Heitkötter. SCS-C: Some Observations using the Simple Classifier System. C source code, Systems Analysis Research Group, LSXI. University of Dortmund, Department of Computer Science, Dortmund, Germany, October 1994. Available at `http://surf.de.uu.net/people/joke`.

[Holland & Miller 91] John H. Holland and John H. Miller. Artificial Adaptive Systems in Economic Theory. *The American Economic Review*, 81(2):365–370, May 1991.

[Holland & Reitman 78] John H. Holland and J. S. Reitman. Cognitive Systems Based on Adaptive Algorithms. In D. A. Waterman and F. Hayes-Roth, editors, *Pattern-Directed Inference Systems*. Academic Press, New York, 1978. Reprinted in: Evolutionary Computation. The Fossil Record. David B. Fogel (Ed.) IEEE Press, 1998. ISBN: 0-7803-3481-7.

[Holland 75] John H. Holland. *Adaptation in Natural and Artificial Systems*. The University of Michigan Press, Ann Arbor, 1975. Second Edition: MIT press, 1992.

[Holland 86] John H. Holland. Escaping Brittleness: The Possibilities of General-Purpose Learning Algorithms Applied to Parallel Rule-Based Systems. In R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, editors, *Machine Learning II*, pages 593–623. Morgan Kaufmann, 1986.

[Holland 92] John H. Holland. *Adaptation in Natural and Artificial Systems*. MIT Press, 1992.

[Holland 95] John H. Holland. *HIDDEN ORDER. How Adaptation Builds Complexity*. Perseus Books, Reading, MA., 1995.

[Holland *et al.* 86] John H. Holland, Keith J. Holyoak, Richard E. Nisbett, and Paul R. Thagard. *Induction: Processes of Inference, Learning, and Discovery*. MIT Press, Cambridge, 1986.

[Holland *et al.* 00] John H. Holland, Lashon B. Booker, Marco Colombetti, Marco Dorigo, David E. Goldberg, Stephanie Forrest, Rick L. Riolo, Robert E. Smith, Pier Luca Lanzi, Wolfgang Stolzmann, and Stewart W. Wilson. What is a Learning Classifier System? In Lanzi et al. [Lanzi *et al.* 00], pages 3–32.

[Horn *et al.* 94] Jeffrey Horn, David E. Goldberg, and Kalyanmoy Deb. Implicit Niching in a Learning Classifier System: Nature's Way. *Evolutionary Computation*, 2(1):37–66, 1994.

[Ibbotson 01] Roger Ibbotson. Interview: Random Talks with Eugene Fama. Public access in: `http://www.ibbotson.com`, Consulted May, 2001.

[Irwin *et al.* 93] S. H. Irwin, T. R. Krukemyer, and C. R. Zulauf. Investment Performance of Public Commodity Pools: 1979-1990. *The Journal of Futures Markets*, 13:799–820, October 1993.

[Jaeger 97] Colin Jaeger. A Complete Trading Strategy. *Applied Derivatives Trading*, March 1997. In Focus - Trading.

[Jang *et al.* 97] Jyh-Shing Roger Jang, Chuen-Tsai Sun, and Eiji Mizutani. *Neuro-Fuzzy and Soft Computing. A Computational Approach to Learning and Machine Intelligence.* Prentice-Hall Inc., Upper Saddle River, NJ, 1997.

[Joshi & Bedau 98] Shareen Joshi and Mark A. Bedau. An Explanation of Generic Behavior in an Evolving Financial Market. Working Paper 98-12-114, Santa Fe Institute, 1998.

[Joshi *et al.* 98] Shareen Joshi, Jeffrey Parker, and Mark A. Bedau. Technical Trading Creates a Prisoner's Dilemma: Results from an Agent-Based Model. Working Paper 98-12-115, Santa Fe Institute, 1998.

[Joshi *et al.* 99] Shareen Joshi, Jeffrey Parker, and Mark A. Bedau. Financial Markets Can Be at Sub-Optimal Equilibria, type = Working paper. Technical Report 99-03-023, Santa Fe Institute, 1999.

[Keynes 30] John Maynard Keynes. *A Treatise on Money*, volume II. Macmillan & Co., London, UK, 1930.

[Kimoto & Yoda 93] Takashi Kimoto and Morio Yoda. Buying ans Selling Timing Prediction System for Stock Based on Modular Neural Networks. *Fujitsu Scientific & Technical Journal*, 29(3):257–264, Autumn 1993.

[Kirman & Vignes 91] Alan P. Kirman and A. Vignes. Price Dispersion. Theoretical Considerations and Empirical Evidence from the Marseilles Fish Market. In *Issues in Contemporary Economics*. Macmillan, London, 1991.

[Kleinman 97] George Kleinman. 7 Secrets Of The Professionals. *Applied Derivatives Trading*, September 1997. in September' 97 Feature.

[Kline 80] Morris Kline. *Mathematics: The Loss of Certainty*. Oxford University Press, New York, 1980.

[Kosko 92] Bart Kosko. *Neural Networks and Fuzzy Systems. A Dynamical Systems Approach to Machine Learning*. Prentice-Hall Inc., Englewood Cliffs, NJ, 1992.

[Kovacs 00] Tim Kovacs. Strength or Accuracy? Fitness calculation in learning classifier systems. In Lanzi et al. [Lanzi *et al.* 00], pages 143–160.

[Koza 92] John R. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection.* MIT Press, 1992.

[Koza 94] John R. Koza. *Genetic Programming II:Automatic Discovery of Reusable Programs.* MIT Press, 1994.

[Kurz 94a] Mordecai Kurz. Asset Prices with Rational Beliefs. CEPR Publication No. 375, Stanford University, February 1994.

[Kurz 94b] Mordecai Kurz. On Rational Belief Equilibria. *Ecomonic Theory*, 4:859–876, 1994.

[Kurz 94c] Mordecai Kurz. On the Structure and Diversity of Rational Beliefs. *Ecomonic Theory*, 4:877–900, 1994.

[Kurz 97] Mordecai Kurz. Asset Prices with Rational Beliefs. In Mordecai Kurz, editor, *Endogenous Economic Fluctuations: Studies in the Theory of Rational Beliefs*, volume 6 of *Springer Series in Economic Theory.* Springer-Verlag, August 1997.

[Lanzi *et al.* 00] Pier Luca Lanzi, Wolfgang Stolzmann, and Stewart W. Wilson, editors. *Learning Classifier Systems: From Foundations to Applications*, volume 1813 of *Lecture Notes in Artificial Intelligence.* Springer-Verlag, Berlin, 2000.

[LeBaron 95] Blake LeBaron. Experiments in Evolutionary Finance. SSRI Working Paper No. 9528, Department of

Economics, University of Wisconsin-Madison, August 1995.

[LeBaron 99] Blake LeBaron. Building Financial Markets With Artificial Agents: Desired Goals, and Present Techniques. Technical report, Graduate School of International Economics and Finance, Brandeis University, February 1999.

[LeBaron 00] Blake LeBaron. Agent Based Computational Finance: Suggested Readings and Early Research. *Journal of Economic Dynamics and Control*, 2000.

[LeBaron *et al.* 99] Blake LeBaron, W. Brian Arthur, and Richard Palmer. The Time Series Properties of an Artificial Stock Market. *Journal of Economic Dynamics and Control*, 23:1487–1516, 1999.

[Lettau 97] Martin Lettau. Explaining the Facts with Adaptive Agents: The Case of Mutual Fund Flows. *Journal of Economic Dynamics and Control*, 21:1117–1147, 1997.

[Leuthold *et al.* 94] R. M. Leuthold, P. Garcia, and R. Liu. The Returns and Forecasting Ability of Large Traders in the Frozen Pork Bellies Futures Markets. *Journal of Business*, 67:459–473, July 1994.

[Li & Tsang 99a] Jin Li and Edward P. K. Tsang. Improving Technical Analysis Predictions: An Application of Genetic Programming. 1999.

[Li & Tsang 99b] Jin Li and Edward P. K. Tsang. Investment Decision Making Using FGP: A Case Study. July 1999.

[Li & Tsang 00] Jin Li and Edward P. K. Tsang. Reducing Failures in Investment Recommendations using Genetic Programming. Barcelona, Spain, July 2000.

[Lo & MacKinlay 99] Andrew W. Lo and A. Craig MacKinlay. *A Non-Random Walk Down Wall Street.* Princeton University Press, 1999.

[Lowe & Webb 91] David Lowe and Andrew R. Webb. Time Series Prediction by Adaptive Networks: a Dynamical System Perspective. *IEE Proceedings-F*, 138(1):17–24, February 1991.

[Lux & Marchesi 99] Thomas Lux and Michele Marchesi. Scaling and Criticality in a Stochastic Multi-Agent Model of a Financial Market. *Nature*, 397:498–500, February 1999.

[Lux & Marchesi 00] Thomas Lux and Michele Marchesi. Volatility Clustering in Financial Markets: A Micro-Simulation of Interacting Agents. *International Journal of Theoretical and Applied Finance*, 3:675–702, 2000.

[Mahfoud 92] S. W. Mahfoud. Crowding and Preselection Revisited. In Reinhard Männer and Bernard Manderick, editors, *Parallel Problem Solving from Nature, 2*, pages 27–36. Elsevier Science Publishers, Amsterdam, 1992. Also published as Illigal Report No. 92004.

[Malkiel 99] Burton G. Malkiel. *A Random Walk Down Wall Street.* W.W. Norton & Company, 6th edition, 1999.

[Marengo & Tordjman 96] Luigi Marengo and Hélène Tordjman. Speculation, Heterogeneity and Learning: A Simulation Model of Exchange Rate Dynamics. *KYKLOS*, 49(3):407–438, 1996.

[Masters 95] Timothy Masters. *Neural, Novel & Hybrid Algorithms for Time Series Prediction.* John Wiley & Sons, October 1995.

[Michalewicz 96] Zbigniew Michalewicz. *Genetic Algorithms + Data Structures = Evolution Programs.* Springer-Verlag, Berlin Heidelberg New York, third edition, 1996.

[Mitchell 96] Melanie Mitchell. *An Introduction to Genetic Algorithms.* The MIT Press, Cambridge, Massachusetts, second edition, 1996.

[Mullins 97] Gerald C. Mullins. Trend following, barings bank and the 1995 yen. *Applied Derivatives Trading*, July 1997. in July' 97 Feature: An Opportunistic Trio.

[Murphy 86] J. J. Murphy. *Technical Analysis of the Futures Markets: A Comprehensive Guide to Trading Methods and Applications.* New York: Institute of Finance, 1st edition, 1986.

[Ols96] Views from the Frontier: Commentary on the New World of Forecasting and Risk Management – What's Ahead? Report, Olsen & Associates Research Group, May 1996. Interview with Andrew Lo, Blake LeBaron, Andreas Weigend, Martin Peter and other experts.

[Olsen *et al.* 92] Richard B. Olsen, Michel M. Dacorogna, Ulrich A. Muller, and Olivier V. Pictet. Going Back to the Basics – Rethinking Market Efficiency. Discussion Paper RBO.1992-09-07, Olsen & Associates Research Group, December 1992.

[Palmer *et al.* 94] R. G. Palmer, W. Brian Arthur, John H. Holland, Blake LeBaron, and P. Tayler. Artificial economic life: a simple model of a stockmarket. *Physica D*, D 75:264–274, 1994.

[Palmer *et al.* 99] R. G. Palmer, W. Brian Arthur, John H. Holland, and Blake LeBaron. An artificial stock market. *Arfif Life Robotics*, 3:27–31, 1999.

[PBS01] Betting on the Market. An Interview with Peter Lynch. Public access in: `http://www.pbs.org`, Consulted May, 2001.

[Peterson 99] Ivars Peterson. Financial Noise. *The Mathematical Association of America*, February 1999.

[Phillips & Weiner 94] G. M. Phillips and R. J. Weiner. Information and Normal Backwardation as Determinants of Trading Performance: Evidence from the North sea Oil Forwards Market. *Economic Journal*, 104:76–95, January 1994.

[Pictet *et al.* 96] Olivier V. Pictet, Michel M. Dacorogna, Rakhal D. Davé, Bastien Chopard, Roberto Schirru, and Marco Tomassini. Genetic Algorithms with collective sharing for Robust Optimization in Financial Applications. Working Paper OVP.1995-02-06, Olsen & Associates Research Group, January 1996.

[Piper 97] John Piper. The Human Brain. *Applied Derivatives Trading*, March 1997. In Focus - Trading.

[Quinlan 93] J. R. Quinlan. *C4.5: Programs for Machine Learning*. Morgan Kaufmann, San Mateo, CA, 1993.

[Rajan & Slage 96] Vijay Rajan and James R. Slage. The use of Artificially Intelligent Agents with Bounded Rationality in the Study of Economic Markets. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI'96)*, volume 1, pages 102–107. AAAI Press, The MIT Press, August 1996.

[Ref95] In Apostolos-Paul N. Refenes, editor, *Neural Networks in the Capital Markets*. John Wiley & Sons, March 1995.

[Ref96] In Apostolos-Paul N. Refenes, Yaser Abu-Mostafa, John Moody, and Andreas Weigend, editors, *Neural Networks in Financial Engineering: Proceedings of the Third International Conference on Neural Networks in the Capital Markets*. World Scientific Publishing Company, June 1996.

[Reynolds *et al.* 95] S. B. Reynolds, J. M. Mellichamp, and R. E. Smith. Box-Jenkins Forecast Model Identification. *AI Expert*, page 15, June 1995.

[Ridley 93] Matt Ridley. Frontiers of Finance. *The Economist*, pages 3–22, October 9th 1993.

[Roosevelt 97] Ruth Barrons Roosevelt. Scenario Trading. *Stocks and Commodities*, March 1997. in Trading Psychology.

[Ross 99] Sheldon M. Ross. *An Introduction to Mathematical Finance. Options and Other Topics*. Cambridge University Press, Cambridge, UK, 1999.

[Routledge 94] Bryan R. Routledge. Artificial Selection: Genetic Algorithms and Learning in a Rational Expectations Model. Technical Report, Faculty of Commerce and Business Administration, University of British Columbia, November 1994.

[Sakurai 88] Miyoko Sakurai. Hitachi Works on a Neural Stock Analyst. *Electronic Engineering Times*, (508):25, October 17 1988.

[Schulenburg & Ross 99] Sonia Schulenburg and Peter Ross. An Evolutionary Approach to Modelling the Behaviours of Financial Traders. In *Genetic and Evolutionary Computation Conference Late Braking Papers*, pages 245–253, Orlando, Florida, 1999.

[Schulenburg & Ross 00] Sonia Schulenburg and Peter Ross. An Adaptive Agent Based Economic Model. In Lanzi et al. [Lanzi *et al.* 00], pages 265–284.

[Schulenburg & Ross 01] Sonia Schulenburg and Peter Ross. Strength and Money: An LCS Approach to Increasing Returns. In Pier Luca Lanzi, Wolfgang Stolzmann, and Stewart W. Wilson, editors, *Advances in Learning Classifier Systems*, volume 1996 of *Lecture Notes in Artificial Intelligence*, pages 114–137. Springer-Verlag, Berlin, 2001.

[Schulenburg & Ross 02] Sonia Schulenburg and Peter Ross. Explorations in LCS Models of Stock Trading. In Pier Luca Lanzi, Wolfgang Stolzmann, and Stewart W. Wilson, editors, *Learning Classifier Systems*, volume Forthcoming of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, Berlin, 2002.

[Sen & Sekaran 95] Sandip Sen and Mahendra Sekaran. Multiagent coordination with learning classifier systems. In *Working Notes of the IJCAI-95 Workshop on Adaptation and Learning in Multiagent Systems*, pages 429–435, August 1995.

[Simon 55] Herbert A. Simon. A Behavioral Model of Rational Choice. *Quarterly Journal of Economics*, 69:99–118, 1955.

[Simon 59] Herbert A. Simon. Theories of Decision-Making in Economics and Behavioral Science. *American Economic Review*, 49:253–283, 1959.

[Smith & Taylor 98] Robert. E. Smith and Nick Taylor. Market-Oriented Programming: Some Early Lessons. In C. Looney and J. Castaing, editors, *Proceedings of the 1998 International Conference on Intelligent Systems*, pages 221–224. ISCA Press, 1998.

[Smith *et al.* 99] Robert E. Smith, B. A. Dike, R. K. Mehra, B. Ravichandran, and A. El-Fallah. Classifier Systems in Combat: Two-sided Learning of Maneuvers for Advanced Fighter Aircraft. In *Computer Methods in Applied Mechanics and Engineering*. Elsevier, 1999.

[Soros 95a] George Soros. A Failed Philosopher Tries Again. Jan Patocka Memorial Lecture, April 1995. Public access in: `http://www.geocities.com/WallStreet/3880/index.html`.

[Soros 95b] George Soros. The Role of World Governance, September 1995. Public access in: `http://www.geocities.com/WallStreet/3880/index.html`.

[Soros 01] George Soros. The Wisdom of George Soros. Quotations, Consulted May, 2001. Public access in: `http://www.geocities.com/WallStreet/3880/index.html`.

[Sutton & Barto 98] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning*. MIT Press, 1998.

[Sutton 91] R. S. Sutton. Reinforcement Learning Architectures for Animats. In J. A. Meyer, , and Stewart Wilson, editors, *From Animals to Animats: Proceedings of the 1st International Conference on Simulation of Adaptive Behavior*, pages 288–296, Cambridge, MA, 1991. MIT Press.

[Tanous 97] Peter J. Tanous. *Investment Gurus*. New York Institute of Finance, 1997. The Interview with Eugene Fama also appears in the web page of Dimensional Fund Advisors Inc. given in [Tanous01].

[Tanous 01] Peter J. Tanous. An Interview with Rex Sinquefield. Public access in: `http://www.dfafunds.com`, Consulted May, 2001.

[Thomas & Sycara 00] James D. Thomas and Katia Sycara. Integrating Genetic Algorithms and Text Learning for Financial Prediction. In Alex A. Freitas, William Hart, Natalio Krasnogor, and Jim Smith, editors, *Proceedings of the GECCO-2000 Workshop on Data Mining with Evolutionary Algorithms*, pages 72–75, Las Vegas, Nevada, USA, 2000.

[Thomas 00] James D. Thomas. Thesis Proposal, Carnegie Mellon University, Department of Computer Science, 2000.

[Toghraie 96] Adrienne Laris Toghraie. Transforming trader self-talk. *Stocks and Commodities*, November 1996. in Trading Psychology.

[Toghraie 97a] Adrienne Laris Toghraie. Dealing with change. *Stocks and Commodities*, May 1997. in Trading Psychology.

[Toghraie 97b] Adrienne Laris Toghraie. The Perfect Fit. *Applied Derivatives Trading*, April 1997. In April'97 Feature.

[Tomek & Querin 84] W. G. Tomek and S. F. Querin. Random Processes in Prices and Technical Analysys. *The Journal of Futures Markets*, 4:15–23, Spring 1984.

[Tordjman 98] Hélène Tordjman. Some General Questions About Markets. Working Paper IR-98-025, International Institute for Applied Systems Analysis, Laxenburg, Austria, May 1998.

[Tri93] In Robert R. Trippi and Efraim Turban, editors, *Neural Networks in Finance and Investing. Using Artificial Intelligence to Improve Real-World Performance*. Probus Publishing Company, Chicago, Illinois, 1993.

[Tsang *et al.* 98] Edward P. K. Tsang, Jin Li, and James M. Butler. EDDIE Beats the Bookies. *International Journal of Software, Practice and Experience*, 28(10):1033–1043, August 1998.

[Tsang *et al.* 00] Edward P. K. Tsang, Jin Li, Sheri Makrose, Hakan Er, Abdel Salhi, and Giulia Iori. EDDIE in Financial Decision Making. *Journal of Management and Economics*, November 2000.

[Vakkur 97] Mark Vakkur. The moving average convergence/divergence histogram. *Stocks and Commodities*, April 1997. in Indicators.

[Viswanath 01] P. V. Viswanath. Efficient Capital Markets and Corporate Finance. Public access in: `http://`

webpage.pace.edu/pviswanath, Consulted May, 2001.

[Wagner & Matheny 97] Gary S. Wagner and Bradley L. Matheny. Market Timing And Candlesticks. *Stocks and Commodities*, March 1997. in Trading Techniques.

[Watkins 89] C. Watkins. *Models of Delayed Reinforcement Learning*. Unpublished PhD thesis, Psychology Department, Cambridge University, Cambridge, United Kingdom, 1989.

[Weisbuch *et al.* 98] Gérard Weisbuch, Alan Kirman, and Dorothea Herreiner. Market Organisation and Trading Relationships. Working Paper, Laboratoire de Physique Statistique, Ecole Normale Supérieure, February 16 1998.

[Wellman & Wurman 97] Michael P. Wellman and Peter R. Wurman. Market-Aware Agents for a Multiagent World. *Lecture Notes in Computer Science*, 1237, 1997.

[Wellman 96] Michael P. Wellman. Market-Oriented Programming: Some Early Lessons. In Scott Clearwater, editor, *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific, River Edge, New Jersey, 1996.

[White 88] Halbert White. Economic Prediction Using Neural Networks: The Case of IBM Daily Stock Returns. In *Proceedings of the IEEE International Conference on Neural Networks*, July 1988.

[Whitley 93] Darrell Whitley. A Genetic Algorithm Tutorial. Technical Report CS-93-103, Colorado State University, November 1993.

[Wilson & Goldberg 89] Stewart W. Wilson and David E. Goldberg. A Critical Review of Classifier Systems. In J. D. Schaffer, editor, *Proceedings of the Third International Conference on Genetic Algorithms*, pages 244–255. Morgan Kauffmann, 1989.

[Wilson 85] Stewart W. Wilson. Knowledge Growth in an Artificial Animal. In John J. Grefenstette, editor, *Proceedings of the 1st International Conference on Genetic Algorithms and their Applications (ICGA85)*, pages 16–23, Pittsburgh, PA, July 1985. Lawrence Erlbaum Associates.

[Wilson 87] Stewart W. Wilson. Classifier Systems and the Animat Problem. *Machine Learning*, 2:199–228, 1987.

[Wilson 94] Stewart W. Wilson. ZCS: A Zeroth Level Classifier System. *Evolutionary Computation*, 2(1):1–18, 1994.

[Wilson 95] Stewart W. Wilson. Classifier Systems Based on Accuracy. *Evolutionary Computation*, 3(2):149–175, 1995.

[Wurman *et al.* 98] Peter R. Wurman, Michael P. Wellman, and William E. Walsh. The Michigan Internet AuctionBot: A Configurable Auction Server for Human and Software Agents. In Katia P. Sycara and Michael Wooldridge, editors, *Proceedings of the 2nd International Conference on Autonomous Agents (AGENTS-98)*, pages 301–308, New York, May 1998. ACM Press.

[Wuthrich *et al.* 98] Beat Wuthrich, D. Permunetilleke, S. Leung, Vincent Cho, J. Zhang, and W. Lam. Daily Prediction of Major Stock Indices from Textual WWW Data. In *Knowledge Discovery and Data Mining – KDD-98*, pages 364–368, 1998.

[Young 97a] Patrick Young. The Internet - Virtual Trading Floors. *Applied Derivatives Trading*, May 1997. In Burgenstock 96, ADT Special Coverage.

[Young 97b] Patrick Young. Who's Afraid of A Little Black Box? *Applied Derivatives Trading*, September 1997. in September' 97 Feature.

[Zulauf & Irwin 97] Carl R. Zulauf and Scott H. Irwin. Market Efficiency and Marketing to Enhance Income of Crop Producers. OFOR Paper Number 97-04, Ohio State University and University of Illinois, October 1997.

[Zurada 92] Jacek M. Zurada. *Introduction to Artificial Neural Systems*. West Publishing Company, St. Paul, MN, 1992.