UNIVERSITE DE LAUSANNE
FACULTE DES SCIENCES

Financial markets as a complex system

Thèse de doctorat

présentée à la

Faculté des Sciences de
l'Université de Lausanne

par

Thomas Ankenbrand

Licencié en Informatique économique
Haute-Ecole de Saint-Gall

Jury

Prof. Dieter Schwarzenbach, Président
Prof. Francois Grize, Rapporteur
Prof. Marco Tomassini, Directeur de thèse
Prof. André-Réné Probst, Expert
Dr. Olivier Pictet, Expert

LAUSANNE
1998

Imprimatur

Université de Lausanne

Faculté des Sciences

Vu le rapport présenté par le jury d'examen, composé de

MM M. le Prof. Marco TOMASSINI, directeur de thèse, M. le Prof. François GRIZE, rapporteur, M. le Prof. André-René PROBST, M. le Dr Olivier PICTET, experts, M. le Prof. Dieter SCHWARZENBACH, président.

le Conseil de Faculté autorise l'impression de la thèse de Monsieur Thomas ANKENBRAND, licencié en informatique économique de la Haute-Ecole de Saint-Gall

intitulée

Financial markets as a complex System

Lausanne, le 20 octobre 1998

pour Le doyen de la Faculté des Sciences:
Professeur Dieter SCHWARZENBACH

D. Schwarzenbach

Table of Contents

1. Introduction

Many phenomena which take place in the economy are poorly understood. The research programme that gave rise to this thesis was concerned with a better quantitative understanding of the evolution of the financial markets as a core of the economy. A definitive answer involving the degree of understanding gives the ability of making correct forecasts for the future behaviour of the financial markets which is the ambitious goal of this thesis. The financial markets are analysed and modelled on two different levels. First statistics and artificial neural networks are used for modelling the economic behaviour of a single agent. Second due to the consequences of having different agents for different markets, the agents interact in artificial but realistic market environment for simulating the whole market behaviour.

An investor, speculator or the artificial agent needs a correct model of the financial markets for optimal behaviour. Optimal behaviour means a correct forecast of the future behaviour of the system which allows one to take the right positions for profitable trading. Therefore, investors and moreover speculators, are interested in forecasting techniques. The academics have denied the possibility of beating the market for a long time. The markets seemed to be efficient and therefore it was impossible to gain any free lunch with the different traditional forecasting methods, but the field was still more an art than a science. The two main families of financial forecasting are fundamental and technical analyses. Spoken in statistical terms, the fundamental analysis is a multivariate quantitative and qualitative analysis; whereas the technical analysis was often a univariate quantitative analysis. Both types of analyses were mainly linear and have ignored possible non-linear relations. Today the traditional technical analysis also includes indicators which imply complicated non-linear relations, and it is also used for multivariate studies. The presented methodology covers all quantitative aspects of forecasting technologies. The approach integrates economic knowledge into a multivariate non-linear artificial neural network model. The core of the work is a feasibility analysis which consists of a series of different univariate and multivariate, linear and non-linear statistical tests. The multivariate non-linear test is original and its is seldom attempted in artificial neural network work. This multivariate non-linear neural network test allows to detect dependencies between input and predictor data which are not detected with traditional statistical analysis. This new detected dependencies enhance the forecasting performance of the models. The feasibility analysis evaluates the "a priori" chance of forecasting the defined system and specifies the input and topology of the artificial neural network in the modelling process. The method is applied to two real-life case studies of the Swiss stock and bond market.

The second level of modelling financial markets represents the macro-economic view after the micro-economic one of the first level. It simulates several interacting financial markets in a realistic manner. The artificial market uses populations of diverse artificial traders and investors, whose typologies are simplified versions of those existing in real markets. The agents can trade on several markets simultaneously and can establish an investment portfolio with several assets. They have also some learning and evolutionary capabilities.

Distributed, decentralised, interacting assemblies of agents are commonplace in biology and ecology. One may think for instance of ant or bee colonies, fish schools or flocks of birds as well-known natural examples of the emergence or ordered, coherent collective behaviour without any central control or regulation. These and other systems are complex, evolving and adaptive. No theory of such systems is yet known; rather, their collective behaviour seems to result from the non-linear aggregation of simple local rules followed by the agents. In other words, the agents have a limited horizon and computation capabilities but, in spite of these limitations, the system as a whole appears to behave according to regular patterns.

Artificial life is a new discipline that strives to use computers, robots, and other artificial means to study life-like phenomena [Langton89]. The approach is synthetic rather than analytic and it works mainly by putting together systems that behave, in some respects, like living organisms. Computer simulation for populations of artificial entities is central to such an approach and has given insight into the behaviour of complex systems that would have been difficult to obtain with classical approaches such as differential equations dealing with aggregate quantities.

Economic, financial and social systems share many features with the aforementioned natural systems and their artificial counterparts. Many important economic processes are complex in the sense that it is difficult to decompose them into separate parts that can be studied in isolation and aggregated to yield a complete to whole picture. Furthermore, economic agents do not seem to possess in reality the perfect rationality and computing abilities that the classical economic theory attributes to them. Issues of beliefs and prediction about the future become important and the individual beliefs and choices when aggregated shape the economic indicators, the market prices and ultimately the world that the agents must deal with. Moreover, this world is dynamic, as a result of the constantly changing collection of beliefs and strategies of the agents in their endeavour to adapt to the system evolution. In the real world, contrary to common economic assumptions, there is also diversity between agents due to different computing abilities and computing procedures as well as different time horizons and risk profiles in the financial field. Computerised collection of diverse, adapting agents should then prove useful, especially in studying the dynamic aspects of economic systems which are often neglected in standard theories. Some recent work has shown that the approach usefully complements the classical ones and can offer new insights of its own [Arthur95; Tesfatsion97]. The artificial market blurring the lack of experiments in economy. In traditional economics, it is difficult to study the consequences of specific policies or rules of behaviour in an orderly way because these systems are dynamic and constantly changing. However, the computational approach allows this to be carried out easily, repeatedly and without risk. This makes them useful as pedagogical devices and also as a source of ideas for new methods of analysis.

The thesis is divided into three parts. The first part provides an underlying analysis for the following sections. It contains five chapters and starts with the second chapter which is an overview of economics and financial markets for readers from fields other than economics. It should help to understand the basic principles of financial markets and its practical applications. The third chapter discusses the possibility of a crystal ball that could foretell the future of financial markets. This is the question of modern finance theory [Samuelson89, 251]. This chapter gives a dialectic discussion of the problems and the history of market prediction. The first paragraph describes efficient market theory which challenges the assumptions supporting market prediction. The following paragraphs question the assumptions and holdings of the efficient market theory and demonstrate how predictions in financial markets are possible. Recent developments in non-linear dynamic systems help to explain some of the natural phenomena in physics and chemistry. Non-linear, dynamic and adaptive systems can perhaps also help to explain the behaviour of the global economy, or financial markets as an important part of it aside from the traditional linear economic models. The fourth chapter explores different possibilities and limits for the application of non-linear dynamic systems in finance and gives the conclusion that the financial markets are adaptive dynamic systems to external shocks. Therefore, statistical tools have to be able to detect the different components in the empirical analysis of financial time series. The fifth chapter will examine a variety of analytical statistical systems. Linear univariate, non-linear univariate, linear multivariate, and non-linear multivariate systems will be discussed. The sixth chapter describes neural networks and their possible application in financial markets. It gives a short overview of artificial intelligence which encompasses artificial neural networks, and its relation to the economy. An overview of artificial neural networks and its possible applications in time series analysis is also given. Lastly, the computation of sensitivity analysis is described. This computation is at the core of the original multivariate non-linear test for the identification of relevant input indicators of the market model which is the final missing part of the forecasting methodology.

The second part contains the micro-economic applications with the analysing and modelling process of a single agent. It starts with chapter seven which presents an integrated approach for analysing and modelling the behaviour of financial markets. The methodology makes use of the assumptions and the elements described in the previous chapters. Its originality lies in the fact that it is based on statistical and macro-economic principles; it integrates fundamental economic knowledge into a multivariate non-linear time- series artificial neural network model. The final model allows for the forecasting of financial time series with good results. This is shown in practice in the next chapter, where the methodology is successfully applied to two case-studies of the Swiss stock and bond market.

The third part with chapter nine and ten contains the macro-economic application of artificial life in the financial markets. It gives additional insights into this mechanism in respect to the models of the previous chapters. The artificial market is built around the behaviour of the real actors of the markets. These actors are represented by computer agents which interact together according to the rules of trading in the artificial market. The first reason of this micro-level approach is the belief that the complex behaviour of financial markets emerges as the result of the collective action of many simple agents in mutual action, rather than due to the complicated and unknown nature of each individual. Chapter nine contains the fundamentals of the artificial market simulation environment and the descriptions of the different trading strategies of the agents. The following chapter ten describes the results and conclusions of the different simulations.

2. Economics and Markets

This chapter allows an overview of economics and financial markets for readers from fields other than economics. It should help to understand the basic principles of financial markets and its practical applications.

2.1 Economics

Economics is a social science whose subjects are the human beings together with their relations. Herein economics concentrate on a special relation between persons, this is usually the exchange relation. The scientific investigation of exchange is economics as a social science. The motivation therefore is drawn from the shortness of goods which is related to supply and demand.

Because economic relations differ very often in numbers and kinds, economics as a science has to abstract the reality. This abstraction is shown in theories about the objects and their relations. A theory consists of axioms and hypothesis. An axiom is the abstract description of the empirical reality, whereas the hypothesis is derived logically from the axioms. The hypothesis must be verified through comparison with actual observations. The value of a theory does lie not only in a correct prediction of the reality, but it also lies in an understanding why the reality goes in the predicted way [Franke86, 1ff.].

Economics as a science has some interfaces to other social sciences or applications. Politics use economic theory to form the reality. Here it is important to distinguish between positive and normative economics. Positive economics describe different facts in the economy, while normative economics involve value judgement, what again must be resolved by political decision, not by economic science. They are never settled by science or by a certain appeal to facts. There is no right answer for example to how high inflation should be. But a common sense value or goal would be the efficient allocation of resources. The original object of investigation or the positive economics is an economy without any interventions of the government. But a modern economic system needs some legal configurations of the government [Samuelson89, 10f.].

There are two theories about interventions of the government in the economic field. One is the theory of Keynes, which was developed during the big depression of the thirties. This theory allows the escape from the depression with interventions of government. The other theory, called Monetary theory, was developed in the sixties, and decreases the inflation. Monetary theory though reject interventions of the government, because it states that the economy regulates itself best [Dornbusch87, 2].

Macroeconomics and microeconomics are the two main parts of economics. Macroeconomics studies the functioning of the economy as a whole, where microeconomics analyses the behaviour of individual components like industries, firms, and households [Samuelson89, 5].

As the very core of economy stands the exchange of goods and a place for exchange, these are markets. The following paragraph describes different kinds of markets.

2.2 Markets

A market is the abstract place where supply and demand meet. The supply and demand of a good determine its price in the market. This price co-ordinates the supply and demand or the production and consume. The price of a good is increasing, if the demand is higher than the supply and the production of the same good is more attractive if the price is higher. The higher production allows to satisfy the higher demand. On the other side, the higher price decreases the demand. Thus the price has two effects: Both of them compensate differences of demand and supply and the market tends to be in an equilibrium [Franke86, 6f.].

There are two types of markets in the economy. The first type is the commodity market in which consumers buy products from the suppliers. The second type is the factor market where suppliers buy labour and capital from the consumer. The factor market has two parts: The labour and the capital or financial market [Franke86, 4].

The goal of a market participant is an optimal satisfaction of his demands. An axiom of the economic theory for the satisfaction of the demands is the rational behaviour of the participants. A market has to offer different opportunities for the participants. The market needs a competition between the different suppliers and consumers. This allows the participants to decide between different partners for the exchange or to reject the possibility of trade and to wait for a better opportunity. It is not sufficient to have enough suppliers and consumers for acting rationally. The market participants need also the information about the offerings and demands of the other participants. The market should be transparent. Rational behaviour, competition and transparency are three axioms in economics. These axioms are not perfectly true in reality. The following paragraphs discuss the single axioms in detail [Franke86, 160f.].

2.2.1 Competition

The competition depends on the number of suppliers and consumers. A market with only one supplier or consumer is a monopoly. The missing competition allows the single participant to reach a better price than in a competition market. He has the possibility to reduce or to increase the supply or demand in the market which allows a direct influence on the price. The situation of only one supplier or consumer is rare, but there are often a few participants of one side of the market. These few participants can find trust, which allows them to reach similar advantages to those of a monopoly.

The price is the only factor of competition in a perfect market. However there are more factors of competition than the price in most markets. The goods in an imperfect market are not exactly the same: there are differences in quality, the market place and the personal relations between the supplier and consumer. An example of buying a new car explains an imperfect market. A buyer of a new car has some preferences about the brand, the style, the type, the colour and others. He also has preferences about the market place. He likes to buy the car near to his home. Perhaps he knows a salesman personally and he would prefer trading with him. The price of the car is also important, but it is not the only factor which influences the decision of the buyer. The differences of the products in an imperfect market can be small and thus it is possible to interchange the different products. They are in the same market and in competition, but there are other factors than price that influence the decision of the participants [Franke86, 160f.].

2.2.2 Transparency

The participants need information for trading and in theory all information is freely available. The market is then said to be transparent. However, obtaining the information has a cost and sometimes the information is not available [Franke86, 161f.].

2.2.3 Rational behaviour

The market participants act rationally in theory. The preferences support the goal of an optimal satisfaction of demand. This behaviour guarantees a single price in a perfect and transparent market. The markets are not perfect and transparent in reality. The adapted behaviour depends on the degree of competition. The participants have other preferences than the price. The other preferences are not always rational to explain. They are founded outside of the economics in the human nature [Franke86, 162f.].

The participation's have also expectations about the future and the behaviour of the other participants. The theory of rational expectations holds that people use all or the best available information and economic theories for their decisions. Forecasts are unbiased and are based on all available information. An unbiased forecast contains no systematic forecasting errors. The critical assumption of rational expectation theory is that people understand any systematic rule by government and any strategy of the other market participants. An implicit assumption is that prices move quickly enough so that markets are always in equilibrium [Samuelson89, 372ff.]. The rational expectations theory is strongly connected to the two previous axioms of the economic theory.

The expectations, the imperfect knowledge of the environment, and the limited computational power support a subjective or bounded rationality of the market participants [Simon82, 8ff.]. The decision process of human is not very clear because it is depending of the decision environment. The human decision making is a theme for future theoretical and empirical research [Loomes98]. The implication of the not rationally acting agents are described in the following chapter in the section about adaptive systems.

2.3 Financial markets

Financial markets transform factories, department stores, and other firms into assets that are tradable in the markets. It is possible to trade these assets with anonymous buyers or sellers, the existence of these assets allows to diversify the risk over a wide group of investors. The financial markets put prices to the assets. These prices allow to allocate money to the most promising use. The financial markets allocate capital to the most productive use, as for example in the field of producing goods. Most economic crisis have originated from abuses in the financial market, this may explain the fear existing about the financial markets. They invite speculation: most people are buying stocks and bonds with the purpose of making money and not for making the allocation of capital more efficient. The financial markets are a place where the people do not only exchange money for stocks and bonds, it is a focal point where individuals, businesses, and even entire economies anticipate the future. Movements of prices reveal how confident people are in their expectations. This allows business to rise capital from the financial markets for their projects. Liquidity, low transaction costs, and the freedom of investors and speculators to act on information are essential to that economic and social function [Bernstein92, 6ff.].

Finance is a part of economic sciences. Financial markets allocate the resources, where the instruments or assets for this allocation are stocks, bonds, warrants, derivatives and more. The management of these instruments seek to optimise the profit or to minimise the risk. Finance is an interdisciplinary science and it is strongly related to economics, mathematics and statistics [Zimmermann96, 1ff.].

The change of finance from an art to a science was supported by the change of investors. If individual investors had dominated the markets during the 1970s and 1980s, the change would have never taken place. The ingenious journal articles from academics would have stimulated more ingenious journal articles, but the new theories of Blake, Fama, Marcowitz, Merton, Miller, Modigliani, Sharp, Tobin, and Samuelson would not have changed a lot in Wall Street or in other financial centres. The revolution was catalysed by the rise of institutional investors such as pension and mutual funds, which have different investment goals, possibilities and restrictions. The individual investor does not trade very often because he is taxable and has higher transaction costs. He prefers a buy and hold strategy. An individual investor does not have the resources available to do deeper research, to learn the newest theories and to use the necessary computing technology, whereas an institutional investor has these possibilities. The institutional investor makes use of all these possibilities in a performance oriented way, facing other restrictions and also a fierce competition. Today the institutional investors and the professional managers are the principal players in the markets. They push financial innovations, but innovations must be preceded by theory [Bernstein92, 8ff.].

The following paragraphs describe different financial instruments, the investment process and different financial markets involved.

2.3.1 Financial Instruments

This paragraph surveys various investment alternatives and their opportunities. Main financial instruments are bonds and stocks. Exchange rates do not only connect the different marketplaces, but also are of importance in an international portfolio regarding the performance of an instrument. Therefore they are also an instrument in this sense [Heri96, 243]. Existing derivatives for these instruments give an alternative to a direct investment.

A bond is a security. An investor lends money to a company or even to the government. The bond as security promises to pay a certain amount of money in interest periodically for a number of years until it matures. At maturity, the borrowing company or government must pay off the principal of the bond at its face value. Payments for interest and principal must be made on time, regardless of whether the company or government has been making profits or not. The investor can also sell the bond to another investor before it matures. This combines the advantages of short term investment with long term credit [Samuelson89, 479].

Common stocks or equities represent shares of ownership in a corporation. The shareholders in principle control the companies they own. They elect directors and vote on many important issues. But actually, no single person or organisation owns a sufficient number of shares to control major companies like IBM, GM or Nestle [Samuelson89, 479]. The motivation to hold stocks lies in the pay-off of dividends and their rising prices. Both depend on the profit of the company and the expectation of the profit.

There are numerous hybrid or combination securities, combining the features of bonds and stocks. A special kind of hybrids are the derivatives, their values originate from the basic of the more traditional investment instruments, like stocks and bonds to which the derivatives are related to in one form or another. The main two derivatives are options and futures. An option represents the right to buy (call option) or to sell (put option) the underlying asset at a specified price within a specified period of time. Underlings for derivatives are named stocks, stock indices, bonds, interest rates, foreign currencies, and commodities. A future is the obligation to buy or sell at a specified price within a specified period of time. The derivatives allow the investor to better control the risk and profit of his portfolio. But the complexity of the investment process increases with these new instruments [Cohen87, 469ff.].

2.3.2 The Investment Process

Investors must decide up on the investment objectives and then select from the available investment alternatives. Decisions must be made in the selection of financial instruments in order to provide a mix with attractive rates of return balanced against acceptable risk for the portfolio [Cohen87, 3f.].

The financial analysis determines if a particular asset is priced properly, or if it is above or below its instrinc value. The best way to measure its instrinct value is the present value of expected future cash flows from the investments. Future cash flows of a stock are dividends. The value of these dividends is not known because it depends on the future profit of the company. In the worst case, the company is not able to pay any dividends, at anytime in the future. The next section describes the possibilities on how to expect future dividends. The value of the sum of these future dividends today is called present value, which means that the present value is the discounted sum of the expected future dividends. The Dividend Discount Model uses this approach for the valuation of stocks. The value of a stock depends on this approach from the future dividends and the interest rates. Both factors are unknown at present therefore the investor must fall back on expectations of the two factors. The fundamental analysis helps to estimate future profits of a company, which are responsible for the payment of future dividends [Cohen87, 341].

The fundamental analysis helps to estimate future earnings, and it also studies the power of the firm for future payments to the investor. These projections are based on profit estimates affected by the economic environment, industry trends, and the company itself. To do this, the analyst judges the balance sheet and calculates different ratios for the comparison of different firms. This allows to decide which one would be the most promising company for buying stocks and here with worthwhile for [Cohen87, 40f.].

The value of a bond is the future discounted cash flow. The future cash flow of a bond is the payback of the principal at maturity and the interest rate every year. Both of them are known. The value depends on the level of the interest rate and on the quality of the issuer. This is why bonds are not as risky as stocks [Cohen87, 436f.]. Macroeconomics models help to explain the level of the interest rates and also the currency exchange rates [Samuelson89, 937ff.]. The valuation of derivatives depends on the instrument and the underlying [Cohen87, 469ff.].

Perhaps as important as the choice what stocks, bonds, or currencies to buy, is the decision when to buy them. The technical analysis helps for investment timing. It studies phenomena which are an integral part of the market mechanism and it does this by looking for recurring patterns of price movement or recurring interrelationships between price movements and other market data like volumes. Price movements are thought to reflect the opinions of millions of different people about everything having a bearing on the market. Some of the technical methodologies are chart analysis, moving averages, theory of contrary opinion, relative strength, and Elliot wave theory [Cohen87, 252ff.].

A diversification of assets in a portfolio diminishes the risk. A traditional investor buys undervalued assets one at a time and sells them later at a profit. The modern investor selects assets in a mix. He uses the Capital Asset Pricing Model to build up an efficient portfolio with the highest return at a given risk level [Cohen87, 11ff.].

2.3.3 Marketplaces

The financial system is composed of a number of markets. There is a market for every financial instrument, there are primary markets and secondary markets, organised exchanges, over-the-counter markets [Cohen87, 44f.] and there are different geographical markets.

The primary market is the new issues market. This simple example should explain the main function of the primary market. The U.S. government needs some capital for the financing of infrastructures. So it places a series of bonds with a duration of ten years and an interest rate of ten percent in the market. Now it is possible for every investor to buy such a bond from the government. The cash flows of the investment are the following: First the investor pays the face value of the bond to the government, then the investor receives a yearly ten percent interest from the government. After ten years the government returns the face value of the bond and the investment period is finished. The market maker in the primary markets are normally investment bankers. They buy new issues form the lender at an agreed-on price and hope to resell them to the investing audience at a higher price. The investment bankers differ from a broker, who acts as an agent without risk. Secondary markets provide the trading forums with existing securities. The simple example at the beginning of this paragraph also helps to explain the function of the secondary market. The investor will receive the money back at the face value of the bond, after ten years. What can be done when the lender needs the money back earlier? The duration of the bond has been set to ten years, so there is no way of getting the money from the borrower directly. But what the investor can do, is sell his bond with the rights of interest and the pay-off of the face value to another investor. The secondary market in this way allows to liquidate long term investments at once. The buyer of the bond also has an advantage. He has not to wait to the next issue of government bonds. He also can invest at once.

The organised exchanges and over-the-counter markets are types of secondary markets. The organised exchanges provide physical or electronic marketplaces, for example the New York Stock Exchange is the biggest stock exchange in the world. The most important stock exchange of Switzerland is in Zurich. The exchange for derivatives in Switzerland is SOFFEX (Swiss Options and Financial Futures Exchange). The over-the-counter market consists of a loose aggregation of brokers and dealers. It is a complex world-wide network of trading rooms linked by telephone and electronic communications [Cohen87, 46f.].

Today global trading is reality. The over-the-counter currency exchange market is always open. The national exchanges still have opening and closing hours, but they already cooperage with other exchanges for expanding the trading hours [Cohen87, 79].

3. Market prediction

A correct prediction of prices increases the possibility of making additional profits and reduces the risk in financial markets. Therefore, investors and moreover speculators, are interested in forecasting technologies. There are many different technologies, but the field is still more an art than a science. The previous paragraph of the investment process described the two main families of financial forecasting technologies: fundamental and technical analyses. Both are well known and widely used in financial markets with varying degrees of success rates.

Most applications in the past have covered the stock picking process. Stock picking means the selection of a specific stock in a given stock market. However, in an international portfolio stock picking is the last step in the investment process. Prior to stock picking is the selection of countries and investment instruments. This is often forgotten in discussions about the possibility of market prediction.

Is there a crystal ball that will foretell the future of financial markets? This is the question of modern finance theory [Samuelson89, 251]. This chapter gives a dialectic discussion of the problems and the history of market prediction. The first paragraph describes efficient market theory, which denies the possibility of market prediction. The following paragraphs critique the assumptions and holdings of the efficient market theory and demonstrate how predictions of financial markets are possible with the bounds of an economic point of view.

3.1 Efficient market theory

Findings of price studies of financial markets from traditional economics and finance professors are controversial and have angered many financial analysts. The academic ideas developed in this area are grouped under the heading of efficient market theory. An efficient market is one where all new information is quickly understood by market participants and becomes immediately incorporated into market prices. The theory of efficient markets holds that market prices contain all available information. It is not possible to make profits by looking at old information or at patterns of past price changes [Samuelson89, 251f.].

The market is assumed to be efficient because it consists of a large number of rational, profit-seeking, risk-averting investors who compete freely with each other in estimating the future value of an asset. Since any new change affecting a given asset is quickly known throughout the entire investment community, it is therefore rapidly reflected in the price of the asset to which it relates [Cohen87, 13]. The price movements in a efficient market are random, because all predictable information has already been built into the price. It is the arrival of new information that affects prices. This information, known as news, by its very definition must be random and unpredictable. Predictable information would have already been reflected in market yields before the event happened [Samuelson89, 252f.].

A price chart in an efficient market is a random walk [Samuelson89, 252f.] because any non-random fluctuations in price (other than a steady upward drift approximating the risk-adjusted rate of return) would be exploited by speculators who would buy before an expected rise in price or sell short before an expected fall in price. This would eliminate any predictable fluctuations and make all price changes random [Cohen87, 14].

The subject of stock market efficiency has been divided into three separate and distinct components: the weak, the semistrong, and the strong forms. The weak form holds that successive changes in stock prices are essentially independent of each other and that the information content of historic market data (prices, trading volume, etc.) is already embedded in the existing price. The semistrong states holds that stock prices adjust rapidly to all new publicly available information (not only market data, but economic, social, and political events) and that action taken after an event is known, will produce no more than random results. The strong form holds that stock prices fully reflect not only public information but also privately held information [Cohen87, 145].

3.2 Source of randomness

Financial time series are often cited as examples of randomness, but the statement of randomness is an empirical one. Behaviour that was previously believed random might become predictable with a better understanding of underlying dynamics, better measurements of market behaviour, or more advanced computational power to process a sufficient amount of information. The cause of unpredictability or randomness is often ignorance. If we do not know the forces that cause something to change, then we cannot predict its behaviour. The motion of roulette balls can be predicted using simple physical laws. With roulette, a sufficient degree of accuracy is achieved, giving a significant advantage to the house. Some classical examples of randomness are no longer random, if the knowledge of the system's dynamics is available and enough information about the state of the system is known [Farmer88, 100f].

In addition to ignorance a cause of randomness in economics is complexity. In an economy there are so many different agents at work that it is impossible to keep track of them all, much less learn the forces that describe their interactions. But even very complex systems may have collective modes of behaviour that are described by simple laws. Randomness can also occur because a small cause at one time determines a large effect at a later time. Imperceptible errors are amplified to macroscopic proportions, so that the system is random over long periods of time, even if it is described by simple laws. This, in fact, is chaos which is described as a sensitive dependence to initial conditions that occurs in a sustained way [Farmer88, 101].

The complexity and the sensitive dependence to initial conditions of financial markets is discussed in detail in the following chapter which focusses on the financial market as a nonlinear dynamic and adaptive system. The next paragraph describes the limits of the efficient market theory from an economic viewpoint.

3.3 Criticism of the efficient market theory

The success and acceptance of the efficient market theory lies in the fact that very few investors or speculators are able to beat the market, meaning that if investors cannot beat the market then is it that their stocks cannot outperform the market index. A broadly diversified mutual fund of common stocks achieves the return with less risk than a single stock

[Samuelson89, 254f.]. This strongly supports that the market portfolio is a very good portfolio, but does not prove the efficient market hypothesis.

Most academics favour the efficient market hypothesis, while practitioners believe that past stock prices show foreseeable trends and therefore help to forecast the market based on past performance [Cohen87, 14]. But academics tend to ignore their arguments for the efficient market theory when they try to build econometric models for forecasting foreign exchange rates and interest rates [Samuelson89, 217ff.]. There is no reason why the efficient market theory should only be valid in stock markets. The main assumptions of the efficient information distribution must be true for the whole economy, or the assumptions are not valid.

The main assumptions of the efficient market theory are rational behaviour of the market agents and transparency. This means that all information is freely available. These two axioms of economics are not perfectly true in reality. The market agents do not have unlimited power of computation and cogitation which would guarantee optimal and rational behaviour, and tall information has a cost. The assumption of rational behaviour can also not be attained because the actors have different time horizons and different risk aversions [Olsen92, 1ff.]. These weak assumptions make it difficult to trust the efficient market theory.

The mixed empirical research results also suggest that the capital markets are not so hyper-efficient as to make modelling of financial markets worthless. The so called very efficient foreign exchange market has a time lag up to 3 minutes between the future price in Chicago and the spot price of the Reuters FXFX page [Martens98] or also in the capital inefficiency can be found [Shin98]. On the other hand, the different studies suggest that the markets are efficient enough to warn any investor that better-than-average results will be difficult to achieve [Cohen87, 150]. This corresponds with the view on Wall Street. Peter L. Bernstein, a respected Wall Street practitioner, suggests this view: "The stock market in particular is highly efficient in rapidly incorporating information that would have an effect on prices in the short run, even it fails to process more complex and longer-run information in an efficient manner. This is precisely why so many investors who follow the news have such disappointing results." [Cohen87, 147]

It seems difficult to forecast financial markets; however, the situation is not hopeless. Forecasting is not possible simply with fundamental or technical technologies, but the use of complex methodologies is very promising for the task of forecasting financial markets. The next chapter discusses the application of system dynamics to financial markets, expanding the view of financial markets and dispelling the shortcomings of the efficient market theory.

4. Markets as nonlinear dynamical and adaptive Systems

Many natural phenomena show a complicated, aperiodic dependence on time. Recent developments in non-linear dynamic systems help to explain some of the natural phenomena in physics and chemistry. Non-linear, dynamic and adaptive systems can perhaps also help to explain the behaviour of the global economy, or financial markets as an important part of it behind the traditional linear economic models. This chapter explores different possibilities and limits of the application of the non-linear dynamic systems in finance.

4.1 System

A natural system is never isolated from the rest of the universe, this statement can also be applied to financial systems. The following example should show some relations between the financial market and the rest of the universe, as in chapter two the economic relation between the commodity and factor markets has been seen. The financial market as part of the factor market depends on the commodity markets, the offers in the commodity market depend on the industrial and agriculture production. The agriculture production depends on the weather, and so on.

Therefore, what we call a natural or financial system is an idealisation, where a certain number of parameters x_i are chosen to describe the system. The rest of the universe appears as a perturbing noise ω. The discrete time evolution of such a system is:

$$x(t+1) = f(x(t), \omega(t))$$

$x(t)$ is a vector with a number of components x_i. The external variables are subsumed in $\omega(t)$. The decision is arbitrary to put the different variables in $x(t)$ or $\omega(t)$. One can eliminate ω by studying the evolution of the entire universe, and put all variables in $x(t)$. The resulting equation would be intractable, or $x(t)$ can be very small, and the external noise $\omega(t)$ (often called shocks in economics) produces the aperiodic time evolution. But it is also possible to explain complex aperiodic behaviour like prices in financial markets in absence of shocks [Ruelle88, 195ff.].

The next paragraph explains the possibilities of non-linear system dynamics in absence of noise.

4.2 Nonlinear deterministic dynamics

This paragraph discusses the possible presence of low dimensional deterministic dynamics in financial markets and analyses some results from empirical research. It also looks for the evidence of chaos in financial markets, as a special case of deterministic dynamics and a fashionable topic in finance.

A dynamical system is said to be deterministic if it can be predicted exactly. The future of a stochastic time evolution is only partly determined by past values, so that exact predictions are impossible and must be replaced by the idea that future values have a probability distribution which is conditioned by a knowledge of past values. A stochastic process can be described as a statistical phenomenon that evolves in time according to probabilistic laws [Chatfield89, 5ff.].

The demonstration of low dimensional chaos in finance would be an important finding. It explains irregular oscillations of the system and limits the predictability of its future. Financial markets can be described as deterministic systems with some added stochastic noise, but separation between the noise and deterministic part of the evolution is ambiguous, because one can interpret noise as a deterministic time evolution in infinite dimension. This description is meaningless unless one put conditions to the deterministic part. The deterministic part of the time evolution must be low dimensional, if one is to be able to reconstruct the dynamics from experimental data [Ruelle94, 466ff.].

Differentiable dynamics is the study of deterministic time evolution of the form

$$x(t+1) = f(x(t))$$

where time t is discrete. The time evolution is deterministic in the absence of noise or external shocks. The point $x = x(t)$ runs over a phase space $\mathbf{R}^n$ with possibly infinite dimensions and the curve $x(t)$ is called a trajectory in phase space. The function f is assumed to be differentiable [Ruelle88, 197].

The phase space trajectory can describe what is called an attractor. The simplest attractor is the *fixed point* attractor, which means the following. $x(t)$ converges to $\bar{x}$ as t tends to infinity, where the average $\bar{x}$ does not depend on time. $\bar{x}$ is a stable equilibrium in physics or mechanics, but not necessarily in economy. The economists use the word equilibrium to denote a state which is not necessarily time independent.

There are also *periodic attractors* in which the phase space trajectory follows a closed curve forever. The distance of $x(t)$ and $\bar{x}(t)$ tends to be zero as t goes to infinity for a periodic attractor.

$$\bar{x}(t+T) = \bar{x}(t)$$

where T is the period.

A more complicated attractor is not a point or a periodic attractor. The orbit of $x(t)$ lies asymptotically on a more complicated set than a point or a loop. A special kind of more complicated attractors is the *strange attractor* which has a sensitive dependence on the initial conditions $x(0)$ [Ruelle88, 198].

The Lyapunov exponent measures the dependence on initial conditions for the time evolution. An infinitesimal change $\delta x(0)$ to the initial condition will be at time t a corresponding change $\delta x(t)$. There is a sensitive dependence on initial conditions if $\delta x(t)$ grows exponentially with t:

$$|\delta x(t)| \approx |\delta x(0)| e^{\lambda t}$$

where $\lambda > 0$. λ is called the Lyapunov exponent. It calculates the rate at which nearby trajectories in the phase space diverge. Chaos is the exponential growth of $\delta x(t)$ for almost all initial conditions x(0), and almost all $\delta x(0)$. Chaos is a sensitive dependence of initial conditions, no matter what the initial condition is [Ruelle90, 241f.].

The experimental estimation of the Lyapunov exponent can also be made from the reconstructed phase space, but it needs a lot of datapoints N. The number of points N needed to estimate Lyapunov exponents is about the square of that needed to estimate the dimension [Eckmann92, 187].

The dimension of the system is called information dimension. The information dimension is dependent on the complexity of the system being studied, and means something like the number of degrees of freedom. The number and possible components of the vector $x(t)$ and the equations of motion are not known for financial markets. One can do experiments in the natural sciences for the definition of the system parameters, but in finance the possibility of conducting experiments is poor.

The missing knowledge of the system parameters makes it impossible to construct the true phase space of the system dynamics, but it is possible to reconstruct the phase space of a dynamic deterministic system from one observable variable [Peters91, 151]. The reconstruction of a phase space from the price time series is possible, because the price is an always observable variable of financial markets. But the reconstruction of a phase space from the price time series is only possible, if financial markets have a low-dimensional non-linear deterministic system dynamics. This means the absence of any shocks or noise in financial markets. This assumption of a closed system dynamic is restrictive for the financial system [Zimmermann94, 20], but some empirical findings of low dimension dynamics in financial markets support this assumption [Scheinkman89, 311ff.; Peters91, 151ff.] what could be an artefact.

The reconstruction of a phase space from the price time series leads to the experimental determination of the dynamic system parameters of the markets. The original system dynamic is embedded with delay co-ordinates of an observable time series of the system [Farmer88, 104]. The idea of the time-delay method for reconstructing the dynamics is [Ruelle88, 200]:

Imagine we know only $u(t) = \psi(x(t))$, where ψ is some scalar observable of a time evolution $t \to x(t)$ in the phase space $\mathbf{R}^n$. The time evolution

$$t \to \tilde{x}(t) = (u(t), u(t+\tau), \ldots, u(t+(m-1)\tau)) \in R^m$$

is often a faithful image of the original time evolution if the delay time τ is well chosen and the embedding time dimension m is large enough. The resulting vector $\tilde{x}(t)$ will contain enough information to uniquely determine the system dynamics [Farmer88, 105]. It is possible to estimate the dimension of the reconstructed attractor. The dimension of the reconstructed attractor becomes equal to the dimension of the true attractor for sufficiently large embedding dimension.

The calculation of the information dimension of the reconstructed phase space of financial markets is now possible from price time series. Instead of calculating the information dimension directly from observable time series, what is possible, the most empirical work in finance calculates the correlation dimension as an approximation of the information dimension rather than the information dimension itself [Ruelle90, 244].

The correlation dimension can be calculated with help of the Grassberger-Procaccia algorithm. This algorithm is widely used in empirical analysis, for example see [Scheinkman89, 311ff.; Peters91, 151ff.], but there are limitations to the algorithm. The Grassberger-Procaccia algorithm finds necessarily:

$$d_{max} = 2 \log N$$

where d is the correlation dimension and N the number of data points in the time series [Eckmann92, 185ff.]. This fact casts some doubts concerning the empirical findings of low dimensional dynamics of 6 for different financial markets, because the used time series consists of about thousand datapoints [Scheinkman89, 311ff.; Peters91, 151ff.]. The empirical findings of low dimensional dynamics in financial markets could be a statistical artefact what is also supported by investigations with large data sets [Guillaume94, 9].

The correlation dimension allows to estimate the complexity of the attractor, but it gives no information about the development of the system through time. The large time behaviour of x shows different characteristic trajectories in the phase space, or different attractors.

The conclusion is: to extract useful dynamic information from time series (dimensions, Lyapunov exponents), long time series of high quality are necessary [Eckmann92, 187], but the time series in economics and finance are short and noisy [Brock91, 106]. Financial markets evade a quantitative description by low dimensional chaos, but these are chaotic in the qualitative sense that a small change in the initial conditions would have important consequences for later evolution. Some well-controlled systems present chaos. It is possible to study scenarios in detail by which this is reached. For some less well-controlled systems, like financial markets, one may confidently infer the presence of chaos, even when a precise demonstration is not practically possible. There is no convincing model of deterministic time evolution in finance, but it is reasonable to think of a sensitive dependence of initial conditions [Ruelle90, 243f.]. A useful model for financial markets could also include noise or shocks and drift of the deterministic dynamics, if there is a qualitative element of low-dimensional dynamics. Drift stands for a change over time in some parameters of the deterministic part of the dynamics [Ruelle94, 466ff.]. By introducing a number of ad hoc parameters (obtained from observation) one could get a more realistic and reliable model of financial markets [Ruelle94, 466f.] than with deterministic techniques [Brock91, 106].

4.3 Adaptive systems

This paragraph discusses qualitatively the origins of the possible drift component of the system description. This drift component of the system is the main difference with respect to physical dynamical systems, and constitutes on major limitation for the unreflected application of non-linear dynamical systems theory in finance.

The overall direction of financial markets is determined by the interaction of many dispersed agents acting in parallel. The action of any given agent depends upon the state of the market which depends upon the previous actions of the agents. There are rarely any central global controls or coordinations which are provided by the market mechanisms. The economy has many levels of organisation and interaction. A global market does not exist, much more are there different markets, which are interconnected with all sorts of interactions. These organisations are revised continually as the system accumulates experience. The adaptation of the system through accumulation of experience is driven by the adaptation of the agents, and innovations. Innovations are continually created by new technologies. Upon the innovation the financial markets operate far from a perfect behaviour, and improvements are always possible, and occur regularly [Holland88, 117f.].

The building blocks of traditional market theories are fixed rational agents that operate in a linear, static, statistically predictable environment, but in reality the interactions of the agents are characterised by limited rationality, adaptation (learning), and increasing returns [Holland88, 118]. Rational economic behaviour is the maximisation of subjectively expected utility, but the market participants are not the mythical heroes called rational men whose power of computation and cogitation are unlimited. Therefore they often use satisfactory rules of thumb as their internal model instead of truly optimal behaviour. Even when the agents know the basic structure of financial markets, they cannot derive reliable long-term predictions, and thereby form any meaningful rational expectation of the future, if there is a sensitive dependence of initial conditions in the system dynamic of financial markets [Hodgson91, 2]. The rules of thumb (or models of the agents) play a role in finance quite similar to the role of a strategy in the theory of games. The feasible strategies in non trivial games are not optimal, because the min-max strategies for these games are unknown [Holland88, 119]. Financial markets selects the market participants through the capital resources of the individual investor or the competition of institutional investors. Loosing investors disappear from the markets, and performing investors grow through the accumulation of money and imitators. Selection drives agents toward fitness maximisation. It is always possible to enhance the behaviour through learning, because the agents act only nearly rationally. But the adaptation of the market participants is slow, because there are structural constraints of the markets and the agents themselves. This constraint, for example laws, infrastructures, and others, slow down the speed of adaptation and can be a selection criterion [Selten91, 4ff.]. The microeconomics adaptation of agents drives the macroeconomics structural evolution of financial markets.

The structural evolution of economy has preoccupied many leading economic theorist since Adam Smith, especially Marx, Keynes, and Schumpeter. The engine of change is the pursuit of profit of the agents. Consequently, the economy not only generates perpetual motion but also continually alters its own structure. It can be considered either as a single agent altering its structure or as selection of new species through competitive survival or disappearance. Schumpeter regards business cycles as a form of economic growth. Now one can try to explain the business cycles and the structural change of economy with endogenous deterministic system dynamic, or exogenous with technical innovations. A definitive empirical conclusion is not possible, because the economic time series are too short. Kondratiev brings history into economics with his long-term business cycles of economic growth, which are in close relation to major technical innovations, and shows a kind of regularity, but normally innovations are not arriving in a deterministic way [Goodwin91, 138ff.].

In physics and chemistry models work well because the individual elements that make up the system must obey fixed laws which govern their behaviour. When molecular collisions are the source of the interaction, the collision invariants determine the observed behaviour of the system. In that case, prediction or explanation is based on the fact that when molecules of the right kind collide fast enough, then the reaction occurs. The mechanisms are fixed, and the molecules never learn or adapt. In the mechanical view, predictions can be made by simply running the equations forward in time, and studying where they lead. Does the system display chaotic behaviour or tends it to an equilibrium? Fascinating stuff, but of course only of any significance if the equation remains a good description of the system. The deterministic system dynamics that runs as a model of the financial markets will only be a good description for as long as there is no change in the significant variables, in the market mechanisms, or in the agents behaviour [Allen94, 6ff.]. Therefore a frequent change of the deterministic part of a financial market model, a trading model based on moving average with bandwidth, and a portfolio management model, has the potential for bringing more profit than a stable model without adaptation [Phang94, 125ff.]. The variation and selection of agents make markets more similar to biological systems with evolutionary concepts than to physical systems [Hanappi94, 7ff.]. Therefore theoretically it is arbitrary to define the economic evolution as endogenous or exogenous variables, but practically it is difficult enough to define a stable deterministic system dynamic without adaptation as an economic model. For that reason a definition of financial markets as a dynamic deterministic system with shocks (noise), like politics and natural events for example, and drift seems to be most promising for the future work.

5. Statistical tools

Financial markets are adaptive dynamical systems with external shocks. Therefore, statistical tools have to be able to detect the different components in the empirical analysis of financial time series. This chapter describes the methods of a series of univariate and multivariate, linear and non-linear statistical tests which identify the relevant indicators for modelling and forecasting financial markets. In a multivariate system it is possible that the variation in one time series explains the variation in another series. This may lead to a deeper understanding of the mechanism which generated a given time series. It is useful to distinguish four different kinds of influences on a given financial time series or financial system. Different statistical tools allow for the detection of different influences. Therefore, this chapter will examine a variety of analytical systems. Linear univariate, non-linear univariate, linear multivariate, and non-linear multivariate systems will be discussed. The first paragraph is a short explanation of the properties of a time series.

5.1 Time series

A time series is a collection of observations made sequentially in time. The special feature of time series analysis is the fact that successive observations are usually not independent and that the analysis must take into account the time order of the observations. Many time series arise in finance, but the most interesting and most often used time series are the price time series. Price time series are discrete, meaning observations are taken only at specific times. The observations are usually taken at equal intervals, for example hourly, daily, weekly or monthly time series. Time series using unequal intervals of time also exist in finance. For example, high frequency time series contain every price or quote of the market and therefore do not adhere to equal intervals. There are more or less active phases in the markets with more or less prices depending on higher or lower activity [Chatfield89, 1ff.]. This work is mainly concerned with discrete time series measured at equal intervals of time.

Plotting the observation against time is the first step in any time series analysis. A straight plot of the predictor time series and the first difference can show trends, cycles, seasonal, outliners and other effects of the series. For some series, the variation is dominated by obvious features, and a fairly simple model, which only attempts to describe trend and seasonal variation, may be perfectly adequate to describe the variation in time series. A graph also enables one to look for outliners which do not appear to be consistent with the rest of the data. The outliners may be a perfectly valid but extreme observation. Alternatively, the outliner may be a freak observation arising, for example, when a recording device malfunctions. The treatment of outliners is a complex subject. A freak outliner needs to be adjusted before further analysis of the data can be made. A robust analysis is a statistical method that is designed to be insensitive to outliners [Chatfield89, 5f.].

Available time windows, which means the lenght of the available time series for analysis, should span most types of system states like bull, bear, and congested financial markets. A bull market means a upward trend of the price time series, a bear market means a downward trend, and a congested market is a trendless market dynamic.

A useful way of describing a time series is to identify the moments of the process. Particularly the first and second moments, which are called the mean, variance, and autovariance functions respectively. A process is called stationary if the moments are constant through the time evolution [Chatfield89, 28].

5.2 Linear univariate analysis

A linear system converts an input series to an output series by a linear operation [Chatfield89, 6]. The *autocorrelation coefficients* measure the linear correlation between observations at different distances. These coefficients often provide insight into the probability model which generated the time series. Given T observation $x_1, \ldots, x_T$ of a discrete time series the *autocovariance* coefficient c at lag k is given by [Chatfield89, 19f.]:

$$c_k = \frac{1}{T}\sum_{t=1}^{T}(x_t - \bar{x})(x_{t+k} - \bar{x})$$

where
$\bar{x}$ is the average value of the time series x_t.

If x_t and x_{t-k} are independent, then the covariance is zero. If they are not independent, then the covariance may be positive or negative depending on whether high values tend to go with high or low values.

Covariance is difficult to interpret as it depends on the units in which x is measured. Thus it is often useful to standardise the covariance. The standardised quantity called the autocorrelation coefficient r at lag k is given by:

$$r_k = c_k / c_0$$

The correlation coefficient must lie between ± 1 and is a useful measure of the linear association between two variables. An aid in interpreting a set of autocorrelation coefficients is a graph called a correlogram in which r_k is plotted against the lag k. If a time series has no linear autocorrelations, then for large T, r_k should be more or less equal to zero. The 95% confidence interval for r_k is

$$-1/N \pm 1.96/\sqrt{N}$$

because r_k is approximately N$(0, 1/T)$ for a random time series. This spotlights one of the difficulties in interpreting the correlogram, in that a large number of coefficients is quite likely to contain one or more significant results, even when no real effects are present [Chatfield89, 20f.].

Stationary time series often exhibit short-term autocorrelation characterised by a fairly large value of r_1 followed by a few further significant coefficients which tend to get successively smaller. The correlogram of an alternating time series, with successive observations on different sides of the overall mean, also tends to alternate. If a series contains a trend, and therefore is non-stationary, then the values of r_k will not decrease. In fact the autocorrelation function is only meaningful for stationary time series, so any trend should be removed before calculating it. If a time series contains a seasonal fluctuation then the correlogram will also exhibit an oscillation at the same frequency. Therefore, the seasonal variation should also be removed before computing the autocorrelation coefficients. It is also necessary to adjust outliners because the autocorrelation function is not very robust [Chatfield89, 21ff.].

5.3 Non-linear univariate analysis

5.3.1 The BDS test

A test of the existence of potentially forecastable structure, nonstationarity, or hidden patterns is the Brock, Deckert, and Scheinkman (hereafter denoted BDS) test. The BDS test is especially useful when one has no idea what kind of hidden structure to expect in the time series [Brock91, 4ff.]. The BDS statistic is a test for nonlinearity. It tests the null hypothesis that a time series x_t is independent and identically distributed against an unspecified alternative using a nonparametric technique [Brock91, 41f.]. Therefore, it is a univariate test that detects linear and non-linear stochastic, and linear and non-linear deterministic systems, however the test is unable to distinguishes between the kind of systems [Hsieh91, 1839ff.].

The BDS test basically measures the statistical significance of the correlation dimension calculations. According to the BDS test, the correlation integrals should be normally distributed if the time series under study is independent. The correlation integral is the probability that any two points are within a certain length e apart in phase space. The correlation integrals are calculated according to the following equation [Peters94, 246f.]:

$$C_m(e) = (1/N^2) \sum_{i,j=1}^{T} Z(e - |x_i - x_j|)$$

where
i not equal j
$Z(x)$ is 1 if $x > 0$; 0 otherwise
T = number of observations
e = distance
C_m = correlation integral for dimension m.

The function Z counts the number of points within a distance e of one another. The phase space is reconstructed by lagging the time series, allowing for the calculation of the correlation integral without knowledge of what the true phase space looks like. The correlation integral fills the space of whatever dimension it is placed in. Therefore the BDS statistic w is normally distributed with a mean of 0:

$$w_N(e,T) = \left| C_N(e,T) - C_1(e,T)^N \right| * \sqrt{T} / s_N(e,T)$$

where
$s_N(e,T)$ = the standard deviation of the correlation integrals

Thus the BDS statistic w has a standard normal probability distribution. The null hypothesis of independence and identical distribution can be rejected with 95 percent confidence, if it is greater than 2.

Monte Carlo simulation recommends to chose e between one-half and three halves of the standard deviation of the data, and m between 2 and 5 for small data sets (200 to 500 observations), and up to 10 for large data sets (at least 2'000 observations). A calculation of the BDS test is difficult and not reliable with less than 200 datapoints [Brock91, 169f.].

The rejection of the null hypothesis of independence and identical distribution claims further tests to distinguish between the different systems. These tests are the autocorrelation function for linear stochastic systems and the cumulated periodogram for non-linear stochastic systems.

5.3.2 The cumulated periodogram

The *cumulated periodogram* is a non-linear parametric test. It discovers hidden periodicities in a time series, and therefore it detects non-linear stochastic systems. It tests the null hypothesis that a time series x_t is from a white-noise process or the alternative hypothesis that a time series x_t is from an unspecified alternative process [Schlittgen87, 266ff.]. The periodogram can be directly calculated from the data by

$$I(\omega_p) = \left[(\sum x_t \cos \omega_p)^2 + (\sum x_t \sin \omega_p)^2 \right] / T\pi$$

where
$\omega_p = 2\pi p / T$ is often called the pth harmonic [Chatfield89, 20f.].

The cumulated periodogram is calculated with the following equation [Schlittgen87, 274]:

$$S_r = \frac{\sum_{p=1}^{r} I(\omega_p)}{\sum_{p=1}^{M} I(\omega_p)}$$

where
$M = [T/2]$

The test statistic of the cumulated periodogram is:

$$C = \max \left| S_r - \frac{r}{M} \right|$$

The null hypothesis is rejected with (1-α) confidence, if C > c, where c is approximated:

$$c = \frac{\sqrt{-0.5\ln(\alpha / 2)}}{\sqrt{M-1} + 0.2 + \frac{0.68}{\sqrt{M-1}}} - \frac{0.4}{M-1}$$

A graph of the cumulated periodogram and c shows clearly and simply the results of the test.

5.4 Linear multivariate analysis

Thus far, this work has been concerned with analysing a single time series. We now turn our attention to the situation where we observe more than one time series and examine relationship between them. Our main interest lies in finding causal relations between different time series. One series is regarded as input to a system, while the other series is regarded as output. The input time series y_t is called indicator time series, and the output series x_t is called the predictor time series. The indicator time series is lagged for the purpose of forecasting. The cross correletion function allows for the detection of linear input-output systems. It measures the linear dependencies between the indicator and the predictor like the autocorrelation function in the univariate case.

The *cross correlation function* $\rho_{xy}(k)$ is defined by

$$\rho_{xy}(k) = \gamma_{xy}(k) / \sqrt{\gamma_{xx}(0)\gamma_{yy}(0)}$$

The cross correlation function is the standardised cross covariance function $\gamma_{xy}(k)$ which is defined by

$$\gamma_{xy}(k) = Cov(x_t, y_{t+k}) = \sum_{t=1}^{T} (x_t - \bar{x})(y_{t-k} - \bar{y}) / (T-1)$$

The interpretation of the cross correlation function is similar to the autocorrelation function [Chatfield89, 136ff.].

5.5 Non-linear multivariate analysis

Non-linear nonparametric multivariate dependencies can be measured with an artificial neural network test [Ankenbrand96, 28f.]. This neural network test is based on a sensitivity analysis which is calculated with the aid of an artificial neural network. Neural networks and the sensitivity analysis are described in the next chapter. The next paragraph introduces the bootstrap methodology which is used for calculating the test statistics of the neural network test.

5.6 The bootstrap

The bootstrap methodology answer the question how accurate are data summaries like the results of our neural network test. This question constitutes part of the process known as statistical inference. The bootstrap is a computer based method of statistical inference that helps to answer many statistical questions. It requires modern computer power to simplify the often complex calculations of traditional statistical theory. The bootstrap is a data-based simulation method for statistical inferences. It gives the direct appreciation of variance, confidence interval, and other probabilistic phenomena and computes the accuracy of estimated ratios with a minimum of mathematical assumptions and without complicated mathematical formulas [Efron93, 1ff.]. The bootstrapped is well suited for statistical inference of different economic problems [Godfrey98; Kanas98; Li98].

The idea behind the bootstrap is an old one. If we wish to estimate a functional of a population distribution function F, such as a population mean for example

$$\mu = \int x dF(x),$$

consider the same functional of the sample (or empirical) distribution function $\hat{F}$, which in this instance is the sample mean

$$\bar{x} = \int x d\hat{F}(x).$$

The empirical distribution is that probability measure that assigns to a set measure equal to the proportion of sample values that lie in that set. The same size-size resamples may be drawn repeatedly from the original sample, the value of a statistic computed for each individual resample, and the bootstrap statistic approximated by taking an average of an appropriate function of these numbers [Hall92, 1].

The bootstrap confidence intervals are directly constructed from real data sets using a simple iterative computer algorithm. This allows for the computation of confidence intervals from every ratio, even if it is not possible with the traditional statistical inference theory. The use of the term 'bootstrap' derives from the phrase, 'to pull oneself up by one's bootstrap', widely thought to be based on the adventures of Baron Munchhausen a fictional character created by Rudolph Erich Raspe. In one story, the Baron had fallen to the bottom of a deep lake. Just when it looked like all was lost, he thought to pick himself up by his own bootstraps [Efron93, 1ff.].

The basic idea behind the bootstrap method for the this work is simple. The bootstrap estimates the standard error of a statistic of interest $s(x)$ from the observed data points $x_1,...,x_T$. A bootstrap sample $x^* = (x_1^*,...,x_T^*)$ is obtained by randomly sampling T times with replacements from the original data points. The bootstrap algorithm generates a number B of independent bootstrap samples $x^{*1},...,x^{*B}$. Typically the number of bootstrap samples range from 50 to 200 for standard error estimation. The approximation improves as the number of resamples increases. Corresponding to each bootstrap sample is a calculation of s. The bootstrap estimate of standard error is the standard deviation of the bootstrap replications of s, given by [Efron93, 12ff.]:

$$se_s = \sqrt{\frac{\sum_{b=1}^{B} (s_b - \bar{s})^2}{B - 1}}$$

For most statistics there is no formula for calculating the standard error, but in fact, no formula is needed. Bootstrap is a nonparametric estimator of standard errors for every statistic directly derived from the data.

This method is used for the estimation of accuracy of the neural network test in the next chapter.

6. Artificial neural networks

This chapter describes neural networks and their possible application in financial markets. The first paragraph gives a short overview of artificial intelligence, which is the super-field of artificial neural networks, and its relation to economy. The second paragraph gives an overview of artificial neural networks and its possible applications in time series analysis. The next paragraph describes in detail the feed-forward neural networks used in this work. The fourth paragraph discusses the learning algorithms for optimisation of the used neural network. Described in the fifth paragraph is the computation of sensitivity analysis. This computation is at the core of the multivariate non-linear test for the identification of relevant input indicators of the market model, which is the final missing part of the forecasting methodology.

6.1 Artificial intelligence

There have been many definitions of artificial intelligence. A particulary useful one states:

"Artificial Intelligence is a field of science and engineering concerned with the computational understanding of what is commonly called intelligent behaviour, and with the creation of artefacts that exhibit such behaviour" [Shapiro92, 54].

The three main fields of artificial intelligence are: computational psychology, computational philosophy, and advanced computer science. Computational psychology tries to understand intelligent human behaviour by creating computer programs that behave in the same or similar way people do. The goal of computational philosophy is to form a computational understanding of intelligent behaviour, without being restricted to the similiarity of the algorithms and data structures actually used by the human mind. Advanced computer science studies push the frontier of what we know and how that knowledge is used to program computers. This goal led to one of the oldest definitions of artificial intelligence: the attempt to program computers to do what until recently only people could do. Another way of distinguishing artificial intelligence as a field in its own right is by noting the interest and greater use of heuristics rather than algorithms. Heuristic is a problem solving procedure that fails to be an algorithm for any reason [Shapiro92, 54].

There are a lot of different definitions and subfields of artificial intelligence. A formal framework for these different approaches varies along two main dimensions. The first dimension distinguishes between thinking and acting. The second dimension differentiates human and rational behaviour, where rationality is defined as an ideal concept of intelligence. A human-centred approach must be an empirical science, involving both hypothesis and experimental confirmation. But there do not exist systems these are similar like humans. A rationalist approach involves a combination of mathematics and engineering. The framework is organised into four categories [Russell95, 4ff.]:

- Systems that think like humans.
- Systems that act like humans.
- Systems that think rationally.
- Systems that act rationally.

Let us look at each system in more detail. The Turing Test was designed to provide a satisfactory operational definition of systems that acts like humans. Turing defined intelligent behaviour as the ability to achieve human-level performance in all cognitive tasks. In order to test the system, the computer is interrogated by a human via teletype. The computer has been successfully programmed if the interrogator cannot tell if there is a computer or a human at the other end. The computer needs to possess the following capabilities: natural language processing to enable it to communicate successfully with the interrogator, knowledge representation to store information, automated reasoning to use the stored information, and machine learning to adapt and to detect patterns. The cognitive approach tries to model systems that think like humans. If we are going to say that a given program thinks like a human, we must have some way of determining how humans think. The General Problem Solver by Newell and Simon was not intended to solve problems correctly. Newell and Simon were more concerned with comparing the path of the computer's steps of reasoning to human subjects and their patterns of reason for solving the same problem. The interdisciplinary field of cognitive science brings together computer models from artificial intelligence and experimental techniques from psychology in an attempt to construct theories on the workings of the human mind. The distinction between thinking like humans and rationally thinking initiates the laws of thought approach. The so-called logistic tradition is based on the field of logic. The rational agent approach tries to create model systems that act to achieve certain goals, given a particular set of beliefs. An agent is defined for these purposes as something that perceives and acts in response to stimuli. In this approach, artificial intelligent is viewed as the study and construction of rational agents.

In the laws of thought approach with respect to artificial intelligence, the emphasis is on correct inferences. Making correct inferences is part of being a rational agent, because one way to act rationally is to reason logically. However correct inference does not make up all of what rationality is, because there are often situations where there is no obvious correct answer, yet something still must be done. Achieving perfect rationality is not possible in complicated environments. The computational demands of creating such environments are too high. Acting appropriately when there is not enough time to do all the computations one might like, is called limited rationality [Russell95, 7f.]. The concept of limited rationality of agents in financial market is well known, and it has been described in previous chapters in this work within different aspects of economics and subjects. Perhaps artificial intelligence is concerned with determining courses of action for problem solving which are too large or complex for the application of constrained optimisation algorithms. So far to date, financial markets seems to be very complex and simple constrained optimisation algorithms fail to offer a an adaquate solution, therefore, it would appear that artificial intelligence techniques offer the possibility of extending the models of financial markets [Moss92, 1].

Although the dream of creating intelligent artefacts has existed for many centuries, artificial intelligence had its birth at a conference held at Dartmouth College in the summer of 1956. Although artificial intelligence itself is a young field, it has inherited many ideas, viewpoints, and techniques from other disciplines. Artificial intelligence is generally considered to be a subfield of computer science, but there are several disciplines outside computer science that strongly impact artificial intelligence and on which artificial inteligence studies have a strong impact [Shapiro92, 55]. Our theories of reasoning and learning have emerged from a 2000 year old tradition in philosophy. As a result of over 400 years of mathematics, we have formal theories of logic, probability, decision making, and computation. From psychology, we have learned the tools with which to investigate the human mind, and have created a scientific language within which to express the resulting theories. From linguistics, we have theories of the structure and meaning of language. Finally from computer science, we have the tools with which to make artificial intelligence a reality [Russell95, 8].

At the beginning artificial intelligence researchers were very optimistic, but different kinds of difficulties quickly arose. The first type of difficulty arose because early programs often contained little or no knowledge of their subject matter and succeeded by means of simple syntactic manipulations. Weizenbaum's Eliza program is an example of such a program. Apparently the program could engage in conversation on any topic, however, it merely borrowed and manipulated the sentences typed into it by a human. Another example occurred in early machine translation efforts, which failed because translation requires general knowledge of the subject matter. It is not sufficient to merely transform the text based on grammatical rules, and to replace words using a dictionary. The second kind of difficulty was the intractability of many of the problems that artificial intelligence was attempting to solve. The fact that a program can find a solution in principle does not mean that the program contains any of the mechanisms needed to find it in practice. The illusion of unlimited computational power was not confined to problem solving programs. Early experiments in evolutionary algorithms were based on the correct belief that by making a series of small mutations to a machine code program, one can generate a program for any task. The idea was to try random mutations and then to apply a selection process to preserve mutations that seemed to improve behaviour. Despite thousands of hours of computing time, almost no progress was demonstrated in fact a combinatorial explosion occured. But today there is revival heuristic of evolutionary methods, and evolutionary artificial intelligence is a reality like genetic algorithms and genetic programming. A third difficulty arose because of some fundamental limitations of the basic structures being used to generate intelligent behaviour. For example Minsky an Papert proved that although perceptrons could learn anything they were capable of representing. They could actually represent very few patterns. A two input perceptron could not be trained to recognise when its two inputs were different. However, their results did not apply to more complex multilayer neural networks [Russell95, 21f.].

Although the difficulties of artificial intelligence may appear daunting, a careful application of its techniques could expand the understanding of financial markets [Monostori97]. This work uses artificial neural networks as an artificial intelligence technique for financial time series analysis.

6.2 Neural networks

Artificial neural networks attempt to simulate the way in which the human brain processes information. The basic computing device in the human brain is the neuron. A neuron consists of a body, branching extensions called dendrites for receiving input, and an axon that carries the neuron's output to the dendrites of other neurones. The junction between an axon and a dendrite is called a synapse. A neuron is believed to carry a simple threshold calculation. It collects signals at its synapses and sums them. If the result exceeds a certain threshold, the neuron sends out its own signal. The output signal is a transformation of the total input signal which is determined by a non-linear function [Refenes95, 6]. Biological neural networks are actually much more complicated in their behaviour but can, as first approximation, be modelled by such simple automata.

A neural network is a structure made up of highly interconnected neurones. Associated with each connection between the neurones is a weight, and associated with each neuron is a state. Together the weights and states represent the distributed knowledge of the neural network. This organisation of neurones and weighted connections creates an artificial neural network system [Refenes95, 7].

A neural network learns by changing the internal connection strengths and possibly its topological structure. This is called weight adaptation, and takes place during a training phase. In this phase, external input patterns have to be associated with specific external output patterns or specific activation patterns. This is accomplished by changing the strength of the connections between the neurones [Refenes95, 7].

The interest in artificial neural network research has been brought about by two driving forces. Firstly, there was the realisation that artificial neural networks were powerful tools for modelling and understanding human cognitive behaviour. This has stimulated research in neuroscience, anatomy, psychology and the biological sciences, whose primary objective is to investigate the physiological plausibility of current artificial neural models and to identify new models which will give a more accurate insight into the functionality of the human brain. Secondly, artificial neural networks have powerful pattern recognition properties, and can outperform contemporary modelling techniques in many applications that make artificial neural networks an useful addition to the statistician's toolbox. Neural networks are analogous to nonparametric, non-linear regression models [Refenes95, 1ff.; Anders97].

Parametric models make assumptions about the properties of the data they are trying to fit. Nonparametric models, on the other hand, make no such a priori assumptions but let the data represent itself. In trying to find an appropriate model, nonparametric model estimators are exploring a much larger search space of functions to fit the observations [Refenes95, 5]. Artificial neural networks have been in use for some years now in order to model financial and economic time series such as foreign exchange market prices, stocks or economic indices [Refenes95, 101ff.]. On the whole the published (and unpublished) results appear to be at least as good as those obtained from customary linear statistical methods of the ARMA/ARIMA type and well known economic models like Tobins regression [Eakins98].

In finance and economics, as well as in the social sciences in general, one does not have equations describing the processes that one would like to study, as is often the case in physics. All that is available are measurements of the variables of interest, i.e. time series, which are often noisy and of dubious quality. It can be shown that a time-delay neural network can map a finite time sequence into the value that the sequence will have at some point in the future [Hassoun95]. That is, given a series of values of the variable x at time step t and at past time steps, $x(t), x(t-1), x(t-2),\ldots,x(t-m)$, we look for an unknown function F such that:

$$x(t+n) = F\left[x(t), x(t-1), x(t-2), \ldots, x(t-m)\right]$$

which gives us an n-step predictor of order m for the quantity x. For linear systems, the above reduces to the well-known autoregressive statistical method. However, most systems of interest in finance and economics are non-linear and for these systems the customary statistical approaches are inadequate. On the other hand, neural nets are intrinsically non-linear, non-parametric approximates, which makes them attractive for such prediction tasks. In practice, the time-lagged values *x(t), x(t-1), x(t-2),...,x(t-m)* are fed as inputs to the network which, once trained in the usual way on many significant input-output pairs, gives as output the predicted value for yet unseen past x values. This approach has been successful for chaotic time series prediction in non-linear dynamics and for financial predictions [Refenes95, 101ff.]. In practice, only a few steps ahead can be predicted with good confidence. As soon as n has a larger value, the prediction quality will degrade.

We will use artificial neural networks to predict future values of financial time series in the stock and bond markets. However, our approach is multivariate and takes into account the influence of important economic variables. In fact, the univariate, pure time-lagged model, that works well for phenomena showing well-defined deterministic chaos, might fail in the less precise economic world [Ruelle90, 241ff.]. Specifically, we look for an approximation of a function f of the form:

$$x(t) = f(I_1, I_2, \ldots, x(t - \Delta_i))$$

where $x(t-\Delta_i)$ are lagged values of the variable to be predicted and the I_j are fundamental indicators, possibly lagged, that might have an influence on x according to economic theory, and our statistical analysis called feasibility analysis.

The modelling of future prices is one possible goal. Traders in financial markets are often interested in risk-adjusted performance indicators what makes the forecasting of volatility and other risk parameters, like Sharpe ratio and stop loss levels very useful [Choey97, Atiya97].

6.3 Feedforward neural networks

This paragraph describes the feedforward neural network with its mathematical foundations. In feedforward networks the individual neurones are disposed into successive layers. Only neurones in neighbouring layers are connected and signals are only processed one-way from a set of input units (the input layer) to the output units (output layer), passing through a variable number of so-called hidden layers.

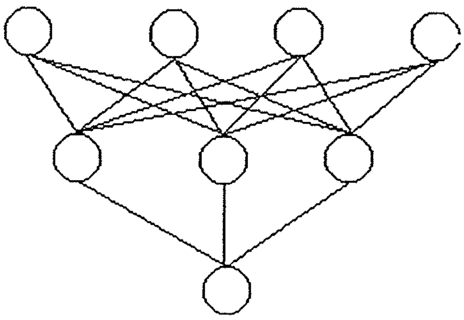

Each neuron may then compute its state once, at most. Each neuron i in the hidden and possibly output layers computes a weighted sum of the m signals coming to it through its input connections and then applies a (usually non-linear) transfer function g to the sum. The result is the output signal v_i of neuron i:

$$v_i = g(\sum_{j=1}^{m} w_{ij} v_j - \theta_i)$$

where the w_{ij} are the weights of the (forward) connections between neurones j in the previous layer and neuron i. θ_i is a threshold value for unit i. It is normally absorbed into the sum of weights by treating it as an additional constant value connection:

$$v_i = g(\sum_{j=0}^{m} w_{ij} v_j)$$

The function g is commonly a hyperbolic tangent or a sigmoid.

If the output units are linear rather than thresholding then it is possible to have arbitrary real numbers as output. This allows the parameterization of an unknown real-valued function, which will be needed in the applications described later. More generally, let us suppose that we want to approximate any continuous real-valued function $F(x_1, x_2, \ldots, x_p)$. A general result [Hassoun95] states that such a family of functions can be approximated to any desired accuracy by a feed-forward network with at least one single hidden layer. That is, there exist real constants α_i, w_{ij} and a monotone-increasing continuous function g (e.g. the sigmoid function) such that:

$$f(x_1, x_2, \ldots, x_p) = \sum_{k=o}^{p} \alpha_k g(\sum_{j=0}^{m} w_{ij} v_j)$$

and

$$\left| f(x_1, x_2, \ldots, x_p) - F(x_1, x_2, \ldots, x_p) \right| < \varepsilon$$

for any $\varepsilon > 0$, where m and p are the number of units in the input and hidden layers respectively. Being only an existence result, the theorem leaves unspecified the optimum number of hidden layers and the number of units in the hidden layers for any given case. This design specification of the neural network will be described in the next chapter which focuses on forecasting methodology.

6.4 Backpropagation

The principal goal of neural learning is to form associations between observed patterns. In supervised learning, the environmental feedback specifies the desired output pattern. The difference between the desired and actual output is known a priori in time series analysis for observed data. The objective of learning is to reduce as much as possible the difference between the actual and desired output patterns by minimising an error function [Refenes95, 7f.].

Backpropagation is the most popular method for performing the supervised learning task. It is a method for calculating derivatives in any large system made up of elementary subsystems or calculations which are represented by known, differentiable functions [Werbos94, 269f.].

Training the network involves systematically changing the weights until the network produces the desired output within a given tolerance. This is repeated over the available time series called the training set. A plausible measure of how poorly the network is performing with its current set of weights is given by the square error E over the training set of the time series in:

$$E = \sum_{t=1}^{T} E(t) = \sum_{t=1}^{T} \sum_{i=1}^{n} \frac{1}{2} \left[\hat{v}_i(t) - v_i(t) \right]^2$$

where $\hat{v}$ is the actual and v the target output [Werbos94, 274].

At the beginning the weights are initialised with small random values, however it may be better to guess the weights based on prior information, in cases where prior information is available. Next, the outputs $\hat{v}_i(t)$ and the errors $E(t)$ for that set of weights are calculated. Then the derivatives of E with respect to all the weights are calculated. If increasing a given weight would lead to more error, that weight is adjusted downward. If increasing a weight leads to less error, it is adjusted upward. After adjusting all the weights up or down, the process starts again, and continues until the rate of errors settles down [Werbos94, 274f.].

Backpropagation calculates the derivatives exactly for all weights with only one complete rotation through the system. The prefix *F_* indicates the ordered derivative of *E* with respect to whatever variable the *F_* precedes. Thus, for example:

$$F\_\hat{v}(t) = \frac{\partial E}{\partial \hat{v}_i(t)} = \hat{v}_i(t) - v_i(t)$$

which follows by differentiating the previous equation [Werbos94, 276f.].

The conventional chain rule for partial derivatives to calculate the derivatives of *E* with respect to all the weights can be a rigorous approach under certain conditions, but its generality is limited, and it requires great care to be taken with the side conditions. Backpropagation is based on a new chain rule for ordered derivatives:

$$\frac{\partial^+ TARGET}{\partial z_i} = \frac{\partial TARGET}{\partial z_i} + \sum_{j>i} \frac{\partial^+ TARGET}{\partial z_j} * \frac{\partial z_j}{\partial z_i}$$

where the derivatives with the superscript represent ordered derivatives, and the derivatives without superscripts represent ordinary partial derivatives. This chain rule is valid only for ordered systems where the values to be calculated can be calculated one by one in the order z_1, $z_2, \ldots, z_n$, *TARGET*. The ordered derivative represents the total impact of z_j on *TARGET*, accounting for both the direct and indirect effects [Werbos94, 275f.].

It is possible to use a simpler notation, because the ordered derivatives of the on target variable is calculated. One can write the ordered derivative of the *TARGET* with respect to z_j as F_z_i. The *TARGET* variable of interest in backpropagation is the error *E*. In that case this changes the appearance of the chain rule to:

$$F\_z_i = \frac{\partial E}{\partial z_i} + \sum_{j>i} F\_z_i * \frac{\partial z_j}{\partial z_i}$$

and in the case of the adapting the network weights

$$F\_x_i(t) = F\_\hat{v}_{i-N}(t) + \sum_{j=i+1}^{N1n} w_{ji} * F\_net_j(t)$$

where

$$F\_net_i(t) = g'(net_i) * F\_x_i(t)$$

where g' is the derivative of g the hyperbolic tangent or sigmoid function, and

$i = N + n, \ldots, m+1.$

This backward propagation of information is what gives backpropagation its name. The usual method to adapt the weights is to set

$$\Delta w_{ij} = -k * F\_w_{ij}$$

where

$$F\_w_{ij} = \sum_{t=1}^{T} F\_net_i(t) * x_j(t)$$

and where the learning rate k is a small chosen constant smaller than 1 [Werbos94, 277].

For correcting the magnitude of the learning impulse, the formula and the used software of [McClelland89, 121ff.] is changed to:

$$F\_w_{ij} = \sum_{t=1}^{T} F\_net_i(t) * x_j(t) * (1 - x_j(t))$$

For practical purposes the learning rate should be as large as possible without leading to oscillation. This offers the most rapid rate of learning. One way to increase the learning rate without leading to oscillation is to modify the backpropagation learning rule to include a momentum term m [McClelland89, 136f.]:

$$\Delta w_{ij}(t+1) = -k * F\_w_{ij} + m * \Delta w_{ij}(t)$$

The momentum term m determines the effect of past weight changes on the current direction of movement in the weight space that filters out high frequency variations of the error surface in the weight space.

The next chapter discusses some practical procedures for determining the learning rate and the momentum term.

6.5 Sensitivity analysis

Regardless of the quality of the prediction model, success in prediction is largely determined by the selection of the model indicators. One of the common criticisms aimed at artificial neural networks is that their modelling behaviour cannot be explained. It is not possible to say which indicator, or combinations thereof, have a greater impact on the outcome. This is purported to be one of the advantages of classical statistical models. In this paragraph it is shown that the modelling behaviour of feedforward neural networks can be explained by using a sensitivity analysis. Another possible way for explaining these networks could be the analysis of the weights [Sen95, 326]. But the present work is restricted to the sensitivity analysis. The analysis of modelling behaviour allows to use feedforward neural networks as a multivariate non-linear test for detecting dependencies between different time series, and for selecting the most promising indicators for forecasting models.

Firstly, as is usual, the feedforward neural network is trained with the possible indicators. To perform the sensitivity analysis, the ranges for the indicator values are first determined. For each indicator, the minimum, maximum, and the midpoint of the range are determined. The value of each indicator is varied one at a time, holding the values of other indicators fixed at the midpoints of their ranges. For each indicator being varied, the values are spread over ten equal intervals over its entire range of value. The neural network is then used to compute the output. The absolute network output is not as important as the changes in the network output. The changes in the network output indicate the sensitivity of the network output to changes in the indicator [Sen95, 334].

A sensitivity index is computed to find out the relative strength of the influence of an indicator time series to the outcome of the neural network. This index is computed by averaging the deviations in outcome for decile changes over the whole range of values of the indicator [Sen95, 336f.]:

$$S = \frac{1}{10}\sum_{i=1}^{10}(v_{i+1} - v_i)$$

The sensitivity index is interpreted to indicate that if all other indicators were held constant at around the average value, ten percent change in this indicator would affect the outcome of the neural network in the magnitude of the sensitivity index. The sensitivity index measures the dependency of the predictor time series from the indicator time series regardless of the kind of dependencies or relations.

The bootstrap of the input data allows computation of the test statistics of the sensitivity index and also immunises the neural network from the initial set-up problem, for example choosing the initial weights, because every bootstrap sample begins with new random weights. We use 200 bootstrap samples B for calculation 200 times the sensitivity index S with different data sets. The bootstrap produces this different data sets based on the original input data set. The expectation of S is calculated with the different sensitivity indices based on the bootstrap samples [Efron93, 14]:

$$E(S) = \frac{1}{B}\sum_{i=1}^{B} S_i = \bar{S}$$

with the standard error

$$se_{\bar{S}} = \sqrt{\frac{\sum_{i=1}^{B}(S_i - \bar{S})^2}{(B-1)}}$$

This allows one to compute test statistics like confidence intervals for the expectation of the sensitivity index also for problems with few data sets, and completes our multivariate non-linear test.

7. Forecasting methodology

This chapter presents an integrated approach for analysing and modelling the behaviour of financial markets. The methodology makes use of the assumptions and the elements described in the previous chapters. Its originality lies in the fact that it is based on statistical and macro-economic principles and it integrates fundamental economic knowledge in a multivariate non-linear time series artificial neural network model. The final model allows for the forecasting of financial time series with good results. This is shown in practice in the next chapter, where the methodology is successfully applied to two case studies.

Artificial neural networks are applied to support technical analysis of financial markets. However, the univariate technical approach could fail for several reasons, as explained in previous chapters. The market could be efficient, driven only by outside indicators, or the available time series could be too short for a significant technical analysis with the chosen forecasting horizon. This modelling method integrates the multivariate fundamental forces with artificial neural networks. It allows one to forecast difficult financial markets under difficult circumstances.

A useful and practical classification of the different economic models includes the following three-dimensional pairs:

- quantitative - qualitative
- univariate - multivariate
- linear - non-linear

The three dimensions build a cube with eight sub-cubes. Now it is possible to classify the various well known traditional economic models of the financial markets into the sub-cubes. The chart analysis of the practitioners of the financial field is a qualitative, univariate, non-linear model, and the similar technical analysis is a quantitative, univariate, linear model. The academic econometric models are quantitative, multivariate, linear models, etc.. The presently examined methodology integrates the whole quantitative part with univariate and multivariate, linear and non-linear parts.

The developed methodology has four steps for constructing forecasting models of financial markets:

- Data pre-processing
- Feasibility analysis
- Model development
- Performance evaluation

Each of these steps will be discussed in the next several paragraphs of this chapter. The construction process is not a single-pass procedure but a repetetive one. At any time during the process, it may be necessary to re-examine the information from a prior step or even to collect new or different data [Klimasauskas94, 14].

The core of the presented work is a methodology for the identification of the relevant market indicators and the appropriate model type for modelling the relations between the indicator time series and the predictor time series, which is desired for accurate forecasting. The method is called feasibility analysis and consists in identifying the relevant indicators with linear and non-linear, univariate and multivariate analysis. The different analyses evaluate "a priori" chance of forecasting the defined system and help to define the different aspects of the forecasting model. The relations between the predictor and possible indicators are identified with traditional statistics and an artificial neural network sensitivity analysis with the bootstrap methodology. The fundamentals of the single elements of the feasibility analysis were described in the previous chapters on statistical tools and artificial neural networks. The results of the feasibility analysis allow for the design of an artificial neural network or a simple linear regression model for the forecast of the predictor.

This chapter begins by describing the data pre-processing steps. The second section examines feasibility analysis. The following sections contain the design steps of a regression model and an artificial neural network. The last section describs the possibilities for performance evaluation. The issue of performance evaluation is theoretically simple, but for practitioners very important as it is distinct from the academic performance evaluation for model design. The goal of this chapter is to clarify the applicability of the theories from previous chapters for the following chapter on case studies. This chapter is perhaps the most valuable for the practitioners from the financial field.

7.1 Data preprocessing

Data is the foundation of quantitative models. This explains the importance of careful data pre-processing before beginning with the steps for the analysing and modelling. Three issues are addressed in data pre-processing:

- data collection
- handling missing or incorrect data
- transformation or scaling

Implicit input for the system and model characterisation is derived from a fundamental knowledge of the economy. This is helpful for a first selection of indicators and models, where the neural network is only one of several different possibilities.

The first step is data collection. Market data on stocks, bonds, foreign exchange, commodities, futures, and options are readily available from various sources. Access is relatively easy, but the price, the available history and the quality of the data varies widely. Data collection involves the acquisition of high-quality data from reputable and in the future available sources, that provide complete and consistent information for the desired data fields. The use of fundamental economic data, like unemployment rates, is much more difficult than the use of market date, because fundamental economic data are frequently revised. This raises the issue of which series to use. Should one use the original series that influenced the market, or the revised series available months or quarters later? An additional complication is that when the revised data are made available, the original data are overwritten. It is essential to maintain a pragmatic engineering perspective, meaning, a balance between research and production requirements. Market data fit that requirement much better than do fundamental economic data [Deboeck94, 30]. The influence of fundamental macro-economic data seems not to be very high in comparison to market date what allows the restriction to market data in the case studies [Williams98].

The selection of the appropriate data frequency depends on the forecasting horizon. Market data are easily available at a daily or lower frequency. The collection of high frequency intraday data needs a specific and expensive infrastructure. The advantage of high frequency data is a larger number of datapoints; however, the time period covered is normally quite short from an economical point of view. That makes it difficult to span most types of market states like bull, bear, and congested markets. Daily and especially monthly data have longer histories, and span better the different states of economic development. However, there are fewer datapoints and thus the statistical significance of the results is more doubtful [HFDF95, 6ff.].

If there are missing data or gaps in the time series, it is desirable to do the following:

- Try to find the missing data potentially from other sources.
- Interpolate or use average values.
- Drop the record in question.

The best solution is if one can return to another data source and find the missing data. But often the missing data is not available because certain markets may be closed at the time. The situation of closed markets makes it impossible to average missing values in production because future values are not known. Another approach is setting the missing value to 0 or to the last available value. In this case the missing values do not influence the weight change of an artificial neural network, because the weight update formula in back propagation is [Klimasauskas94, 14]:

$$\Delta w_{ij} = -k * \sum_{t=1}^{T} F\_net_i(t) * x_j(t)$$

The last possibility is to eliminate patterns with the missing data from the sliding time window.

Raw data are seldom adequate for analysing and modelling financial markets. Transformations are necessary to enhance information that provides a better descriptor of the processes present in the raw data. The objective of data transformations is to pull the data apart and thus simplify classifications by artificial neural networks or other forecasting models. Neural networks are not powerful enough to find a "needle in a haystack", as was believed at the beginning of neural network research. Neural networks are math, not magic. Data transformation should be made to help the neural networks with the process of prediction [Deboeck94, 31f.].

Several other issues emerge with regard to data transformations when different types of data are used by multivariate modelling. Plotting the observations of the predictor and indicator time series is the first step in any time series analysis and offers a preliminary idea of the behaviour of the time series.

In general, there are four major classifications of data: nominal, ordinal, interval, and ratio. Nominal data are typically symbolic, with no order relation. Order data permit order relations between data, but not exact differences between data elements. Interval data permit ordering and difference relations. Ratio data allows all comparing functions between data elements. The market time series usually consist of the same type, interval or ratio, which makes integration easier, rather than integrating different data types into a single model [Deboeck94, 33].

Unlike most technical forecasting systems, neural networks usually cannot deal with the wide range values found in raw data. The time series should be scaled in such a manner so that they are consistent with the activation function used in the neurone. This rule may be overridden by the issue of outliners described later. Since the neural network has a sigmoid output node and the indicators are of different orders and magnitude, there is a need for scaling of the predictor and the indicators. The predictor must lie between 0 and 1, because the sigmoid function produces a value in the range [0...1], but it is difficult to reach the extremes of the range. For this reason and that of being aware of exceeding the limit by future values, the predictor is scaled among 0.2 and 0.8.

It is important to scale the data such that no particular variable will influence the error estimates based on its magnitude. Thus, if the variables range comparably, errors will more likely be a function of misfit rather than of magnitude discrepancies among variables [Deboeck94, 35]. The indicator time series are scaled in the range [-1...1]. The scaling formula for the predictor and the indicators is [Klimasauskas94, 13f.]:

$$y = \frac{x(T_{\max} - T_{\min})}{x_{\max} - x_{\min}} + T_{\min} - \frac{x_{\min}(T_{\max} - T_{\min})}{x_{\max} - x_{\min}}.$$

where
y is the rescaled value.
x is the raw or unscaled data value.
T_{min} = target minimum (0.2 for predictor, -1 for indicators)
T_{max} = target maximum (0.8 for predictor, 1 for indicators)
x_{min} = raw minimum
x_{max} = raw maximum.

Outliners can substantially affect the ability of a neural network to use the input data. One of the characteristics of the numeric algorithms used in backpropagation is that they tend to smooth out noise. This is good in that it allows for effective modelling of noisy systems. However, if all of the data is concentrated in a very small portion of the input range, that input to the network may have little effect on the resulting model [Klimasauskas94, 14f.]. The distribution of values in a time series may be skewed towards either low or high values. Thus it may be worth considering a transformation that reduces the asymmetry. The square root transformation is effective in this respect [Azoff94, 23]. It is also possible to drop the outliners or to use log transformations. For example, the movement of the stock market on October 19,1987, may not be relevant to the model. In general, when the tails of a distribution of the data are very wide then log or square root transformations should be considered. These transformations compact the large values into smaller values so that more logical distances are created for the neural networks to learn [Deboeck94, 36]. But there exist a trade-off between an optimal transformation and a possible loss of information, because every transformation can destruct the fragile structure inherent in the original time series [Azoff94, 27].

The use of differences detrends the predictor. Another possibility is the adjusting of stock returns with the inflation rate [Peters91, 165].

An additional pre-processing step is the aggregation or further transformation of input indicators. Examples of statistical transformations are regression slopes that measure the direction and acceleration of trends in the data. More complex transformations include Fourier transformations and wavelet transforms. Other transformations for financial applications are derived from technical analysis. Examples are, moving averages to identify trends, volatility measures to identify trading markets, and oscillators to identify over-sold or over-bought conditions [Deboeck94, 32]. The difference between complex data transformations and models is not very clear. In this work only scaling transformations and the calculation of spreads are used. Spreads are differences between the values of similar markets. A common spread is the difference between the long and the short interest rate. This spread measures the expectation of the market participants for the future interest rates.

7.2 Feasibility analysis

The feasibility analysis contains the identification process of the significant indicators and the relevant time lags of the market indicator time series. Traditional univariate and multivariate statistics and an artificial neural network test with a sensitivity analysis are used together in order to identify linear and non-linear relations between predictor and indicators. The multivariate non-linear part with the artificial neural network sensitivity analysis with bootstrap is original [Ankenbrand96a, 27ff.]. Much of the application development with neural networks has been done in an ad-hoc basis without due consideration for model identification. An exception is a methodology including a model selection procedure, a variable selection procedure based on variable significance testing and a model adequacy procedure based on residuals analysis [Refenes97]. The methodology is similar to ours and independently developped.

Multivariate analysis is important for a number of reasons. The assumption of closed system dynamics is probably too restrictive for economic systems. The following multivariate tests do not depend on an assumption about the system dynamics or market efficiency. However, if the financial markets have a deterministic explanation, the feasibility analysis does not increase the complexity of the forecasting model. Due to the significance of the univariate tests and the failure of the multivariate tests, one can design an univariate forecasting model. An economic feedback loop gives information about the relevance of the results of the feasibility analysis, and helps to explain the results of the final forecasting model.

Feasibility analysis is an analytic technique for the selection of variables for the forecasting model. Other approaches directly use the neural network to decide which inputs are appropriate. In addition, a sensitivity analysis can be used to determine which inputs are important, to rank each input based on importance, and to eliminate the bottom of the list [Klimasauskas94, 18f.].

The four steps of the feasibility analysis and the used techniques are:

- univariate linear analysis: autocorrelation function
- univariate non-linear analysis: cumulated periodogram and BDS test

- multivariate linear analysis: crosscorrelation function
- multivariate non-linear analysis: original neural network sensitivity analysis with bootstrap

The results of the different analysis determine the selection of the model type and its topology. Regression models allow for the modelling of linear correlation, which are simpler than neural networks. Non-linear dependencies are modelled with neural networks. The feasibility analysis also determines the inputs for the different models.

7.3 Linear regression model design

7.3.1 Univariate model

A useful class of models for univariate linear time series is formed by combining a moving average and autoregressive processes. A mixed autoregressive/moving average process is an ARMA process. The importance of ARMA processes lies in the fact that a stationary time series may often be described by an ARMA model involving fewer parameters compared to a pure moving average or autoregressive process used by itself. But in practice most time series are non-stationary. In order to fit a stationary model it is necessary to remove non-stationary sources of variation. If the observed time series is non-stationary in the mean then we can difference the series. Such a model is called an integrated model because the stationary model which is fitted to the differenced data must be summed or integrated to provide a model for the non-stationary data. A general autoregressive integrated moving average process, abbreviated as an ARIMA process, is widely used [Chatfield89, 40ff.].

A univariate linear forecasting procedure based on ARIMA models is known as the Box-Jenkins approach. The main steps in setting up a Box-Jenkins forecasting model are as follows:

- model identification
- estimation
- diagnostic checking

The first stage of model identification examines the data to see which member of the class of ARIMA processes appears to be most appropriate. The second step estimates the parameters of the chosen model. The last step examines the residuals from the fitted model to see if the model is adequate. If the model appears to be inadequate for any reason, then other ARIMA models may be tried until a satisfactory model is found [Chatfield89, 71f.].

The first step for the estimation of the ARIMA model parameters is to difference the observed time series until they are stationary. Now an ARMA model is appropriate for the differenced time series. The problems of estimation for an ARMA model are similar of those for a moving average model in that an iterative procedure must be used. The residual sum of squares can be calculated at every point on a suitable grid of the parameter values, and the values which give the minimum sum of squares may then be assessed. Nowadays exact maximum likelihood estimates are often preferred, despite the extra computation involved [Chatfield89, 59f.].

7.3.2 Multivariate model

Multivariate linear procedures are multiple linear regression models, where the variable of interest y is linearly related to one or more other variables $x_1, \ldots, x_p$ which are called explanatory variables. Lagged values of the response and explanatory variables can be included. Economists sometimes describe a regression model as an econometric model [Chatfield89, 76ff.]. A detailed description of multiple regression can be found in many statistics textbooks and will not be repeated here.

The explanatory variables are determined by the crosscorrelation function of the feasibility analysis.

7.4 Neural network design

The accuracy of a neural network design is measured by two metrics:

- convergence - accuracy of model fitness in-sample
- generalisation - accuracy of model fitness out-of-sample

Convergence is concerned with the problem of whether the learning procedure is capable of learning the classification defined in any data set, under what conditions it does so, and what the computational requirements are for convergence. Fixed topology networks prove convergence by showing that in the limit, as training time tends toward infinity, the error minimised by the gradient descent method will tend toward zero. Generalisation is the main property that should be sought. It measures the ability of a network to recognise patterns outside the training set. A network having a structure which is simpler than necessary cannot give good approximations even of patterns in the training set. A structure more complicated than necessary overfits, in that it leads to good fit for the training set but performs poorly with unseen patterns [Refenes95, 16f.].

The following control mechanisms influence the above performance measures:

- the choice of cost function
- network architecture
- learning algorithms
- learning times

The choice of the cost function is believed to play an important role in determining the convergence and generalisation characteristics of supervised learning procedures. The most commonly used cost function is the family of quadratics, for example the least mean squared error. Several researchers suggest that instead of minimising a cost function between observed and actual values and subsequently using its output as an input to the ultimate objective function, which typically in economics tries to maximise profits according to a specific investment strategy. It may be more useful to consider both processes in a single step and attempt to maximise the ultimate objective function in the beginning. Since the mechanics of the implementation of this approach depend on the specific application with its own objective function, this work will not deal with this approach any further and instead will use the common cost function described in the artificial neural network chapter [Refenes95, 21].

The architecture or design of a neural network system is defined by the arrangement of its units, that is, the set of all weighted connections between neurones. The used single hidden layer feedforward network is a layered or hierarchical network whose neurones are hierarchically organised into disjointed layers. The neurones only affect neurones at a higher layer. The lowest layer of units is the input layer, the intermediate layer is the hidden layer, and the highest layer is the output layer [Refenes95, 27].

The topology of the network is important in determining the cognitive characteristics of the network and particularly its ability concerning generalisation. Generalisation is the main property that should be sought from a neural network system. It determines the amount of data needed by the system so that a correct response is produced when the system is presented with patterns outside the training set. The conditions in which good generalisation performance can be obtained are yet not well understood. A common sense rule is to minimise the number of free parameters in the network. However, this must be done without reducing the size of the network to the point where the desired function can no longer be computed and the convergence of the neural network is reduced [Refenes95, 27].

Since reducing the size of the network will also reduce its generality, some knowledge about the task domain is necessary in order to insure the preservation of the network's ability to solve the problem. Such a priory knowledge is readily available from the feasibility analysis. The results of the feasibility analysis determine the type of input.

There are hardly any systematical investigations for the optimal design. The degrees of freedom of the neural network are determined by the number of connections. An important rule of thumb to remember is the relation of degrees of freedom of the neural network to the number of data sets [Refenes95, 25ff.]:

$$c < \frac{n}{10}$$

where
c is the number of connections.
n is the number of data sets.

More degrees of freedom allow the artificial neural network to adapt to noise. Overlearning decreases the capabilities of generalisation. The neural network does not recognise the general structure of the market, it adapts only to the past.

A third important mechanism for controlling the performance of neural learning is the choice of the learning procedure and the choice of the control parameters for the gradient descent onto the error surface. An important term involved in the control of the gradient descent is the learning rate. It is also used as an additional momentum term. The learning rate and the momentum term influence the magnitude of weight changes and, therefore, are crucial for learning performance. However, it is difficult to find appropriate learning rates and momentums: because rates and terms of small magnitude imply small changes of the weights even when greater weight changes are necessary. There are two reasons for requiring greater weight changes: The first is speed of conversion and the second is network stability. High learning rate values help the network to escape from a local minimum. On the other hand, rates and terms of greater magnitude imply greater weight changes which can lead to oscillation [Refenes95, 26].

The problem of finding appropriate learning rates and momentum terms can be viewed as a specific aspect of the stability-plasticity dilemma which is expressed by the problem of how a learning system can be designed to remain flexible in response to significant new events, yet also remain stable in response to irrelevant events. Using modifiable learning rates and momentum terms leads to the distinction between two types of adaptation rule [Refenes95, 26f.]:

- weight adaptation rule
- learning rate and momentum adaptation rule

The weight adaptation rule is the well known backpropagation. The learning rate and momentum term adaptation rule is simple. If the error of the cost function grows, the learning rate and the momentum decrease [Miller94, 129ff.].

The backpropagation learning procedure requires unequal initial weights. The initial weight matrix defines the starting point on the weight-error surface. Usually, weights are initialised to random values. Typically several training runs with random initial weights are required to test the stability [Refenes95, 29f.].

The number of presentations of the data set to the network during training is referred to as the training time. A related term is the epoch, which stands for the number of training cycles after which an update of the connection weights is performed. It is typically fixed to 1. The training time and the network size are important factors in controlling network overfitting and overlearning. During the first few passes over the data set, the backpropagation algorithm extracts the main features of the data set. As training progresses the procedure will start overfitting the presented data, provided that the network has a sufficiently large number of free parameters. The common method of controlling overfitting with the use of training time is through cross-validation and the premature termination of training. The available data are divided into a learning or training set and a validation set. The training set is used in the normal way, while the validation set is used to test the validation set performance of the estimator. Training is terminated once the validation errors begin to rise while the learning rate and the momentum are equal to zero [Refenes95, 29].

Typically ten percent of the available data are in the validation set [Groot92, 21]. The selection of an input pattern for the validation set can influence the number of learning cycles. A different distribution of the learning and validation set stops the learning too early. This problem can be controlled with the bootstrap of the feasibility analysis. Several runs with different random splits for the training and validation sets are carried out to ensure statistical stability [Refenes95, 29].

The final implicit knowledge of an artificial neural network is a stable quantitative fundamental non-linear multivariate model.

7.5 Performance evaluation

When using neural networks or any mechanical trading system, it is essential to define how the performance will be measured. Criteria typically used in neural network development, such as root mean square error of the cost function, may have little validity in terms of the profit objectives of the system. For example, a neural network was developed that correctly predicted the next-day value of the US Treasury bond 85% of the time, but it was not profitable for trade. The system correctly predicted the market direction when the moves were small, but it was wrong on almost every large move. The fact that most markets did not have a Gaussain distribution, but had long tails, was responsible for the discrepancy between accuracy measured by the root mean square error and that measured by the profit. Moreover, transaction costs made the small moves unprofitable. The neural network consistently lost money [Klimasauskas94, 8ff.].

The performance of a trading model must also be evaluated in light of potential risks. Accordingly, the profit must be substantial enough to justify the unlikly loss of all investments plus additional trading capital. The risk must be well defined, thoroughly understood, and acceptable in light of the profit. Stringent risk management principles must be applied in addition to the forecasting model in order to further defend against catastrophic loss [Pardo92, 121].

A formal performance indicator for the forecasting model is the normalised mean square error. For practitioners in the financial field the goal of any testing strategy is to measure the profitability and the risks of the system. It is therefore important to design a testing strategy and metrics from the outset. There are many well-known metrics for measuring estimator performance [Refenes95, 67ff.; Pardo92, 122ff.]. In this work only the following simple metrics are used:

- residual analysis
- trend direction
- profit
- model efficiency
- maximum drawdown
- reward-to-risk ratio

The reason for the simple non-parametric performance indicators is due to the difficulties of measuring the performance of non-linear systems with parametric indicators, which makes the results doubtful. For measuring system stability and consistency; the performance of the system should be measured with an out-of-sample data set. The in-sample data is used for the development of the system and consists of the training and validation sets. The out-of-sample time window is kept for the entire analysing and modelling process, and is called the test set. It is also sometimes called the walk-forward analysis. The walk-forward analysis judges the performance of a trading or forecasting system strictly on the basis of postoptimization or out-of-sample trading. It is the closest simulation possible of the way in which a developed system is used in real time [Pardo92, 20]. The relation of the in-sample and out-of-sample performance results show the stability, consistency, and generalisation capability of the developed model. From this distinction of in-sample and out-of-sample data follows a second reason for the use of simple performance indicators. As economic problems are often undersampled, the statistical significance of complex indicators is doubtful. Next, the different performance indicators will be described in detail.

When a model has been fitted to a time series, it is advisable to check that the model really does provide an adequate description of the data. As with most statistical models, this is usually done by looking at the residuals, which are defined by:

residual = observation - fitted value

The tools for the residual analysis are the same tests as in the univariate feasibility analysis, for example the autocorrelation function and the cumulated periodogram. The residuals should be white noise. The residual analysis gives information about the accuracy and power of the developed model. If there are strong relations left in the data, then there are systematic errors which can be analysed and modelled in a new iterative cycle of the forecasting methodology [Chatfield89, 61f.].

Although predicting the levels of price change is desirable, in finance, foreseeing the sign of change is equally important. The following metric measures the success of trend prediction:

$$d = \frac{1}{n}\sum_{i}^{n} a_i$$

where

$$a_i = \begin{cases} 1 & if\,(x_{t+1} - x_t)(y_{t+1} - x_t) > 0 \\ 0 & otherwise \end{cases}$$

x is the time series
y is the predicted time series
n is the number of test data sets.

The d statistic has a simple interpretation: $d = 1$ implies that the estimator is predicting 100% of the directional changes; $d = 0$ implies 0% prediction in directional changes: therefore, any estimator with $d > 0.5$ offers a better result than tossing a coin. The statistic should be used with care, because it is relatively easy to obtain a higher value for d in a trending market [Refenes95, 69].

The profit calculation also takes into consideration the magnitude of the change. Profitability is always calculated in the context of a trading rule. The estimate of expected returns must be turned into investment actions. Consider a simple strategy in which positive expected returns are executed as long positions, and negative expected returns are executed as short positions. The net return of such strategy is given by:

$$r = \sum_{i}^{n} p_t (x_{t+1} - x_t)$$

where

$$p_t = \begin{cases} 1 & if(y_{t+1} - x_t) > 0 \\ -1 & if(y_{t+1} - x_t) < 0 \\ 0 & if(y_{t+1} - x_t) = 0 \end{cases}$$

assuming away transaction costs and slippage. Slippage is the difference between the desired trading price of the model and the real trading price in the real financial markets [Refenes95, 70].

The profitability of the system can be evaluated against two benchmarks: a simple buy-and-hold strategy and its distance from the ideal net profit. The buy-and-hold test is a good benchmark in traditional bond and stock markets where it is difficult to take short positions. In derivative or forex markets it is not very useful, because it is possible to take short positions. A short position means to sell an assets that will be buy in the future. Speculators like short position because it is possible to profit from decreasing prices. They sell assets without owning and buy it late for a cheaper price.

More useful is the second benchmark, the distance from the ideal net profit or model efficiency. It measures the returns of the trading system against that of a perfect predictor *d*:

$$r_d = \frac{\sum_{i}^{n} p_t (x_{t+1} - x_t)}{\sum_{i}^{n} \left|(x_{t+1} - x_t\right|}$$

with the same assumptions and definitions as in the previous formula.

The success of net profit and its derivatives has been overrated for several reasons. First, there is the possibility that a few large trades will skew the results. It is unrealistic to use a system whose success depends on the probably of non-recurring events. Second, it gives no indication of the risks involved [Refenes95, 71].

The characteristics of the cumulated profit or equity curve are important in assessing the acceptability of any trading model. The key characteristic is the maximum drawdown. Net profit and maximum drawdown are the ultimate expression of risk and reward. Maximum drawdown is expressed as the greatest retracement figured along the axis of cumulative returns. It is probably the most important in daily operations. It requires a lot of loyalty and deep pockets to trade in a system with a big drawdown. A smooth equity curve is much more desirable and harder to obtain than a high net return. Most often, the systems that give the largest net profits have the largest drawdowns [Refenes95, 72f.].

The reward-to-risk ratio compares maximum reward, represented by annualised profit, with maximum risk, represented by maximum drawdown. As a rule, the bigger the reward-to-risk ratio the better. It must be noted however, that the validity of this number is in direct proportion to the statistical validity of its components, net profit and maximum drawdown [Pardo92, 124f.].

An improvement in the results is possible with an antithetical combination of different models, especially when the predictors are not correlated [Groot93, 46f.]. The single models have threshold values if they at least have nominal output. The different models could also be weighted differently. The evaluation method for optimising could be the net profit and maximum drawdown. Also, additional improvements are possible with additional rules like [Pardo92, 40]:

- risk management
- profit targeting
- profit management
- money management
- position management

The search method is a grid search. An enhancement with evolutionary algorithms can be very promising [Pardo92, 63ff.].

A feedback loop from mathematics to the economy is important. The results can be explained with economic knowledge. The economy must be a filter for the model and the results. A system without any economical background is too dangerous, because it is a true black box.

8. Case Studies in Financial Markets Prediction

Both case studies are examples of middle term forecasting models of one month, but the methodology is also applicable to daily short term forecasting for trading models. The application of artificial neural networks is not restricted the stock or bond markets [Steurer97].

8.1 SPI

The first application is the monthly forecasting of the Swiss Performance Index (SPI). The main difficulties of this case study are the undersampling and the trend over the whole training period. Primary versions of this case study are published in [Ankenbrand95a, 257ff.; Ankenbrand95b, 365ff.].

The SPI is a Swiss stock index. It shows the development of a broad diversified portfolio. It contains all titles of the three Swiss bourses. The SPI is a member of the SwissIndex family, which also contains the Swiss Market Index (SMI), the Swiss Bond Index (SBI), and different subindicies. The calculations of all these indices are based on the Laspeyres formula with the weight of capitalisation [SwissIndex, 1ff.]:

$$I = k\sum_{i=1}^{M} p_i a_i$$

where
I is the value of the index.
k is the capital factor, which adjusts the index value to the desired base.
M is the number of index titles.
p is the price of the index title.
a is the number of the specific index title.

The data is available monthly from January 1987 to December 1996. An expansion of the time window is not possible, because the calculation of the SPI starts in 1987. A back calculation of the SPI with the aid of longer available stock indices, like the Vontobel stock index or the Swiss Bank Cooperation Swiss stock index, is not possible because the index formulas and the index titles are different. The expansion of the time window through the calculation of the SPI based on the single stocks is not possible, because the data base for this work does not consist of the prices and numbers of the single stocks. The problem of undersampling can not be solved in this work, and this must be taken into account in the forecasting methodology.

The selection of possible indicators for analysing depends on fundamental economic knowledge and the availability of time series. The indicator used here have in general an economic significance, although this need not to be strong from the economics point of view. However, they must be somewhat influent in practice. The following list shows the possible indicators for analysing and modelling:

- SPI
- 3 M CHF
- Bond CH

- Bond D
- Bond US
- Bond J
- DAX
- S&P 500
- Nikkey 225
- CHFDEM
- DEMUSD
- YENUSD

An influence of the past values of the predictor on the present value of the predictor is possible, and makes an analysis of the lagged values of the SPI as an indicator reasonable. The influence of the interest rates on the stock market is broadly accepted. The use of long and short interest rates gives a more precise picture of the Swiss interest market and its expectations. The 3 M CHF indicator is the 3 month Euroswiss yield, and the Bond CH indicator is the average yield of the Swiss Bond government bonds.

The financial markets of Switzerland are not independent from the rest of the world. For example, Switzerland has strong relationships with Germany. Possible indicators from the German financial markets are the DAX (German stock index) and Bond D (average yield of German government bonds). The two biggest economic areas in respect to Europe are represented by USA and Japan. The possible indicators are the S&P 500 (stock index of America), Nikkey 225 (stock index of Japan), Bond US (average yield of US government bonds), and Bond J (average yield of Japanese government bonds).

The connections and relations between the different geographic financial markets are realised through the foreign exchange rates: between Switzerland and Germany with CHFDEM, between Germany and America with DEMUSD, and between America and Japan with YENUSD.

The predictor and the indicators are available monthly from 1987 to 1996. The indicator values are the closing prices of the latest trading day of the month.

8.1.1 Data pre-processing

The time series are collected and verified from different sources. All indicators are market data and are available at the time of production without any time lag. The data are complete without any errors. The data are manually inspected which is easy because there are few data points.

The analysing and modelling steps are performed on data from 1987 to 1994. The data of 1995 and 1996 are held as a test set for the evaluation of the out-of-sample performance or the walk forward analysis.

The use of the monthly differences of the SPI as predictor time series detrends the predictor. The other indicators are also detrended by the first difference, but the interest rate and forex indicators are in addition investigated with their absolute values. The following list shows the minimum and maximum values and the scaling formula of the analysed indicators:

	min	max	scaling
3 M CHF	1.75	9.63	$y = 0.25x - 1.44$
Bond CH	3.77	6.98	$y = 0.62x - 3.35$
Bond D	5.70	9.10	$y = 0.59x - 4.35$
Bond US	5.97	9.60	$y = 0.55x - 4.29$
Bond J	3.24	7.92	$y = 0.43x - 2.38$
CHFDEM	80.95	92.85	$y = 0.17x - 14.6$
DEMUSD	1.39	1.98	$y = 3.39x - 5.72$
YENUSD	97.35	158.84	$y = 0.03x - 4.17$

The differences of the indicator time series are skewed towards either low or high values, which were affected by the stock market crash on October 19, 1987. The square root transformation is effective to reduce the asymmetry and the width of the tails of the distribution. Null differences remains at a value of zero, because a division through zero is not allowed.

	min	max	scaling
SPI	-288.90	141.47	$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{20x}$
3 M CHF	-1.13	1.50	$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{1.5x}$
Bond CH	-0.58	0.52	$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{0.6x}$
Bond D	-0.50	0.80	$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{0.8x}$
Bond US	-0.38	0.62	$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{0.65x}$
Bond J	-1.63	0.90	$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{1.65x}$
DAX	-322.79	212.94	$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{20x}$
S&P 500	-70.04	41.87	$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{10x}$
Nikkey 225	-5057.29	4210.60	$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{100x}$
CHFDEM	-2.67	2.37	$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{1.5x}$
DEMUSD	-0.14	0.17	$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{2x}$
YENUSD	-11.20	9.76	$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{3x}$

The predictor is the difference of the SPI, because the SPI has a strong upward trend like all stock market time series. The scaling formula for the predictor in the range [0.2...0.8] is:

$$y = \frac{\sqrt{|x|} * |x|}{40x} + 0.5$$

8.1.2 Feasibility analysis

Plotting the predictor against time is the first step in any time series analysis. A straight plot of the predictor time series and the first difference can show trends, cycles, seasonal, outliners and other effects on the series. The following figure shows a plot of the SPI and its first differences:

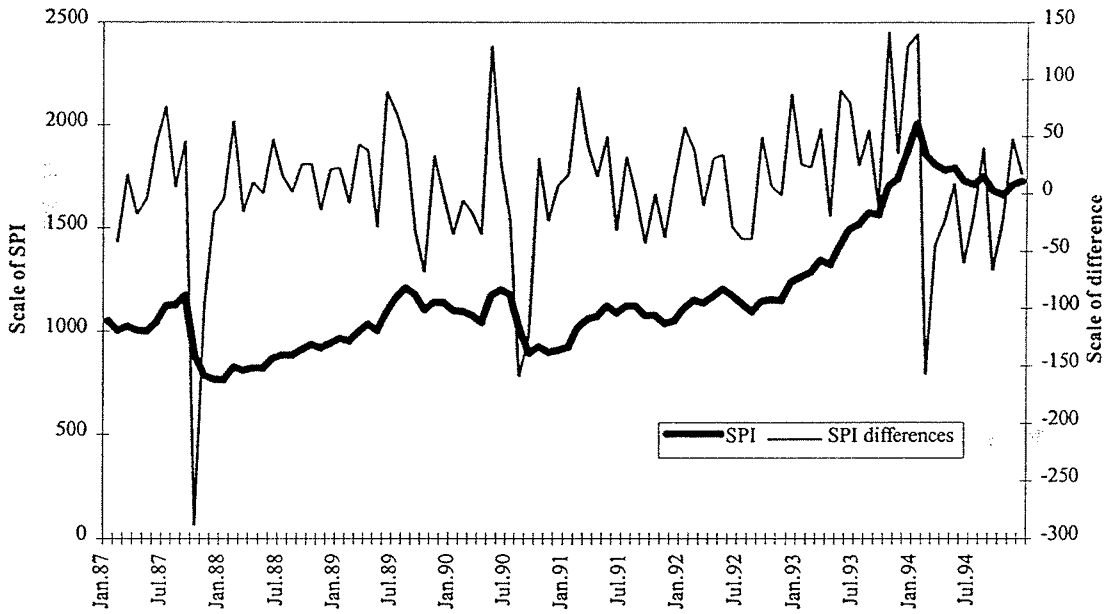

The figure clearly shows the upwards trend of the SPI. The available time window spans the types of market states like bull, bear, and congested markets. Tests on stationarity are not calculated, because the time window is too small to compute and compare the different moments of the sub-windows.

8.1.2.1 Linear univariate analysis

A useful measure of the linear association between lagged variables is the autocorrelation function. A detailed description of the autocorrelation function is found in chapter 5 of this work. The correlogram allows a fast interpretation of the results where the horizontal axis shows the time lags.

The monthly differences of the SPI do not show any significant linear autocorrelations at the 95% significance level, which lies at 0.2. Therefore, lagged SPI moves appear to be linearly independent.

8.1.2.2 Non-linear univariate analysis

A calculation of the BDS test (see chapter 5 for details) is difficult and not reliable with less than 200 datapoints [Brock91, 169f.], but the time series of the case studies contain only 95 datapoints. A less general test is the cumulated periodogram (see also chapter 5 for details), which may discover non-linear periodicity in the time series. Like the correlogram, the cumulated periodogram allows for a fast visual interpretation of the results like the correlogram.

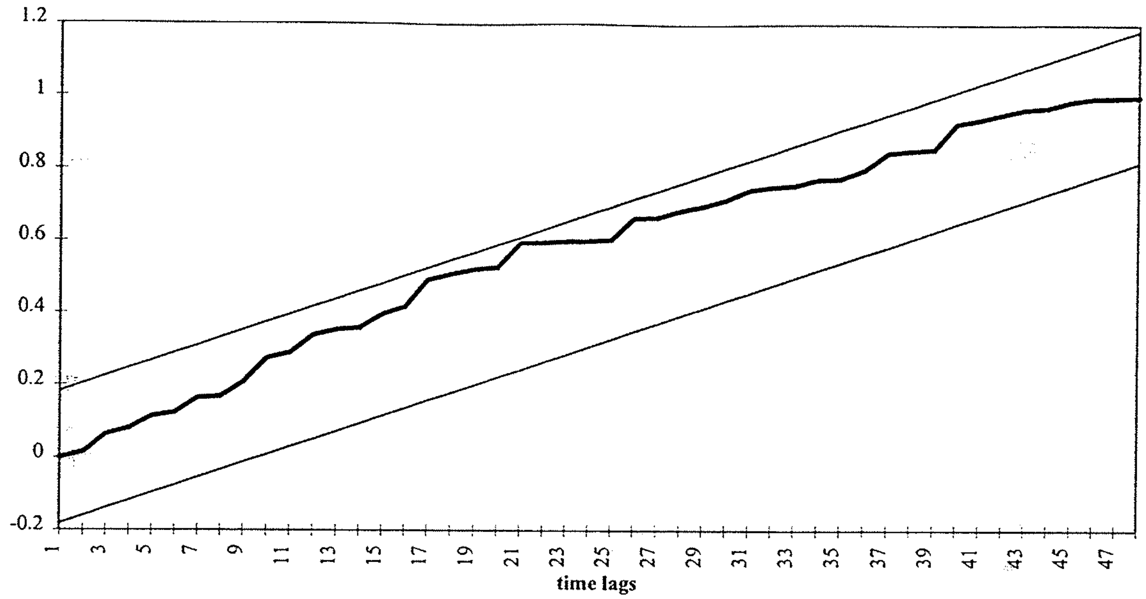

According to the cumulated periodogram test, SPI appears to be a random walk. The cumulated periodogram function line does not cross the 95% significance line.

8.1.2.3 Linear multivariate analysis

The cross correlation function measures the linear dependencies between the indicator time series and the predictor. The following list shows the cross-correlation's coefficients for the possible indicators at time lag 1:

3 M CHF	-0.03
Bond CH	-0.04
Bond D	-0.08
Bond US	-0.25
Bond J	-0.22
CHFDEM	0.17
DEMUSD	-0.04
YENUSD	-0.14

and the first differences of the possible indicators at time lag 1:

3 M CHF	-0.11
Bond CH	-0.26
Bond D	-0.18
Bond US	-0.25
Bond J	-0.24
DAX	0.22
S&P 500	0.18
Nikkey 225	0.07

CHFDEM	-0.11
DEMUSD	0.16
YENUSD	0.09

The 95% significance level is 0.2. Therefore showing that there are some significant but weak linear correlations.

8.1.2.4 Non-linear multivariate analysis

The non-linear dependencies are measured with the neural network sensitivity analysis with bootstrap according to the given description in previous chapters. The test is calculated with three steps. The first step measures the dependencies of the absolute indicator values. The second measures the dependencies of the interest rate differences, and the last step measures the influence of the other indicator differences. A separation is useful with respect to the degrees of freedom of the neural network and the number of data samples. The absolute values of the stock indices are not analysed because they are trendy and not stationarity.

The first neural network for the absolute indicators has eight input units, four hidden layer units and one output unit. The learning set contains 85 samples, and the validation set contains 10 samples. The expectations of the sensitivity indices *S* and its standard errors are:

	sensitivity index	standard error
3 M CHF	0.014	0.017
Bond CH	0.017	0.019
Bond D	0.015	0.017
Bond US	0.018	0.019
Bond J	0.019	0.021
CHFDEM	0.017	0.019
DEMUSD	0.016	0.015
YENUSD	0.018	0.018

The second neural network for the differences of the interest rate indicators has five input units, three hidden units, and one output unit. The learning set contains 84 samples, and the validation set contains 10 samples. The results of the sensitivity analysis are:

	sensitivity index	standard error
3 M CHF	0.016	0.013
Bond CH	0.014	0.013
Bond D	0.019	0.014
Bond US	0.017	0.014
Bond J	0.022	0.016

The third neural network has seven input units, four hidden units, and one output unit. The results are:

	sensitivity index	standard error
SPI	0.023	0.019

DAX	0.022	0.021
S&P 500	0.021	0.020
Nikkey 225	0.023	0.018
CHFDEM	0.029	0.023
DEMUSD	0.019	0.016
YENUSD	0.028	0.020

8.1.2.5 Conclusion and economic review

The results of the feasibility analysis are slightly different from previous published results [Ankenbrand95a; Ankenbrand95b], where the differences of Bond CH, the differences of S&P 500, and the DEMUSD are the most important indicators. However, all these indicators are significant in this analysis. The reasons for the differences are:

- expansion of analysed indicators
- different time window
- different indicator scaling
- improved non-linear feasibility analysis

The performance of the different forecasting models will be compared and discussed in the results section of this chapter.

The important influences are dynamic, only indicator differences strongly influence the development of the SPI. The dependencies between the bond market or interest rates, and the stock market are well known. Additionally, the connection between the Swiss and the German markets is well known in the financial field. Surprising from an economic viewpoint is the strong influence of the Japanese markets.

The details of the linear and non-linear model design are described in the following section.

8.1.3 Model design

The advantage of the multivariate linear regression model and the artificial neural network model is the combination of the influences from the indicators on the SPI. That allows for the development of a good model if the single influence of the single indicator is weak. Multi-factor models perform better than single factor models because it is not possible to identify one main explanatory variable. But the different indicators with their small influences allow for the satisfactory modelling of the SPI. The result of the feasibility analysis suggests the need to develop a linear regression and a non-linear neural network model.

8.1.3.1 Linear regression model design

Significant correlations show the following indicators assuming the multivariate linear analysis:

- Bond US
- Bond J
- Bond CH difference
- Bond US difference
- Bond J difference
- DAX difference

It is possible to exclude the absolute indicators Bond US and Bond J without loosing prediction power. The final regression model is:

$$x(t+1) - x(t)$$
$$= 0.0008(d(t) - d(t-1)) - 0.7(ich(t) - ich(t-1)) - 0.4(ius(t) - ius(t-1)) - 0.1(ij(t) - ij(t-1))$$

where
x is the SPI
ich is the Bond CH
ius is the Bond US
ij is the Bond J
d is the DAX.

The explanatory variables and the predictor are not scaled, because linear regression models do not have the same limitations as neural networks regarding the magnitude of input data.

The regression model is optimised until the end of 1992. The remaining of the data is the validation set. The optimised value is the risk-to-reward ratio which contains the net return and the maximum drawdown. The out-of-sample performance is discussed in the paragraph on results. The in-sample performance of the learning and validation set, which means the convergence of the model, is presented in this paragraph. The following chart shows the forecasted SPI difference, the realised SPI difference, and the net return curve of the learning set.

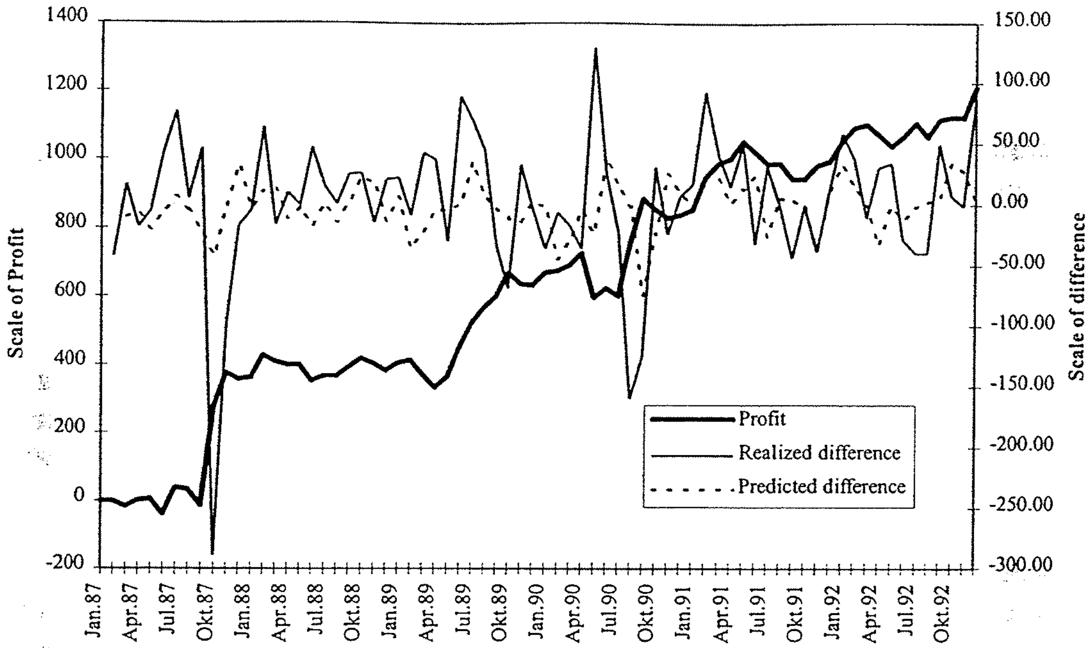

The same graph of the validation set performance is shown with more details:

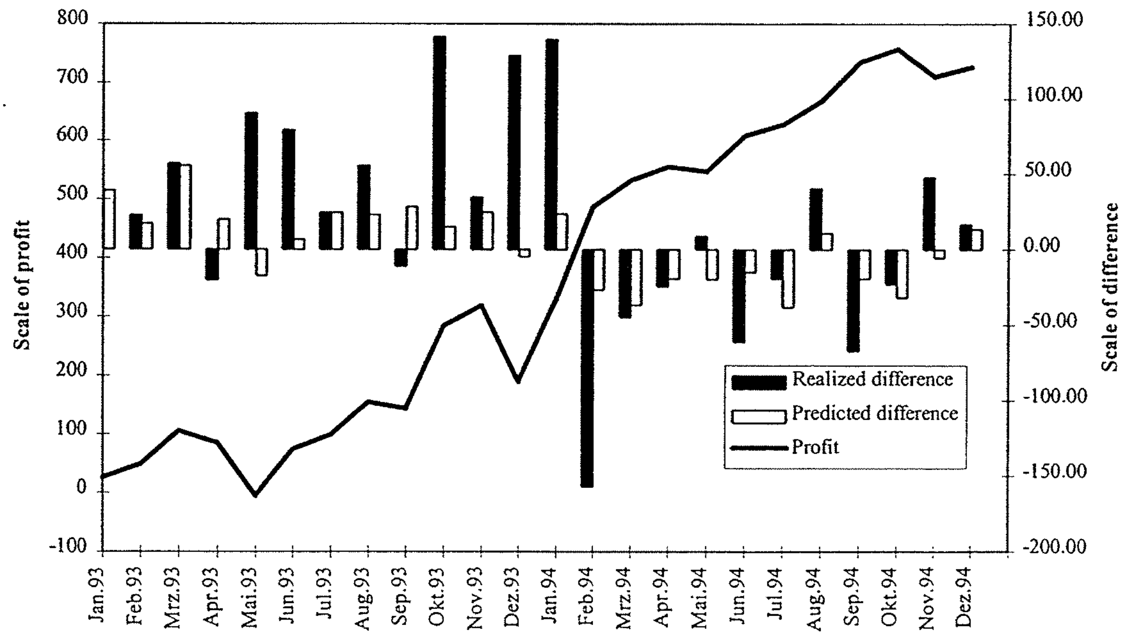

The residuals of the learning and validation set are white noise according to the autocorrelation function and the cumulated periodogram. The trend conformity in the learning set is 74 percent and in the validation set 75 percent. The profit, model efficiency, maximum draw-down and reward-to-risk ratio are shown in the following table:

	learning set	validation set
annualised profit	201	361
model efficiency	44%	54%

drawdown	130	129
reward-to-risk	1.55	2.81

The model appears to be stable, and the generalisation capabilities of the model should be good, as there are only four degree of freedoms, however, a final judgement is only possible after the discussion of the out-of-sample performance.

8.1.3.2 Neural network design

The SPI time series model only has 94 data points. An expansion of the time series is impossible because the calculation of the SPI starts in 1987. The maximum size for an artificial neural network is 3 input, 2 hidden and 1 output units with 10 connections or degrees of freedom. The small size of the neural network reduces the danger of overfitting, and improves the generalisation capabilities. This limitation makes it necessary to choose only the most significant indicators and not to use all the significant indicators. The most significant non-linear indicators regarding the feasibility analysis are the monthly differences in long interest rates of Germany and Japan and the forex exchange rate between America and Japan.

The artificial neural network model is:

$$x(t+1) - x(t) = f\left[d(t) - d(t-1), j(t) - j(t-1), y(t) - y(t-1)\right]$$

where
d is the Bond D
j is the Bond J
y is the YENUSD

The learning set contains the first 84 data samples, the validation set contains the remaining 10 data samples. The backpropagation algorithm needs 1'340 learning cycles with a decreasing learning rate and a momentum term. The following figure shows the development of the error term over the learning and the validation set.

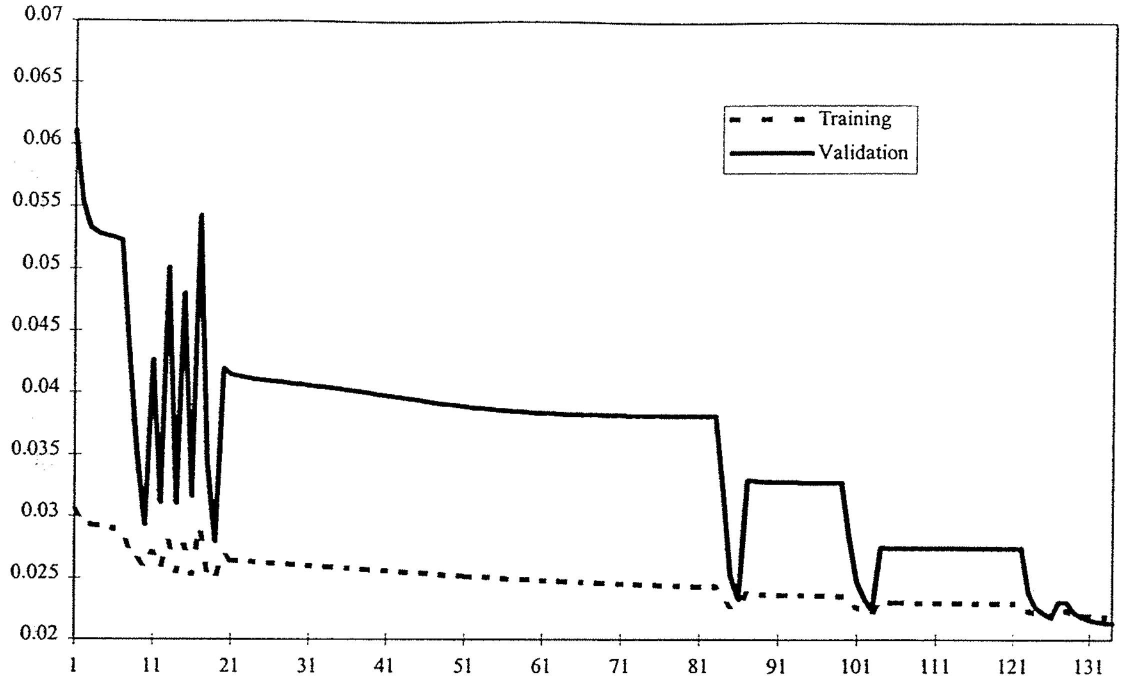

The next figure shows the development of the learning rate and the momentum term over the same period, and explains the jumps in the error curves in the figure above.

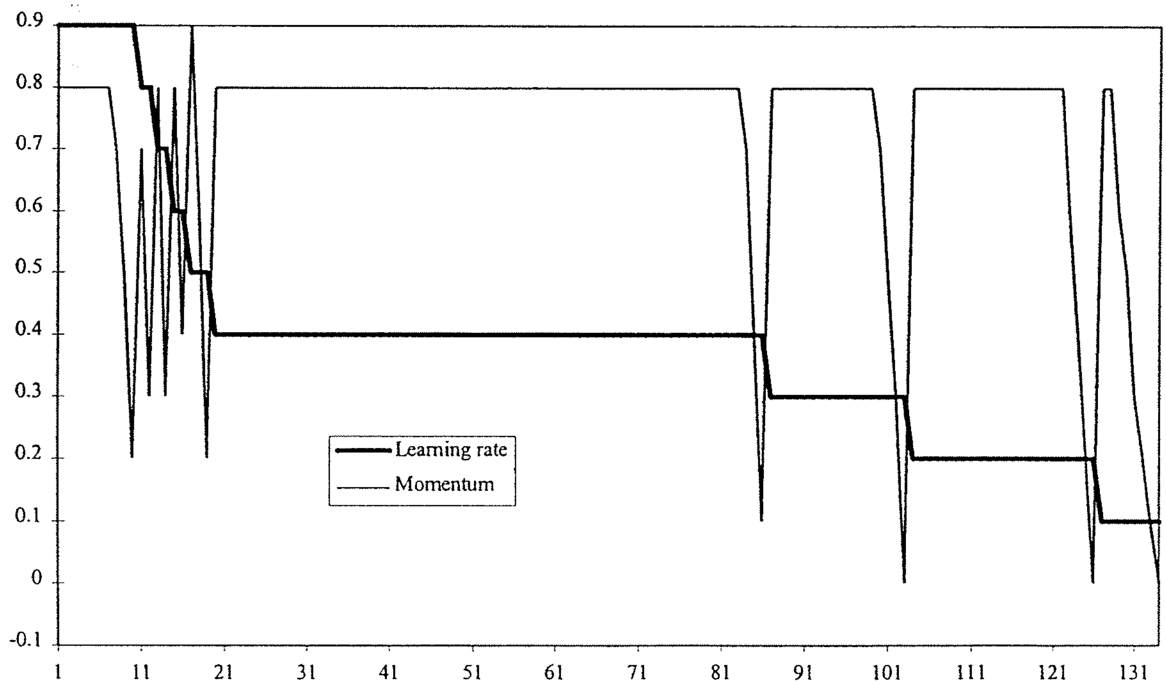

The in-sample performance and the convergence of the neural network are also presented in this section. The following chart shows the performance of the learning set.

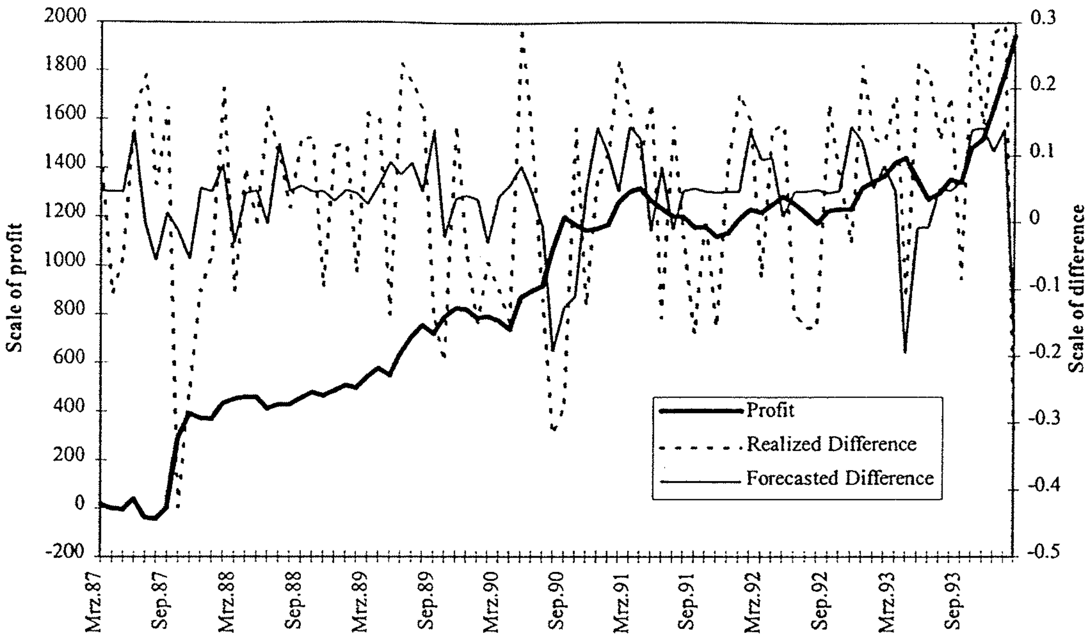

The next graph shows the performance of the validation set.

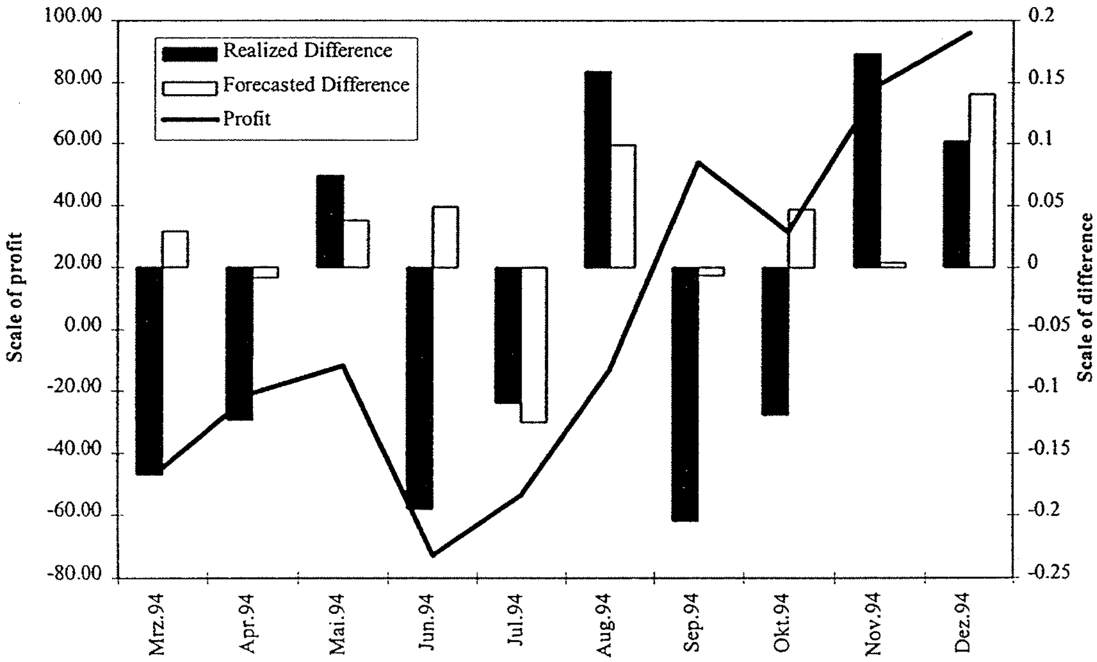

The residuals of the learning and validation set are white noise. The trend conformity of the learning set is 62 percent and for the validation set 70 percent. The other performance measurements are:

	learning set	validation set
annualised profit	315	115

model efficiency	52%	27%
drawdown	197	73
reward-to-risk	1.60	1.58

This model also appears to be stable. An interesting fact is that the trend conformity of the neural network is lower than that of the regression model, but the realised profit is higher, especially in the learning set. However, the results are not exactly comparable, because the training and the validation set of the regression model and neural network model are different. An exact comparison will be done in the next section with the test data sample.

8.1.4 Results

Firstly, the out-of-sample performance of the previous published model [Ankenbrand95a; Ankenbrand95b] will be presented.

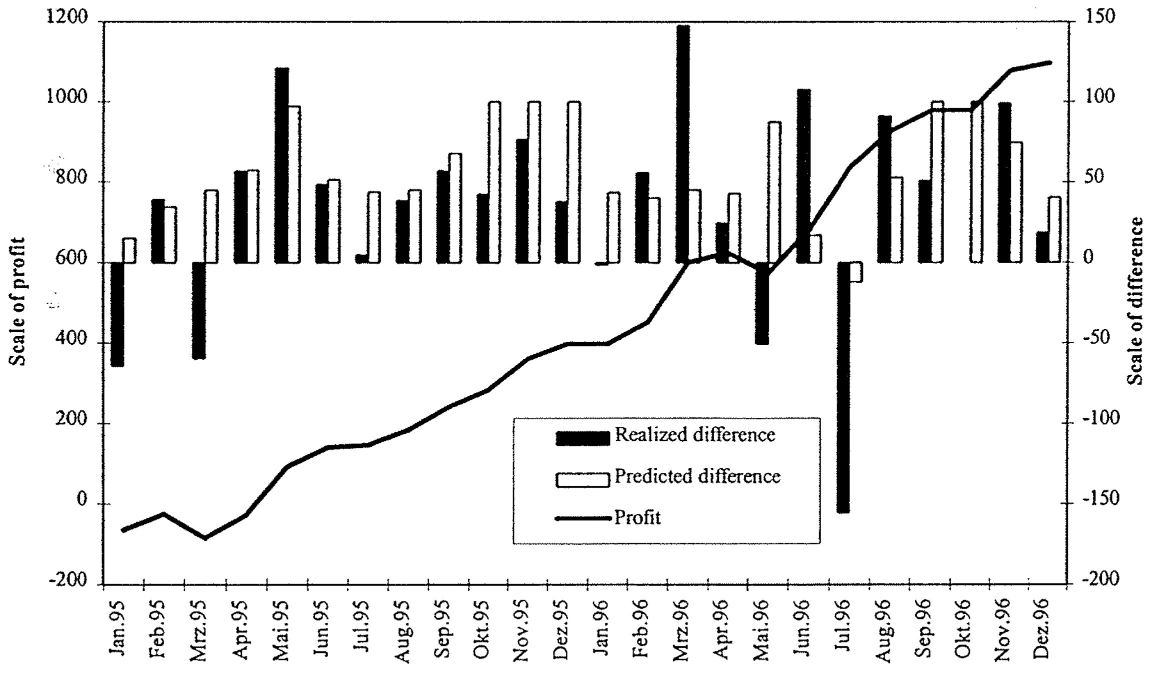

The next figure shows the out-of-sample performance of the linear regression model over the same time window.

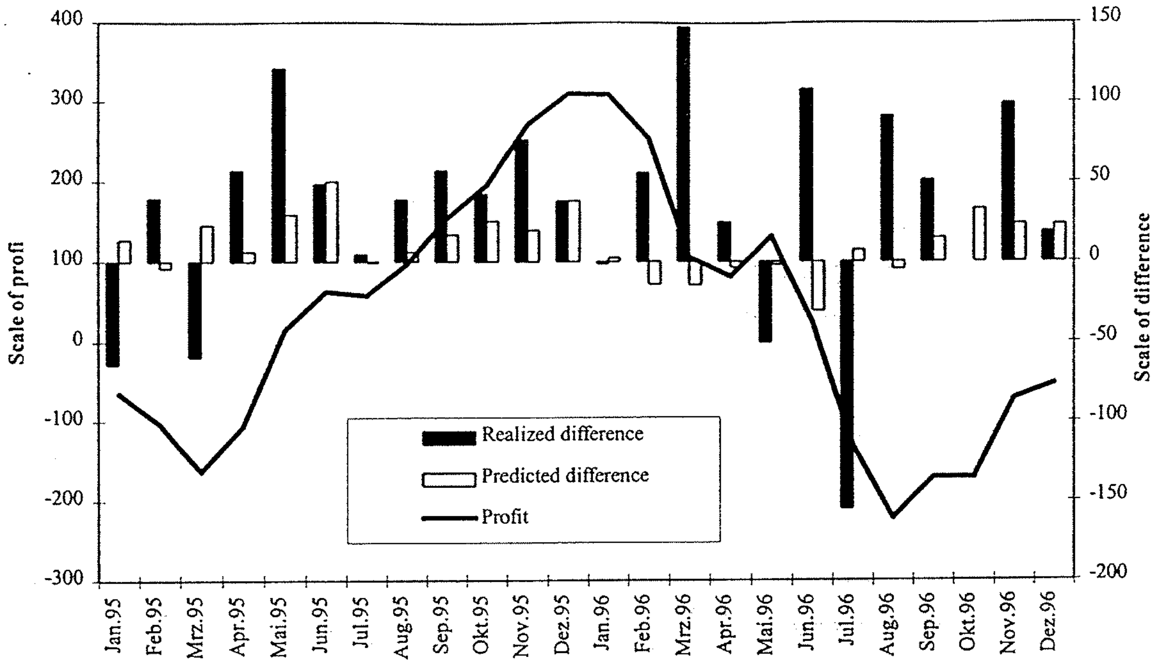

The neural network based on the enhanced feasibility analysis presented in this work has the following out-of-sample performance in the years 1995 and 1996.

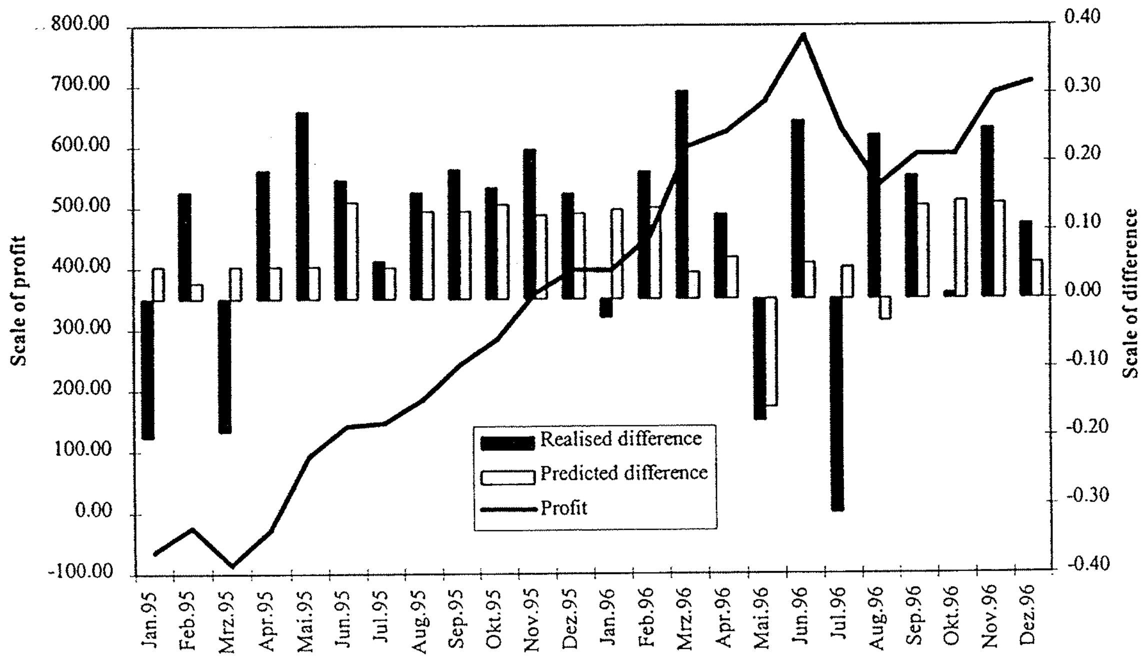

A residual analysis for the out-of-sample data sets are not performed due to the fact that the time series are too short for significant analysis. The following table compiles in a conclusive manner the results of the other performance indicators for the three models:

	previously published	linear regression	neural network
trend conformity	79%	54%	79%
annualised profit	548	-26	353
model efficiency	76%	-3%	49%
drawdown	60	532	246
reward-to-risk	9.13	-0.04	1.43

The previously published model has a very good performance record, which makes it difficult to judge objectively the stable and good performance of the neural network presented here. The two similar values of trend conformity are interesting, but the other performance measurements are different. The linear regression model performs badly, but over-optimisation is not common, because the model has few degrees of freedom, and the performance of the validation set is stable. Another possibility is a change in the linear market structure. It seems to be a market drift; the market has evolved and is more efficient in a linear manner.

A combination of the different models improves the results, but the danger of overfitting is very high, because the model is undersampled and the number of degrees of freedom increases with a combination of the models.

8.2 Bond market

The second case study involves the analysis of the monthly difference of the Swiss bond yield (Bond CH). Previous versions of this case study are published in [Ankenbrand96a, 27ff.; Ankenbrand96b]. There is a longer history of data available than in the first case study on the SPI, and there is not a clear trend in this data as interest rates have more cycles than stock markets.

The predictor is the average yield of a basket of Swiss governments bonds with a duration of between 5 and 10 years. This predictor is a widely used measurement for the long interest rates in the Swiss capital market. The Bond CH is calculated and published daily.

The selection of possible indicators for analysis depends on fundamental economic knowledge and the availability of time series. The following list shows the possible indicators for analysing and modelling:

- Bond CH
- Bond D
- Bond US
- Bond J
- 3 M CHF
- 3 M DEM
- 3 M USD
- DAX
- S&P 500
- Nikkey 225
- CHFDEM
- DEMUSD

The possible indicators and the economic reasons are basically the same as in the first case study. The additional integration of the Japanese related indicators, based on the results of the first case study, was that the Japanese financial market has shown a strong influence on the Swiss market. The data is available from 1982 to 1996 on a monthly basis (on the close of the last trading day of the month). The longer time window is the reason for not using the SPI. This is possible because the SPI and the DAX are strongly correlated, and the Swiss stock market has not shown any influence on the Swiss capital market in previous works [Ankenbrand96a, 27ff.; Ankenbrand96b].

The spreads or differences of different indicators are useful aggregated indicators and can improve forecasting performance. The following aggregated spread indicators are also tested based on the first cycle of the iterative feasibility analysis and previous work [Ankenbrand96a, 27ff.; Ankenbrand96b]:

- SCH
- SD
- SUS
- SCHD
- SUSD
- SUSJ

The first three spreads are the differences of the long and short interest rates of Switzerland (SCH), Germany (SD), and America (SUS). The spreads show the expectations of the market participants. If the interest rate structure is negative, the market expects higher interest rates. A negative structure signifies higher short term rates than long term rates. The other promising difference indicators are the spreads or differences in analogy to the forex exchange indicators:

- SCHD (Bond CH - Bond D)
- SUSD (Bond US - Bond CH)
- SUSJ (Bond US - Bond J)

8.2.1 Data pre-processing

The time series are collected and verified from different sources. All indicators are market data and available at production time without any time lag. The data are complete and without any errors. The data are manually inspected which is easily possible because there are few data points.

The analysing and modelling steps are performed on data from 1982 to 1994. The data of 1995 and 1996 are held as test set for the evaluation of the out-of-sample performance or the walk forward analysis.

The use of the monthly differences of the Bond CH as predictor time series is more adequate, because the main interest in financial markets lies more in the further development of the interest rate than in the absolute value. The other indicators are also detrended by the first difference, but the interest rate, forex indicators and the spreads are also investigated with their absolute values. The following list shows the minimum and maximum values and the scaling formula of the analysed indicators of the learning and validation set:

	min	max	scaling
Bond CH	3.77	6.98	$y = 0.62x - 3.35$
Bond D	5.70	9.80	$y = 0.5x - 4$
Bond US	5.97	14.22	$y = 0.25x - 2.5$
Bond J	3.24	8.72	$y = 0.4x - 2.5$
3 M CHF	1.75	9.63	$y = 0.25x - 1.44$
3 M DEM	3.19	10.13	$y = 0.3x - 2$
3 M USD	3.13	15.75	$y = 0.2x - 2$
CHFDEM	79.52	92.85	$y = 0.15x - 13$
DEMUSD	1.39	3.36	$y = x - 2.4$
SCH	-3.47	2.18	$y = 0.35x + 0.2$
SD	-1.71	3.31	$y = 0.4x - 0.3$
SUS	-2.05	4.25	$y = 0.3x - 0.3$
SCHD	-4.65	-1.12	$y = 0.5x - 1.3$
SUSD	-0.85	5.14	$y = 0.4x - 1$
SUSJ	1.03	6.14	$y = 0.4x - 1.4$

The differences of the indicator time series all have some outliners. The square root transformation is effective to reduce the width of the tails of the distribution.

	min	max	scaling
Bond CH	-0.58	0.52	$y = \frac{\sqrt{\|x\|} * \|x\|}{0.8x}$
Bond D	-0.50	0.80	$y = \frac{\sqrt{\|x\|} * \|x\|}{0.8x}$
Bond US	-0.80	0.62	$y = \frac{\sqrt{\|x\|} * \|x\|}{0.8x}$
Bond J	-1.63	0.90	$y = \frac{\sqrt{\|x\|} * \|x\|}{1.3x}$
3 M CHF	-2.44	1.50	$y = \frac{\sqrt{\|x\|} * \|x\|}{1.5x}$
3 M DEM	-1.19	1.06	$y = \frac{\sqrt{\|x\|} * \|x\|}{1.1x}$
3 M USD	-2.75	1.38	$y = \frac{\sqrt{\|x\|} * \|x\|}{1.7x}$
DAX	-322.79	212.94	$y = \frac{\sqrt{\|x\|} * \|x\|}{20x}$
S&P 500	-70.04	41.87	$y = \frac{\sqrt{\|x\|} * \|x\|}{10x}$

Nikkey 225	-5057.29	4210.60	$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{100x}$
CHFDEM	-2.68	3.77	$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{2x}$
DEMUSD	-0.24	0.18	$y = \frac{\sqrt{\lvert x\rvert} * 2 * \lvert x\rvert}{x}$
SCH	-1.25	1.91	$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{1.4x}$
SD	-0.91	0.99	$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{x}$
SUS	-1.22	2.18	$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{1.5x}$
SCHD	-0.52	0.62	$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{0.8x}$
SUSD	-0.78	0.82	$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{0.9x}$
SUSJ	-1.08	1.39	$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{1.2x}$

The scaling transformation of the differences of the indicators contains symmetry because it does not contain an offset. The scaling formula for the predictor in the range [0.2…0.8] is:

$$y = \frac{\sqrt{\lvert x\rvert} * \lvert x\rvert}{2x} + 0.5$$

8.2.2 Feasibility analysis

The following figure shows a plot of the Bond CH and its first differences:

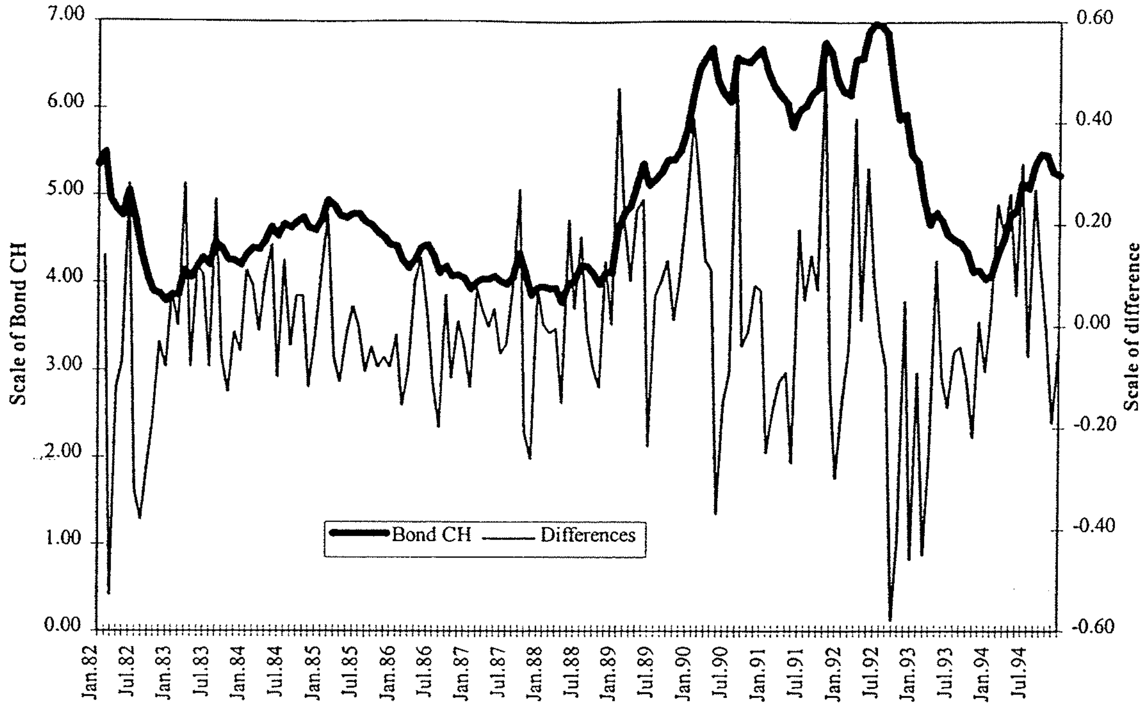

The figure shows the cyclic behaviour of the Swiss government bond. The available time window spans most types of market states like bull, bear, and congested markets. Tests on stationarity are not calculated, because the time window is also too small for computing and comparing the different moments of sub-windows.

8.2.2.1 Linear univariate analysis

A useful measure of the linear association between lagged variables is the autocorrelation function. The correlogram allows a fast interpretation of the results:

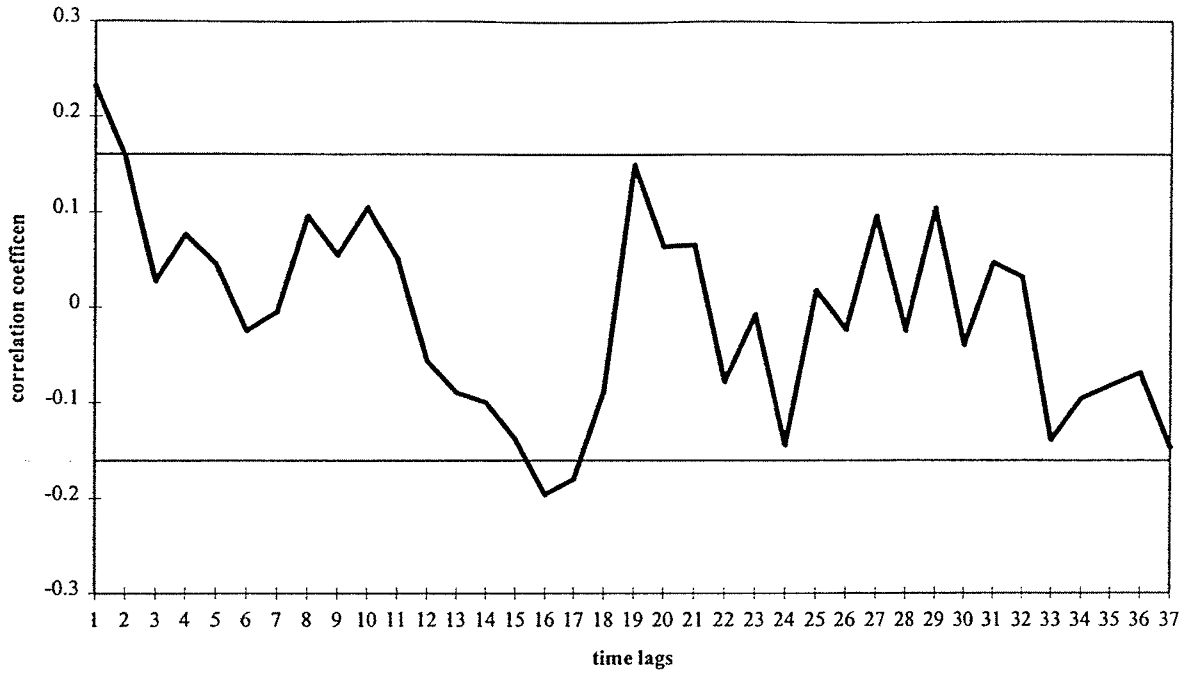

The monthly differences of the Bond CH shows significant linear autocorrelations at lag 1 and 16.

8.2.2.2 Non-linear univariate analysis

A calculation of the BDS test is difficult and not reliable with less than 200 datapoints [Brock91, 169f.], but the time series of the case studies contains only 155 datapoints. A less general test is the cumulated periodogram, which discovers non-linear periodicitis in the time series. Like the correlogram, the cumulated periodogram allows for a fast visual interpretation of the results like the correlogram.

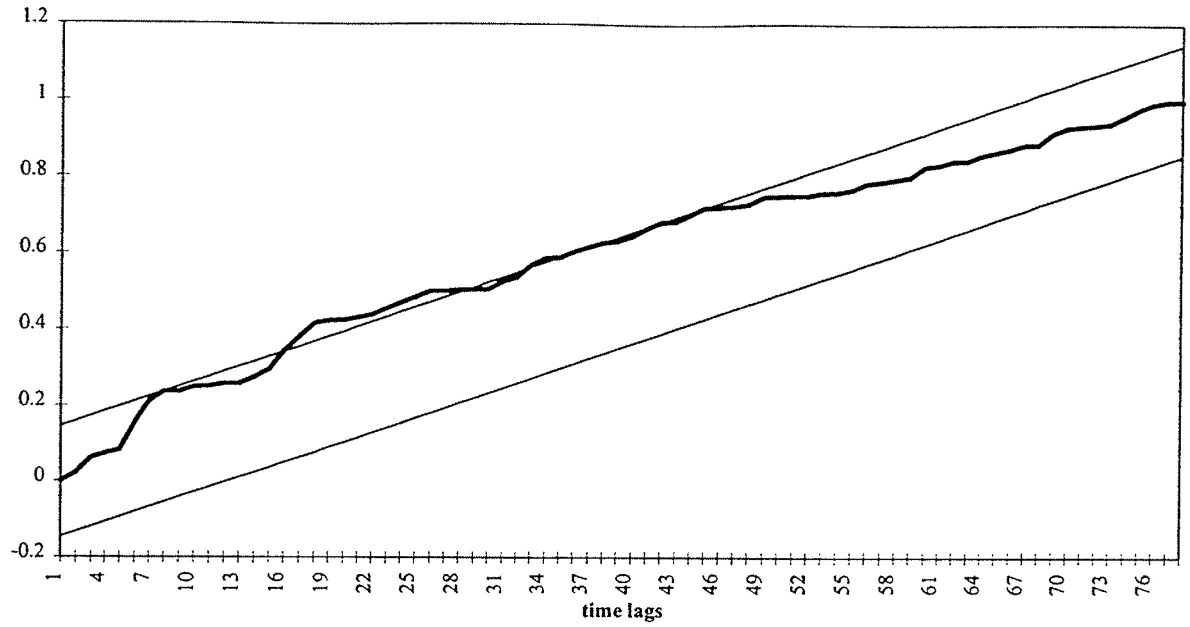

According to the cumulated periodogram test, Bond CH appears not to be a random walk, because the line crosses the significance level. This result is consistent with the results of the autocorrelation function, which also shows significant univariate dependencies.

8.2.2.3 Linear multivariate analysis

The cross correlation function measures the linear dependencies between the indicator time series and the predictor. The following list shows the cross correlation's coefficients for the possible indicators at time lag 1:

Bond D	-0.13
Bond US	-0.04
Bond J	-0.08
3 M CHF	0.00
3 M DEM	-0.15
3 M USD	0.02
CHFDEM	0.00
DEMUSD	0.00
SCH	-0.09
SD	0.11
SUS	-0.09
SCHD	0.05
SUSD	0.04
SUSJ	0.03

and the first differences of the possible indicators at time lag 1:

Bond D	0.14
Bond US	0.29
Bond J	0.07
3 M CHF	0.25
3 M DEM	0.08
3 M USD	0.08
DAX	0.18
S&P 500	0.13
Nikkey 225	0.06
CHFDEM	0.16
DEMUSD	0.07
SCH	-0.19
SD	0.01
SUS	0.08
SCHD	0.07
SUSD	0.17
SUSJ	0.15

The 95% significance level is 0.16. There are some significant linear correlations.

8.2.2.4 Non-linear multivariate analysis

The non-linear dependencies are measured with the neural network sensitivity analysis with bootstrap according to the description given in previous chapters. The test is calculated with six steps. The first step measures the dependencies of the absolute indicator values of the long interest rates. The second measures the dependencies of the short interest rates and the absolute values of exchange rates. The third step analyses the interest rate differences and the fourth, the differences of the least indicators. The fifth and sixth steps measure the absolute values and the differences of the spread indicators. A separation is useful with respect to the degrees of freedom of the neural network and the number of data samples. The absolute values of the stock indices are not analysed because they are trendy and not stationary.

The first neural network for the absolute indicators has four input units, three hidden units, and one output unit. The learning set contains 140 samples and the validation set contains 15 samples. The expectations of the sensitivity indices S and its standard errors are:

	sensitivity index	standard error
Bond CH	0.016	0.016
Bond D	0.019	0.018
Bond US	0.018	0.017
Bond J	0.020	0.018

The second neural network for the absolute values of indicators has five input units, three hidden units, and one output unit. The results of the sensitivity analysis are:

	sensitivity index	standard error
3 M CHF	0.014	0.017

3 M DEM	0.019	0.021
3 M USD	0.022	0.023
CHFDEM	0.019	0.020
DEMUSD	0.018	0.021

The third neural network for the difference indicators has seven input units, four hidden units, and one output unit. The results are:

	sensitivity index	standard error
Bond CH	0.025	0.019
Bond D	0.026	0.020
Bond US	0.031	0.020
Bond J	0.020	0.019
3 M CHF	0.023	0.017
3 M DEM	0.023	0.018
3 M USD	0.026	0.019

The forth neural network has five input units, three hidden units, and one output unit. The results are:

	sensitivity index	standard error
DAX	0.016	0.014
S&P 500	0.013	0.012
Nikkey 225	0.018	0.015
CHFDEM	0.012	0.012
DEMUSD	0.011	0.011

The fifth neural network of the absolute spread indicators has six input units, three hidden units, and one output unit. The results are:

	sensitivity index	standard error
SCH	0.023	0.019
SD	0.019	0.018
SUS	0.016	0.018
SCHD	0.017	0.017
SUSD	0.024	0.020
SUSJ	0.016	0.017

The sixth neural network of the differences of the spread indicators has six input units, three hidden units, and one output unit. The results are:

	sensitivity index	standard error
SCH	0.018	0.015
SD	0.018	0.016
SUS	0.016	0.015
SCHD	0.016	0.016
SUSD	0.017	0.015
SUSJ	0.018	0.017

8.2.2.5 *Conclusion and economic review*

The results of the univariate feasibility analysis show a trend in the behaviour of the predictor. The Swiss bond market appears to be inefficient. There are also different significant linear correlations.

The results of the multivariate non-linear feasibility analysis are slightly different from previous published results [Ankenbrand96a; Ankenbrand96b]. Possible reasons for the differences are:

- expansion of analysed indicators
- different indicator scaling

Similar to the SPI stock market case, the important influences are not only differences of the indicators. Most influences have different interest rates and aggregations. The exchange rates have no influence, and the stock market only a weak influence. The money and capital markets seem to have a leading function in the financial world.

The details of the linear and non-linear model design are described in the following section.

8.2.3 Model design

The result of the feasibility analysis suggests the need to develop a linear regression and a non-linear neural network model.

8.2.3.1 *Linear regression model design*

Significant correlations show the differences of the following indicators assuming the multivariate linear analysis. The linear correlations are only dynamic in nature.

- Bond CH difference
- 3 M CHF difference
- Bond US difference
- DAX difference
- SCH difference
- SUSD difference

It is possible to exclude the aggregated indicator SCH and SUSD without loosing prediction power, because they can transform linearly in the raw indicators. The final regression model is:

$$x(t+1) - x(t) = 0.98(x(t) - x(t-1)) + 0.14(ich3(t) - ich3(t-1)) + 0.4(ius(t) - ius(t-1)) + 0.0012(d(t) - d(t-1))$$

where:
x is the predictor
ich3 is the 3 M CHF
ius is the Bond US
d is the DAX.

The explanatory variables and the predictor are not scaled because linear regression models do not have the same limitations as neural networks regarding the magnitude of input data.

The regression model is optimised until the end of 1992. The rest of the data is the validation set. The optimised value is the risk-to-reward ratio which contains the net return and the maximum drawdown. The convergence results of the learning and the validation set are presented in this section.

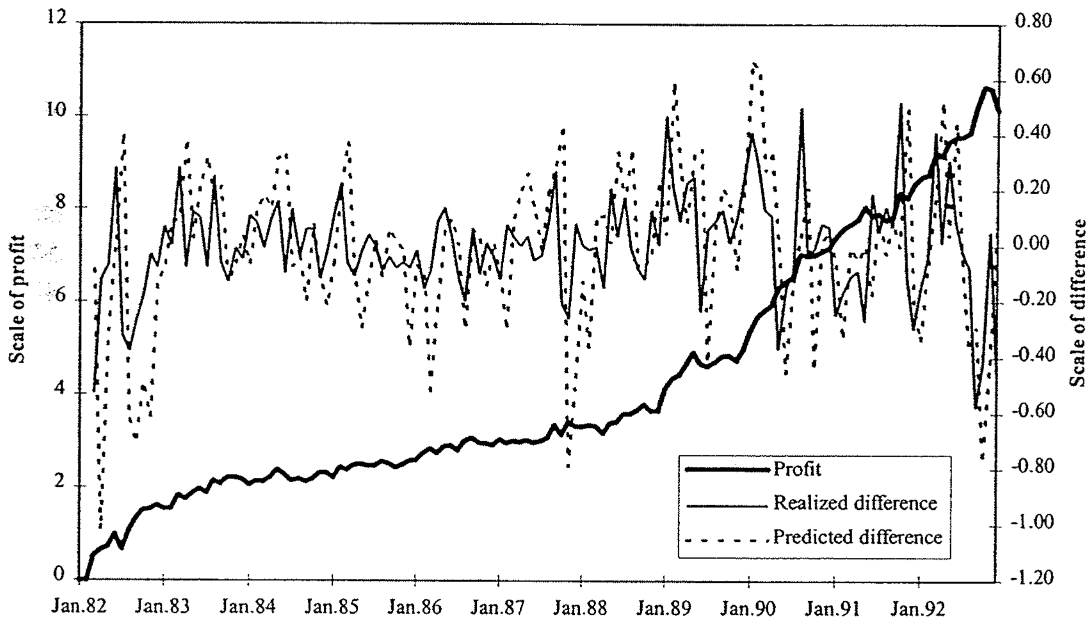

The same graph of the validation set performances is shown with different details:

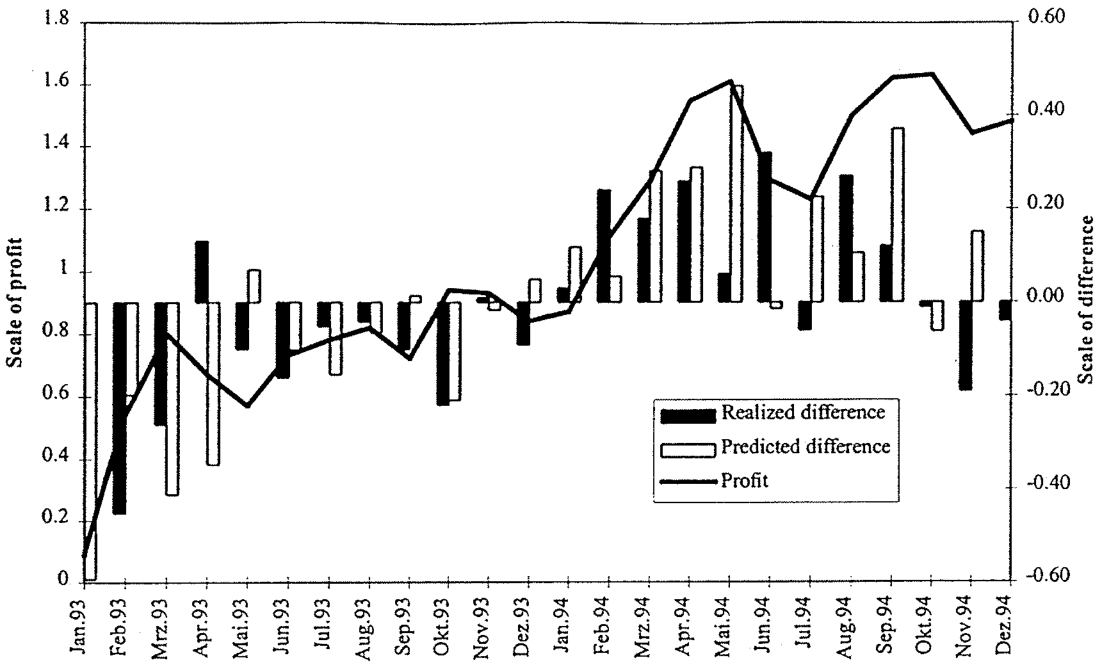

The residuals of the learning and validation sets are white noise according to the autocorrelation function and the cumulated periodogram. The trend of conformity in the learning set is 69 percent and in the validation set 66 percent. The profit, model efficiency, maximum drawdown and reward-to-risk ratio are shown in the following table:

	learning set	validation set
annualised profit	0.92	0.74
model efficiency	57%	42%
drawdown	0.51	0.38
reward-to-risk	1.80	2.64

The model appears to be stable, and the generalisation capabilities of the model should be good, as there are only four degree of freedoms, however, final judgement is only possible after the discussion of the out-of-sample performance.

8.2.3.2 Neural network design

The predictor time series model has 154 datapoints. A good size for an artificial neural network is 3 input, 2 hidden and 1 output units with 10 connections or degrees of freedom. This small size of the neural network reduces the danger of overfitting, and improves the generalisation capabilities. This limitation makes it necessary to choose the most significant indicators and not to use all the significant indicators.

The artificial neural network model is:

$$x(t+1) - x(t) = f\left[x(t) - x(t-1), ius(t) - ius(t-1), iusd3(t) - iusd3(t-1)\right]$$

where:
ius is the Bond US
iusd3 is the 3 M USD.

This model uses the indicator differences instead of the absolute values [Ankenbrand96b]. The previously published model with absolute indicators will be compared in the result section with the presently discussed models.

The learning set contains the first 140 data samples and the validation set contains the remaining 14 data samples. The backpropagation algorithm needs 290 learning cycles with a decreasing learning rate and a momentum term.

The following figure shows the development of the error term over the learning and the validation set.

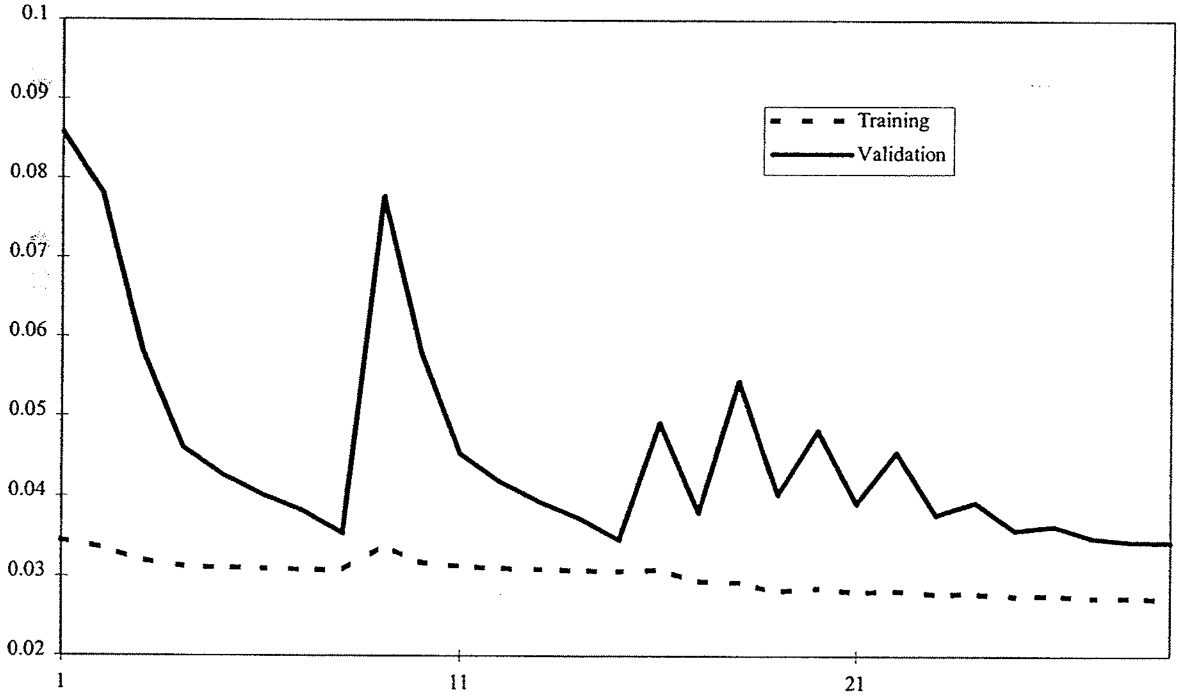

The next figure shows the development of the learning rate and the momentum term over the same period, and it explains the jumps in the error curves in the figure above.

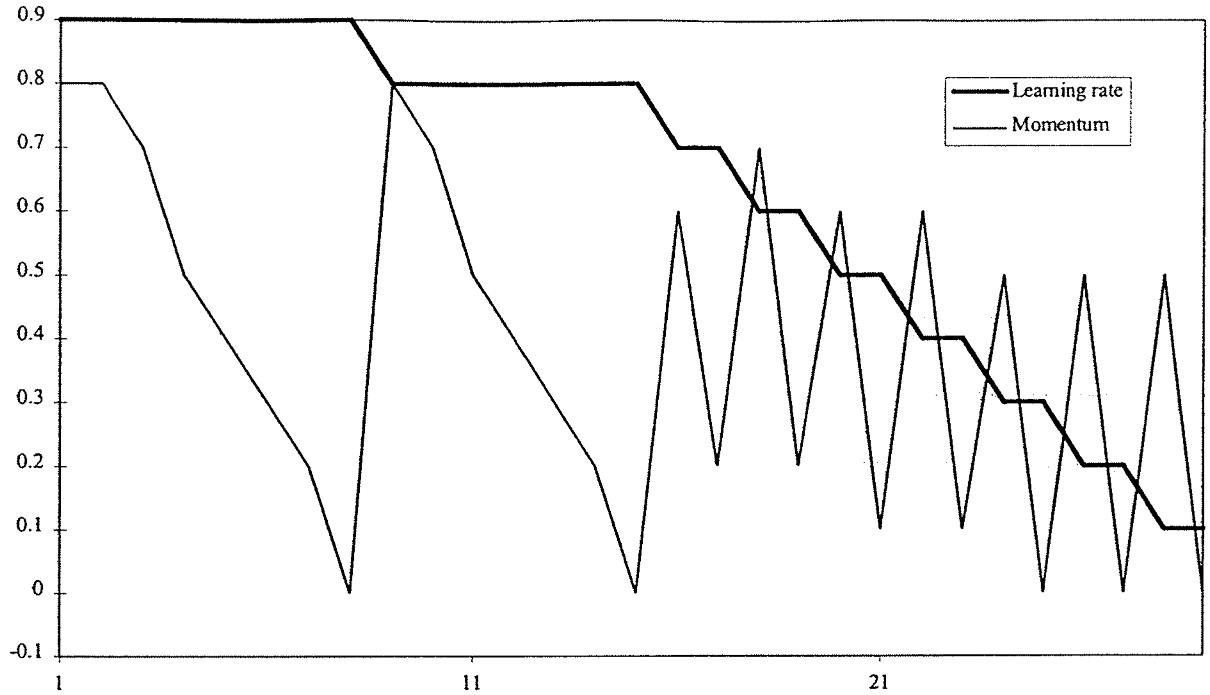

The in-sample performance and the convergence of the neural network are also presented in this section. The following chart shows the performance of the learning set.

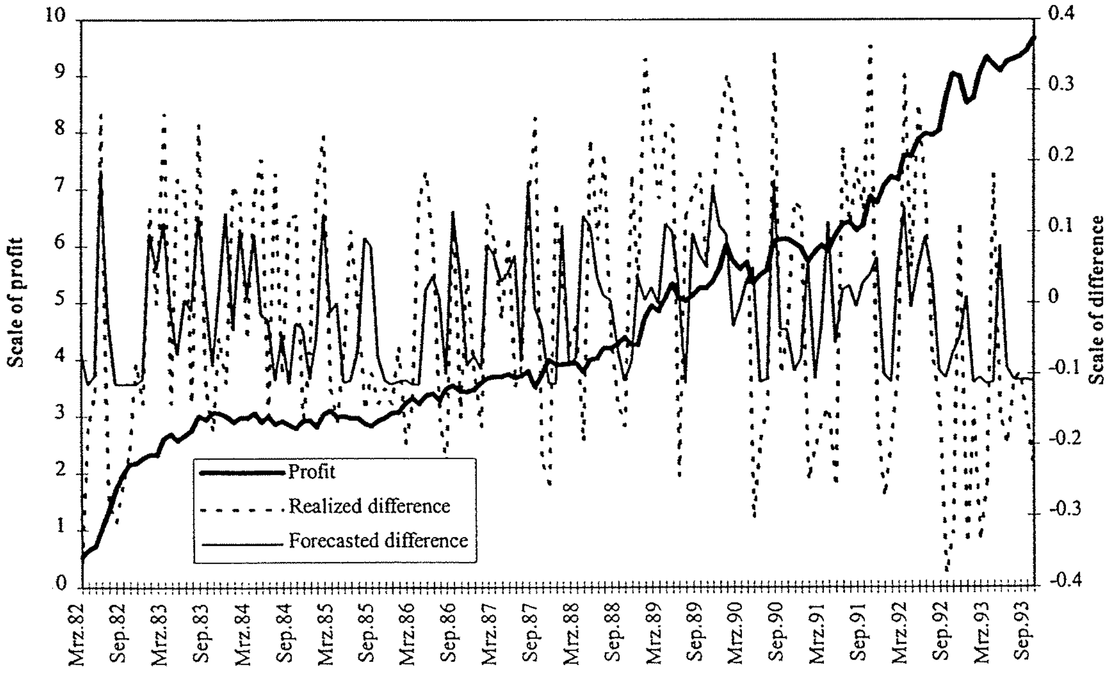

The next graph shows the performance of the validation set.

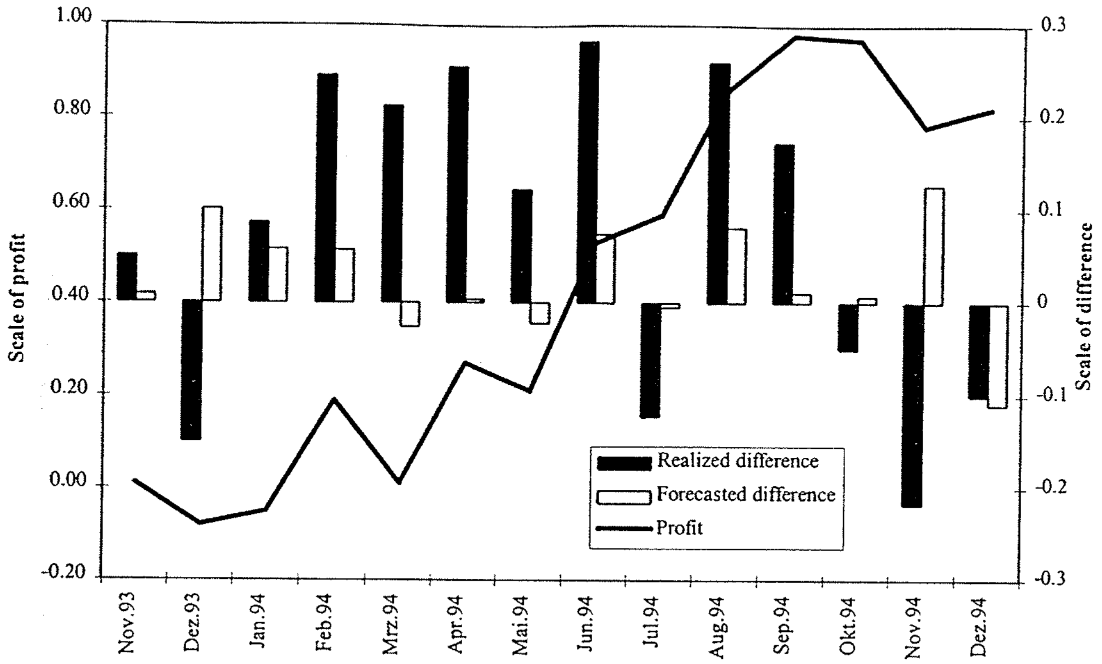

The residuals of the learning and validation set are white noise. The performance measurements of the training and validation sets are:

	learning set	validation set
trend conformity	66%	64%
annualised profit	0.84	0.70
model efficiency	50%	44%
drawdown	0.68	0.20
reward-to-risk	1.24	3.50

The model seems to be stable, but the performance of the validation set decreases slightly. A final judgement will be given in the next section with the out-of-sample data.

8.2.4 Results

Firstly, the out-of-sample performance of the previous published model [Ankenbrand96b] will be presented.

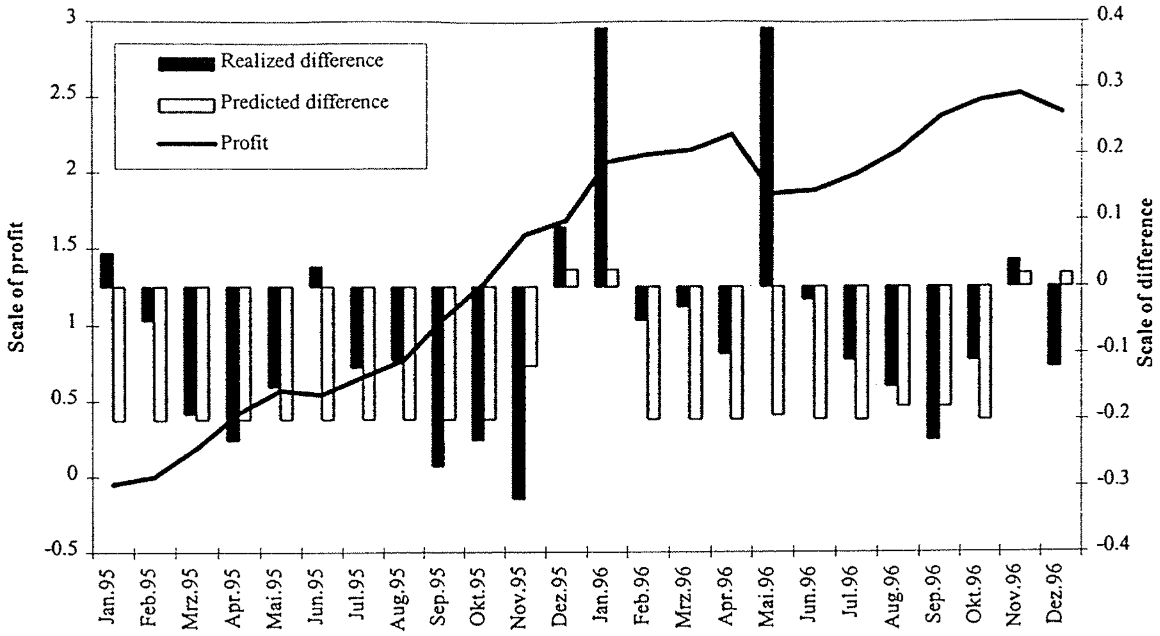

The next figure shows the out-of-sample performance of the linear regression model over the same time window.

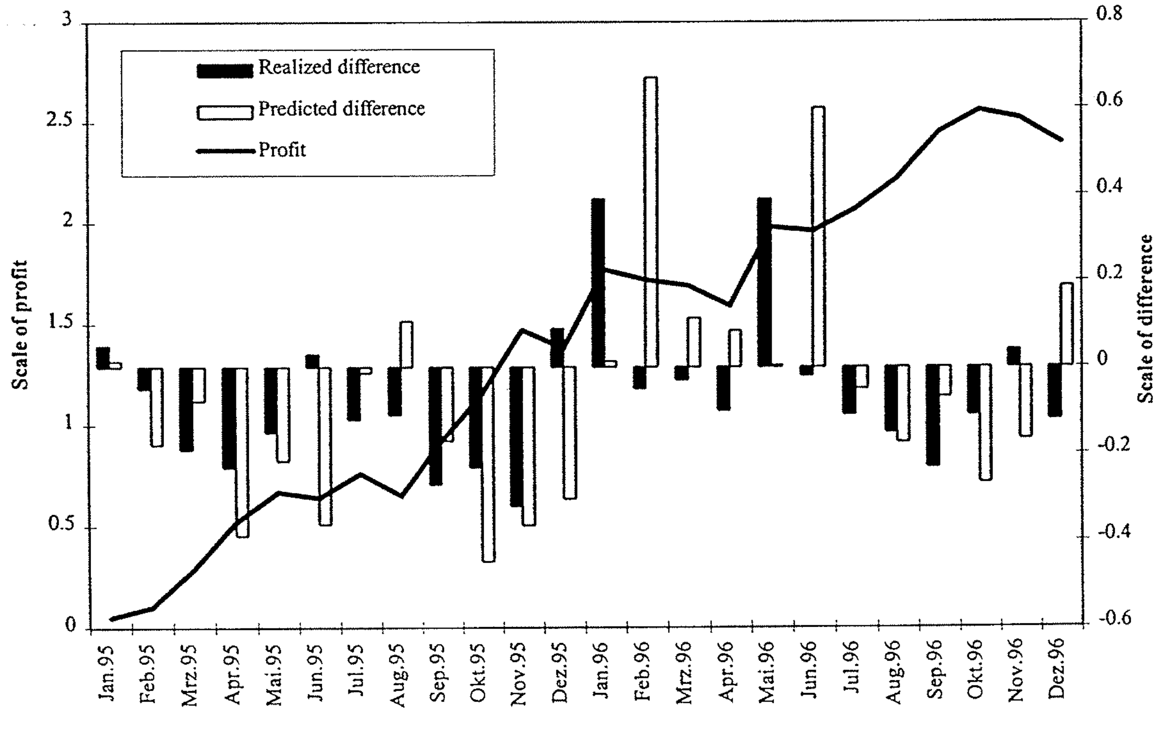

The neural network presented in this work has the following out-of-sample performance.

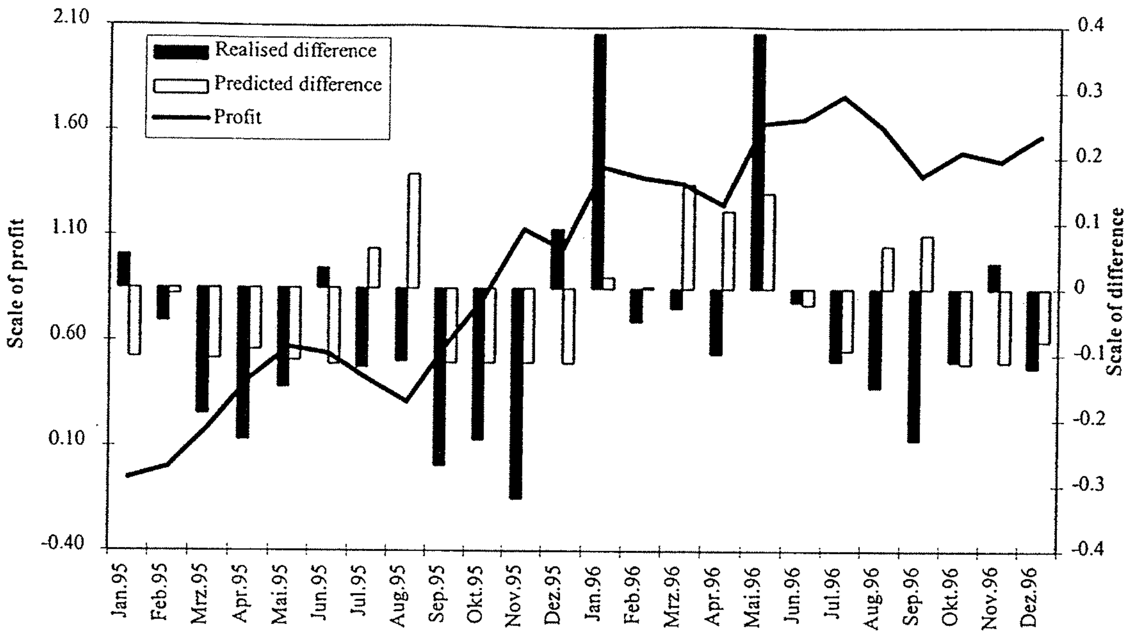

A residual analysis for the out-of-sample data sets are not performed because the time series are too short for significant analysis. The following table concludes the results of the other performance indicators for the three models:

	previous published	linear regression	neural network
trend conformity	83%	63%	54%
annualised profit	1.20	1.20	0.79
model efficiency	67%	67%	44%
drawdown	0.39	0.18	0.38
reward-to-risk	3.08	6.67	2.08

The presented linear and non-linear neural network models are stable and their performance is good. The linear model shows the best out-of-sample performance. The previous published model has a better generalisation capability than the new dynamic model based on the indicator differences. The in-sample performance of the dynamic neural network is higher, but it appears that the market dynamic changes faster than its structure. The out-of-sample time window is small so one must be careful with extensive interpretations because the ranking of the models could be different in another year.

A combination of the different models improves the results, because the three models have different architecture and databases. The linear combination of the three models is optimised until the end of 1994. 1995 and 1996 are held as an out-of-sample performance. The performance of the combined models is:

	in-sample	out-of-sample
trend conformity	74%	80%
annualised profit	1.05	1.27
model efficiency	64%	71%
drawdown	0.30	0.09
reward-to-risk	3.50	14.11

The model combination appears to be stable, and the performance is impressive. However, one must be careful because there are more degrees of freedom in this model combination than in a single model.

9. Artificial Market Environment

Distributed, decentralised, interacting assemblies of adaptive agents are commonplace in biology and ecology. One may think for instance of ant or bees colonies, fish schools or flocks of birds as well-known natural examples of the emergence or ordered, coherent collective behaviour without any central control or authority. These and other systems are all complex, evolving and adaptive. No theory of such systems is yet known; rather, their collective behaviour seems to result from the non-linear aggregation of simple adaptive local rules followed by the agents. In other words, the agents have a limited horizon and limited prevision and computation capabilities but, in spite of these limitations, the system as a whole appears to behave according to regular patterns without any central control. The macro-scale foraging behaviour of an ant colony emerges as a result of the micro-scale behaviours of ants without central control [Nakamura97]. The complex behaviour of financial markets can also emerge as the result of the collective action of many simple agents in mutual action, rather than due the complicated and unknown nature of each individual. Such a modelling approach of markets is going back to the roots of economics; it is going back to the invisible hand of Adam Smith.

Artificial life is a new discipline that strives to use computers, robots, and other artificial means to study life-like phenomena [Langton89]. The approach is synthetic rather than analytic and it works mainly by putting together systems that behave, in some respects, like living organisms. Computer simulation of populations of artificial entities is central to such an approach and has given insight into the behaviour of complex systems that would have been difficult to obtain with classical approaches such as differential equations dealing with aggregate quantities. Different works show an interesting collective behaviour of populations of artificial entities. Populations of artificial ants can show cyclic behaviours, but a slight change of the behaviour of a single ant can destroy the cycles. This can been seen as a manifestation of Wilson's multiplication effect whereby even a small change in the individual's behaviour may cause large social effects at the level of the collective [Beuret97]. Artificial life simulations can find new problem solving methods like the work of [Yokoi97] where the simulation of the behaviour of a multi-cellar amoeba leads to the idea of a new intelligent searching methodology. In most formal simulations, causes of fluctuations or dynamics are planned when the simulation model is built. However in artificial worlds, fluctuations or dynamics are caused by each agent's behaviour, evolution and interaction. Multi agents models and learning algorithms, like artificial neural networks and genetic algorithms make these methods possible. Some theoretical studies which pay attention to local behaviour or isolated interactions have already been done [Shinjoh97]. The artificial markets presented adopt the concepts of agents behaviour and interactions and integrate it in a realistic market structure. The application of the artificial markets simulation has the goal to model the behaviour of the real financial markets.

Artificial life and computer science in general will transform our culture. It is because it will help us to have a better understanding of brains. It will help us to learn what knowledge and learning are. This will change our science and especially the understanding of financial markets where psychology and collective behaviour of the market participants is a very important part beside the hard fundamental facts [Minsky97].

The first decade of artificial life as a research field has been characterised by a great deal of exploration into the possibilities inherent in the synthesis and simulation of life, and by the development of new technological approaches. These approaches are available for the application of practical simulation of real life systems. Despite the successes in computer engineering, adaptive computation, bottom-up artificial intelligence, and robotics, artificial life should not become a one-way street, limited to borrowing biological principles to enhance the engineering efforts in the construction of "life-as-it-could-be". The real value is the development of tools and methods which helps to understand "life-as-it-is" [Langton97].

Economic, financial and social systems share many features with the aforementioned natural systems and their artificial counterparts. Many important economic processes are complex in the sense that it is difficult to decompose them into separate parts that can be studied in isolation and aggregated to give the whole picture. Furthermore, economic agents do not seem to possess in reality the perfect rationality and computing abilities that the classical economic theory attributes to them. Issues of beliefs and prediction about the future become important and the individual beliefs and choices when aggregated shape the economic indicators, the market prices and ultimately the world that the agents must deal with. Moreover, this world is dynamic, as a result of the constantly changing collection of beliefs and strategies of the agents in their strive to adapt to the system evolution. In the real world, in contrast to common economic assumptions, there is also diversity between agents due to different computing abilities and computing procedures and different time horizons and risk profiles in the financial field. Computerised collection of diverse, adapting agents should then prove useful, especially in studying the dynamic aspects of economical systems which are often neglected in standard theories. Some recent work has shown that the approach of artificial life usefully complements the classical ones and can offer new insights of its own [Arthur95; Tesfatsion97; Foss98].

Of course, there are limitations to the simulation approach. The first question is: how simple can agents and the market structure be? Although very simple reactive automata have been employed for simulating economic scenarios, we feel that some of them are a poor approximation to the complex decision-making process that real agents in real economies undergo. On the other hand, too complex agents are difficult to come by, are computationally expensive and unnecessarily complicated. We have chosen a middle ground in which agents possess a certain amount of computational and forecasting abilities, are of diverse type and are able to learn and evolve. Another criticism concerns the simulation itself. A simulation implies that a choice has to be made of what the important components of the world are to be simulated. There are no unique answers to that question and a simulation can always be criticised on this ground. Likewise, it is difficult to draw quantitative conclusions from simulations: only trends, general behaviour and statistics are obtainable.

Nevertheless, simulations can be valuable as they may also give a partial solution to the hard problem of doing experiments with real socio-economic systems under controlled conditions [Davis93, 14]. It is difficult to study the consequences of specific policies or rules of behaviour in an orderly way because these systems are dynamic and ever changing. However, the computational approach allows to do so easily, repeatedly and without risk. This makes them useful as pedagogical devices and also as a source of ideas for new methods of analysis.

An example of such an environment is the realtime simulation of foreign exchange trading rooms at the Swiss Federal Institute of Technology (ETH) in Zurich. It operates on risk adjusted operational time scale derived from the volatility of the markets. It allows the different microeconomic agents to test its trading strategies. The system monitors the agents profit, risk and a money management with dynamic leverage. The system comprises the full functionality for a successful trading in the international foreign exchange rate [Schnidrig97]. The difference of the approach presented in this work differs from the one of the ETH by the following facts. The approach of the ETH is restricted to the foreign exchange market wherein the following approach presented also integrates the stock and bond markets. The following system tries to explain the behaviour of the financial markets from a macroeconomic point of view. The artificial agents fix a new artificial price with their buy and sell orders based on their trading strategies. The environment of the ETH does only monitor the behaviour of the artificial agents in the real markets what makes it for a suitable tool of testing new trading strategies.

Simulations can also be useful for normative economics. A simulation model of an artificial economic evolution shows that some policy intervention can improve the welfare. This refutes the laissez-faire criticism of evolutionary theory. The model bases on Schumpeterian ideas of competition in the presence of technological uncertainty and bounded rationality [Wakeley98].

9.1 The market simulation environment

In this work, we are interested in the generic-dynamic properties of financial markets. The artificial market is build around the behaviour of the real actors of the markets. These actors are represented by computer agents which interact together according to the rules of trading in the artificial market. There exist some old experiments with multiple homogeneous adaptive robots [Walter51]. We are not interested in the exact model of a single agent like in the previous micro-economic forecasting models but in the different robot applications from the artificial life field, for example see [Miura97].

The originality of the approach presented here resides in its multi-market nature. Until now, market simulations have been done in a single market context and with a single good traded (see for instance [Palmer94]). Such a restricted single dynamic market simulation can help to improve the understanding of the behaviour of exchange rates without the classical economic assumption of homogenous and rational economic agents [Tordjman97]. While this can be helpful to simplify the simulation and to control its parameters, real markets are strongly correlated, especially financial markets. Therefore, we strive for including inter-market effects in the simulation, albeit in a very simplified form, by linking several markets and allowing trading in several of them. In addition we try to stay as close as possible to real markets by using a plausible market typology and geographical distribution.

A simplified multi-market simulation is the simulation of the currency system of the European Exchange Rate Mechanism (ERM). The simulation environment contains five participants: small and large speculators, hedgers, central bank and bid/ask traders. The market participants have different information, algorithms, models, constraints and utility functions. The artificial market contains five currencies: the Italian lire, the British pound, the French franc, the Swiss franc and the Deutsche mark. The agents can not change their strategies and the lot size is also not depending on the success. The artificial agents use real data for their decisions and the artificial prices based on the bid and ask orders are compared to the real financial markets. The results of the simulation depend of the initial setup because the artificial market environment is not adaptive [Maza97].

This market environment presented contains eight markets. There are three stock markets and three bond markets in three different currency and two foreign exchange markets for the interconnection of the different currency regions. This is similar to the three main actual financial regions: Asia represented by Japan, Europe represented by Germany and the United States of America. This allows to simulate international diversified portfolios and enhance the variability of the agents, because the trading profile of an investor is different from the trading profile of a speculator. Another advantage of the diversity of markets is the possibility of analysing the relations between different regions and different assets [Poddig94, 368ff.]. This advantage is very important for practical applications because market participants are rarely active in one market only.

Only one good is traded in every market. This is realistic and easily understandable for the currency exchange markets. The bond and stock markets are simplified because there is only one asset traded in every market, whereas in reality different stocks and bonds are traded in one single market. The bonds are assumed to have the same duration and the same coupon. The assumptions are simple but realistic because they amount to a situation very similar to trading of an index or to trading the derivative financial instrument of the future market. The future markets today are very important with very high and increasing volumes. The market environment seems to be realistic and allows to simulate inter market relations and to compare the simulation with time series of the real financial markets around the world.

The trading process is iterated, and consists of the following steps:

- Order placement of the agents.
- Calculation of the price.
- Accounting of the orders in the book of the agents.
- Adaptation of the trading strategies of the agents.

Firstly, every agent places his buy or sell orders in the order book of the desired market based on its trading strategy. An agent can use the past real data of the eight markets which allows to make a realistic trading decision. The size of the orders is depending from the present net asset value. The net asset value of the agent is given by the sum of all the assets held by an agent multiplied by their respective real prices. The order book contains all orders of a specific asset. The agents can only use market orders, which means that the order is executed at the best possible price that can be obtained at the time. The next step is the calculation of the artificial price for every asset. The artificial price of the asset for the actual trading cycle is calculated according to the following formula:

$$y_t = x_{t-1} + 0.1 * \frac{x_{t-1} * (o_{buy} - o_{sell})}{|o_{buy}| + |o_{sell}|}$$

where:
y is the artificial price
x is the real price
t is the time or trading cycle number
o_{buy} is the number of buy orders
o_{sell} is the number of sell orders.

The third step of a trading cycle is the accounting of the executed orders: the assets and the cash are added or removed to or from the account of the agents based on the real price, and the new value of the cash and asset portfolio is calculated for every agent.

All orders placed are executed every trading cycle because there exists a market maker for every market, who has the obligation to execute the difference of buy and sell orders for his own account. The assumptions of the market environment are simple but realistic and comparable to real financial markets [Cohen87, 48ff.].

The forth and last step of every trading cycle is the adaptation of the local trading rules of the agents through the use of a genetic algorithm. This phase is described in the next chapter in the section of the strategy evolution. Then the next cycle can begin with the placement of the orders of the new generation of the agents based on their trading rules.

9.2 Trading strategies of agents

In general object-oriented techniques are used for modelling such advanced information systems. Object oriented approaches are successful in capturing the properties and behaviour of all kind of real world entities. The object represents a snapshot of the entity at a given time. But the properties and behaviour of objects change fast in highly dynamic environments like the financial markets. Therefore restrictions on the object model in dynamic environments lead to a new semantic model of the agent. Agents differ from conventional objects in several ways. First of all, agents are not only reactive. They are active in the sense that they are able to act goal driven without external stimuli. In addition, agents are flexible. They may change their behaviour dynamically during system run-time. They are also able to deal with unexpected, unpredictable situations [Türker96]. The properties mentioned above are useful for modelling the behaviour of financial market participants which are an important part of the artificial market simulation.

Agents are interpreted as dynamic and intelligent objects. They are able to be autonomous or cooperative, and reactive or proactive. Agents have an internal state which is based on their history which influences their behaviour. The internal state reflects the knowledge of an agent at a given time. This knowledge could be changeable during the lifetime of an agent. It includes an internal, imperfect representation of the world or environment of the agent. The agents are assumed to have knowledge acquisition and revision capabilities in order to be able to changes their knowledge in the dynamic environment. Agents have also goals which they try to achieve under given constraints. Since goals are part of the internal state of agents, they may be changed during an agent's lifetime. Each agent is obliged to satisfy his goal by executing a sequence of actions. The actions can be reactive or proactive. The agents can also cooperate with other agents to achieve their goals [Türker96].

There are various types of agents in the different simulations. The different parameters of the trading strategies are coded in the genes of the agents in the dynamic simulations. The adaptive agents contain seven genes. The first gene contains the type of the agent. The other six genes depend from the agent type and represent the parameters of the trading strategies. The information of the genes are represented by integers between -10 and 10. They are transformed from every agent type for fitting the parameters of the respective trading strategies. The agents are simple reactive agents in the present model. The input to the agents are the past prices of the assets and their output is the order which they place in the order book on every trading cycle. The agents do not have any information about the behaviour of the others agents with exception of the price behaviour of the different assets. The importance of an agent increases with his success. The lot size of the trades increases with the accumulated profit of the agent. This should simulate the flows of trading and investment capital to the successful players in the financial markets. It mimics the money management with a dynamic leverage and reinvesting of the profits which is quite common in practice.

The actions of the agents are not influencing the other agents directly. The influence can happen indirectly over the evolution of the strategies of the agents with the genetic algorithms. The agents cannot cooperate or influence the decision or the success of the other agents. The success of the agent depends on the prices in the real financial markets which are also used for the accounting of the orders. This is a main difference to the common simulations of game situations like the prisoner dilemma [Akira97; Akiyama97, Beaufils97, Karandikar98]. The artificial market presented is an evolutionary game in the sense that higher payoff strategies tend to displace lower payoff strategies over time and that the adaptation is an evolution and not a revolution [Friedman98].

The next sections describes the trading rules of the different kinds of static and dynamic, adaptive agents.

9.2.1 Technical traders

Technical analysts study market data in an attempt to gain insight in the future behaviour of the asset. They do this by looking for recurring patterns of price movement. In our artificial population, technical trading agents are univariate traders. Many technical analysts believe in a trend behaviour of the markets and try to use moving averages or momentum analysis for detecting trends. The disadvantage of the trend theory is that the traders buy assets when they are high or expensive and sell when they are low or cheap. The contraries traders do the opposite: they buy assets which are out of favour and they run against the crowd [Cohen87, 252ff.].

There are two types of technical trading agents for the static simulation. The first kind are the trend followers: they buy if the price difference of the last trading cycle and the previous of the last trading cycle is greater than a threshold value, they sell if the price difference is smaller than a threshold value. They do nothing if the price is between the upper and lower threshold value. The second type are the contraries traders which do the opposite: they sell if the price difference of the last trading cycle and the previous of the last trading cycle is greater than a threshold value and they buy if the price difference is smaller than a threshold value. They do nothing if the price is between the upper and lower threshold value.

The technical trading agents for the dynamic simulation are more flexible. The second gene codes the market in which the agent is active. The static technical trading agent is active in every of the eight markets. To remember the first gene represent the type of the agent like technical trader, evolutionary multivariate trader or artificial neural network trader. The third gene represents the time window for calculating the price difference or momentum. It is transformed to always be positive and to lie between 5 and 20 days. Gene five represents the threshold value. If the threshold value is positive and the difference greater than the threshold value, the agent buys the assets. If the threshold value is positive and the difference smaller than the negative threshold value, the agent sells the assets. If the threshold value is negative, the agent does the opposite. A positive fifth gene represents a trend follower agent, a negative gene means a contraries trader. The fourth gene is a short term time window between 1 and 4 with gene 6 as the threshold value. The decision process is the same as in the previous long term time window. The seventh gene makes the distinction between the short and long term decision. The final trading decision is multiplied with the lot size and transmitted to the order book of the desired market. Technical traders can take long and short position. Short positions mean to sell assets and buy them back later for profiting from falling prices.

The use of genetic algorithms for finding technical trading is also successful for its own [Neely97] what makes the chosen configuration very reasonable as a part of the artificial market environment.

9.2.2 Investors

Actually, a static investor has more long-term oriented objectives than those of the previously described artificial traders. The difference between investors and traders or speculators is the portfolio approach. The investor does not concentrate on some specific markets or assets because he believes in the efficient market theories according to which it is impossible to forecast markets for profitable trading. The investor diversifies his capital in different markets and tries to reduce the risk with an efficient portfolio without loosing opportunities of profit offered by the different assets. He is thus interested in long-term multivariate investment decision. The investors base on the theory of Marcowitz.

Most investment consultants make a distinction between tree types of investors. Firstly, the risk averse investor who is interested in protecting his capital. The second type is interested in making a profit with an acceptable risk level. The third type believes in a high profit and accepts the risks of the aggressive strategy.

The risk is measured by the variance of a portfolio. The general formula for the variance of a portfolio is as follows [Cohen87]:

$$\sigma = \sum_{i=1}^{N} \sum_{j=1}^{N} w_i w_j \sigma_{ij}$$

where:
N is the number of assets in the portfolio.
i, j are subscripts used to refer to particular assets.

w is the weighting of each asset invested in the portfolio.

The first type of investors, the risk averse investors try to minimise the variance of their portfolio in the simulation environment. The second growth-oriented type tries to optimise the reward of the risk ratio which is the coefficient of the profit and the risk. The third type seeks to optimise the profit. The investors adjust their portfolios monthly based on the past profits, variances and covariances of the assets. They can only buy stocks and bonds and they buy only currencies for buying stocks or bonds in foreign markets. The classical investors are not allowed to take short positions of stocks or bonds and they can only sell assets that they own. The investors have different home currencies and the strategies could be different depending on the home currency of the investor. Some investors are also limited in investing in foreign markets. They only invest in their home market. The different kind of investors are fix coded in equal numbers. The only adaptation happens over the lot size which improves the importance of successful kinds of investors.

9.2.3 Arbitrage traders

Arbitrage is the simultaneous purchase and sale of assets in order to profit from distortions in price relationships. Variations include the simultaneous purchase and sale of bonds or stocks in different markets. Arbitrage is riskless because the buy and sell orders are simultaneous; it exploits transient inefficiencies in the financial markets. Arbitrage trading is important for a realistic interconnection of the different markets in the simulation environment. Arbitrage traders are static multivariate short-term traders.

The first type of arbitrage traders connect the three bond markets. They buy the bond with the lower price, and simultaneously sell the bond in another market with a higher price. The second type of arbitrage traders connect the three stock markets. They buy the stock in a cheap market and sell it in an expensive one. By valuing a stock, a trader must also take into account the currency exchange rate of the two stock markets. The reason of different stock prices in different markets can also be a mispricing of the currency exchange rate. The arbitrage traders of the stock markets have to sell and buy the related currencies. The arbitrage traders of the bonds do not need to take into account the currency exchange rate because the bond prices are the discounted cash flows of the future and these cash flows are in the same currency as a straight bond. The different types of arbitrage traders are fix coded in equal numbers with the same simple lot size based adaptation method like the static investor.

9.2.4 Evolutionary multivariate trader

The evolutionary multivariate trader holds a portfolio whose composition is coded in his chromosome. The second gene multiplied with the lot size represents the number of assets of the first bond market. The third gene is responsible for the holdings of the second bond market. The other four genes fix the numbers of assets of the remaining bond and three stock markets. The evolutionary multivariate trader adapts the portfolio if he is involved in artificial recombination or mutation during the adaptation of the strategies of the agents or if his net value is increasing, allowing him to change his position. He can also take short positions based on his genes. The evolutionary multivariate agent is a very simple adaptive model of the static investor and the static arbitrage trader described in the previous two sections.

The investment of the evolutionary multivariate trader is not based of any assumption like the static investor and arbitrage trader. He can hold the same portfolios like these two types of agents. He represents the classic agent of a genetic algorithms which is searching for an optimal portfolio. The evolutionary multivariate agent is able to replace the static investor and arbitrage trader in the dynamic simulation for that reason. This step simplifies the simulation environment and reduces the a priori assumptions of the artificial market environment.

9.2.5 Artificial Neural Network trader

The artificial neural network trader is based on the previous described feedforward networks with the backpropagation learning algorithm of chapter 6. The neural network has three input units, two hidden units and one output unit. The predictor time series and the three input time series are presented from the second gene to fifth gene. The length of the learning set is coded in the sixth gene. The seventh gene is responsible for the relearning interval. The neural network can relearn before every trading cycle or the weights are fixed after the initial learning. The genes represent the genotype of the agent. The phenotype of the agents is mapped by the weights of artificial neural network from the genotype. The weights can be adapted without a modification of the genes. Individual lifetime learning can guide an evolving population to areas of high fitness in genotype space through an evolutionary phenomenon known as the Baldwin effect [Baldwin96; Hinton87]. It is accepted wisdom that this guiding speeds up the rate of evolution because it increases an individual's survival chances or fitness. Another interaction between learning and evolution, will be termed the Hiding effect. The Hiding effect shows that learning can reduce the selection pressure between individuals by hiding their genetic differences in the mapping of the phenotype. The differences between the genotypes are reduced which slows down the rate of evolution. There is a trade-off between the Baldwin effect and the Hiding effect depending on the cost of learning. The Baldwin effect dominates when the costs of learning are high. Low costs favour the Hiding effect. [Maley97]. The artificial neural network trader have also no a priori assumptions like the multivariate evolutionary trader.

9.2.6 Market makers

The market makers have the obligation to execute the difference of buy and sell orders of the other agents for their own account. This means that the agents find always a counter part for their trades. This is a realistic assumption for the trading in financial markets. Especially the organised exchanges have nominated market makers which clear the markets and guarantee liquid markets. The market makers profit in reality from smaller transaction costs for their service. Transaction costs are assumed away in the presented artificial market simulation environment. The market maker has not a benefit for his service in the simplified artificial market.

There exist one market maker for every market. His strategy is not evolving and he has no genes for that reason. He has unlimited credit for caring out his task. He will not be killed in the dynamic simulation if his net asset value is under the limit of -10, and he is not depending of the lot size for trading. The market maker has not an active strategy on his own. He only executes the remaining orders of the other agents without any intention.

9.3 Strategy evolution

The goal of the trading strategy evolution is to improve the local trading rules of the agents for maximising the profit. Genetic algorithms are well suited for solving such optimising problems. The metaphor of the genetic algorithms is the natural evolution. In evolution, the problem each species faces is one of searching for beneficial adaptations to a complicated and changing environment. The idea behind genetic algorithms is to do what the nature does. Let us take a population of rabbits. The smarter and faster rabbits are less likely to be eaten by foxes, and therefore the smartest and fastest survive to make more rabbits. But some of the slower and dumber rabbits also survive by luck. The surviving population of rabbits is a good mixture of rabbit genetic material with high diversity of fast, slow, smart and dumb rabbits. The resulting baby rabbits will on average be faster and smarter than their parents because more faster, smarter parents survived the foxes. Additionally, the genetic material of rabbits population is randomly changed which brings new properties in the rabbit population. The foxes are undergoing similar process. Otherwise the rabbits might become too fast and smart for the foxes to catch any of them. A genetic algorithms follows closely the story of the rabbits with a step-by-step procedure [Michalewicz96].

Genetic algorithms use a vocabulary borrowed from natural genetics. One would talk about individuals, genotypes, strings or chromosomes. This might be a little bit misleading: each cell of every organism of a given species carries a certain number of chromosomes. The individuals of genetic algorithms are one-chromosomes individuals. Chromosomes are made of units arranged in linear succession. The units are called genes and control the inheritance of one or several characters. Each genotype or individual would represent a potential solution to a problem. An evolution process run on a population of chromosomes corresponds to a search through a space of potential solutions. Searching strategies requires balancing two objectives: exploiting the best solutions and exploring the search space. Hillclimbing is an example of a strategy which exploits the best solution for possible improvement, but it neglects exploration of the search space. Random search is an example of a strategy which explores the search space ignoring the exploitation of the promising regions in the search space. Genetic algorithms are a general purpose search methods which balances the exploration and exploitation of the search space [Michalewicz96].

Genetic algorithms are more robust than existing directed search methods because they combine elements of directed and stochastic search. Another important property of such genetic based search methods is that they maintain a population of potential solutions. All other methods process a single point of the search space. Genetic algorithms perform a multidirectional search by maintaining a population of potential solutions. The relative good solutions reproduce, while the relative bad solutions die at each generation. An objective or fitness function is used to distinguish between different solutions [Michalewicz96].

The genetic algorithm maintains iteratively a population of potential solutions. Each solution is evaluated to give some measure of its fitness. A new population is formed by selecting the fitter individuals after their evaluation. Some members of this new population undergo alterations by means of crossover and mutation to form solutions. Crossover combines the features of two parent chromosomes to form a similar offspring by swapping corresponding segments of the parents. The selection of the parents is random with the probability of the fitness. A fitter chromosome has the greater chance for a selection what improves the distribution of the best chromosomes. Some chromosomes can be selected more than once. The intuition of the crossover operator is information exchange between different potential solutions. Mutation arbitrarily alters one or more genes of a selected chromosome. The change is random with a probability equal to the mutation rate. The intuition of the mutation operator is the introduction of some variability into the population. The best chromosome represents an optimal solutions after some number of generations when no further improvement is observed [Michalewicz96].

Genetic algorithms theory explains that the population converges to the sought optimal point. But practical applications do not always find the optimal solution because there is a limit on the hypothetically unlimited number of generations or unlimited population size. The coding of the problem often moves the genetic algorithms to operate in a different space than that of the problem. Genetic algorithms can fail like the other searching methods by converging to a local optimum. If convergence occurs too rapidly, then the valuable inform developed in part of the population is often lost. Selective pressure supports the premature convergence and decrease the diversity. A weak selective pressure can make the search ineffective. It is important to strike a balance between these two factors [Michalewicz96].

The application of evolutionary principles in economics is not a new idea [Veblen98]. The evolutionary phase is intended to mimic natural selection in that relatively successful strategies tend to be preserved in the population. Artificial evolution is used to obtain desired functions subject to mutation and selection under selection pressure of the real financial market. Artificial evolution is directed evolution to obtain the desired agents through survival of the fittest by the environmental conditions [Weller97]. The evolution is not focused on an optimisation to a fixed goal because the fitness function depends of the success of the agents in the real financial markets which is changing itself. The fitness landscape is under evolution itself. The evolution tends more to a relative adaptation of the artificial market represented by the agents than to an absolute optimisation [Pollack97]. Rapid evolution in response to changing environmental conditions has been observed in several ecological systems so far, but the mechanism of such adaptations on fitness and the rapidity of their evolution is not sufficiently clear [Kikuchi98].

Each individual is represented by a chromosome which encodes the agent type, the time horizons, thresholds and the trend-follower or contrarian nature for the technical traders, while the multivariate traders have genes representing their positions on the different markets. The artificial neural network agents have genes encoding the input and predictor time series, the number of learning and validation sample and the relearning frequency. Every gene is coded as an integer in a given domain. The integer representation is faster, more consistent, and provides a higher precision than the classical binary representation of the genes [Michalewicz96].

The population evolves according to a standard genetic algorithm [Michalewicz96] and goes through

- fitness evaluation

* selection
* recombination
* mutation.

The individual's fitness is given by the net asset value of the agent, which is the sum of all the assets held by an agent multiplied by their respective real prices. The fitness function is changing over time because the financial markets and its prices are also changing over time. Selection is based on relative fitness, new individuals are obtained through standard one-point crossover with a crossover probability of 0.5 and mutation is applied to every chromosome field with a rate of 0.02. Up to one tenth of the population is replaced at each generation, discarding the agents with net asset value smaller than -10. The elite 5% traders go to the next generation unchanged.

The evolution of the different trading strategies of the agents mimic the selection and learning of the participants of the real financial markets. It is not comparable to the human gene pool because this pool is more or less constant over the against run of a decade or so [Friedman98]. This adaptation and the use of non-linear and adaptive trading strategies reflects the assumption of non-linearity and drift of the financial markets. The next chapter will show that the use of non-linear and adaptive agents improves the similarity of the artificial market and the real financial markets which is the ambitious goal of the simulation.

10. Artificial market simulation

10.1 Simulation environment

Unlike most artificial life studies we use financial market data to verify the assumptions and the design of the artificial market environment. The primary goal of this work is to show that minimal conditions of real markets are satisfied with the assumptions and design of the artificial markets presented to observe real life phenomena. The similarities of artificial markets and real markets are measured by comparing the price time series of the artificial y and the real financial markets x. Introducing a number of ad hoc parameters one could get a more realistic and reliable model of artificial markets without loosing the advantages of pure artificial markets. The parameters are obtained from past time series observation in real financial markets which the agents can use for their trading decisions. This allows to integrate different partial theories especially coming from the micro or agent level or the macro level of the markets and to test this hypothesis about the structure and the development of the real markets by making predictions about the trend and volatility of the different markets [Ferber94, 15f.].

There exists a single artificial market simulation without the integration of real prices for the trading decisions of the agents. The two classes of agents, fundamental investors and technical speculators, are able two simulate a quasi real distribution of returns with fat tails [Lux98]. It shows the usefulness of such an approach for modelling exchange relations. But this single market simulation neglects completely important relations between different markets which is the original approach of the artificial market environment presented.

The real market data used are daily time series from 1. October 1991 to 30. June 1995. This time window is a good representation of the different market types, like bull, bear and consolidation markets, for the different stock, bond and foreign exchange markets. The length of the time series allows to simulate around 900 trading cycles. The three stock markets are represented by the following stock market indices:

- Nikkei 225 for the Asian region
- DAX for Europe
- S&P 500 for the United States of America

The bond markets are represented by the following interest rate futures with a similar duration and coupons:

- Japan Government Bond for Asia
- German Government Bund Future for Europe
- 10 Year US Treasury Note for USA

The currency exchange rates are:

- YENUSD
- DEMUSD

The time series used span the most important financial markets of the world and represent a large part of the daily trading activities around the financial world. The time series are linearly scaled between 0 and 1.

A simulation environment should allow to make experiments which are impossible in reality because the parameters of the environment can not be controlled. This is a general dilemma of social sciences and also economy. The experiments of the simulation environment are only meaningful for the reality if the simulation environment and the reality have the same properties. The environment is realistic and useful if the behaviour of the simulated prices are similar to the observed prices of the different financial markets. If the simulation is qualitatively faithful with respect to the reality, then a further step would be to use the simulation to qualitatively forecast the actual financial markets which makes artificial markets to a useful tool for investment and risk management.

A simulation environment should be as simple as possible from an engineering point of view and in respect to the degrees of freedom which are an indicator to the generalisation capability of the artificial world. This fact allows to simplify the trading process and assuming away the transaction costs. The desired characteristics are the use of simple and identical agents, adaptability to environmental changes, reliability and robustness of the simulation environment [Holland97]. The artificial market environment presented has a very simple structure without a matching technique. Instead there are the market makers similar to the ones in the organised exchanges. This is not really true for foreign exchange markets, but this simplification can be neglected for analysing daily data.

The static artificial market environment is developed with MS VisualBasic for Excel. This implementation environment is simple with a prototyping character which is very useful at the beginning. The disadvantage is the slow execution on the computer which makes it difficult to simulate more numbers of more complex agents. For that reason the dynamic and complex simulation are implemented with Borland C++. There also exist different tools for the simulation of artificial worlds [Ackley97]. We have evaluated different tools but they have not fulfilled the requirements. The expenditure for the adaptation was similar to the effort of the development from the scratch. The development from the scratch has the advantage to control all parameters and there is not a necessity to make a compromise by the implementation of the ideas.

10.2 Static simulation

The first two simulations are very simple. There is no adaptation of the trading strategies of the agents. The agents do not evolve. The only possibility of adaptation is the fact that better performing agents improve their lot size in relation to the improvement of their net asset value.

10.2.1 Configuration

There are two simulation. The first simulation contains 150 agents and the second simulation contains 300 agents. There are the following types of traders:

- Technical trader
- Investors
- Arbitrage trader
- Market maker

The technical traders are trend followers and contrariens based on the parameters which are described in the previous chapter. There are also some very simple adaptive traders, which optimise their parameters, like time horizon and threshold, based on the past price time series. The agents have fixed parameters. The different arbitrage traders make arbitrage between different couple of markets. The classical investors believe in the theory of Marcowitz. They divert in the risk and profit profile and the home currency [Ankenbrand97]. There is an equal number of the different types of agents with the different parameter and market profiles.

10.2.2 Results

The performance measurements are the same than for the neural network models. Chapter 7 describes them in detail. It measures the similarity between the real prices and the artificial prices which are calculated based on the buy and sell orders of the different agents. The *d* statistics for the different markets are [Ankenbrand97]:

	150 agents	300 agents
DEMUSD	53%	53%
YENUSD	49%	48%
Treasury Note	45%	53%
S&P 500	44%	50%
Bund	43%	53%
DAX	49%	51%
JGB	46%	53%
Nikkei 225	50%	52%

The returns are calculated in annualised percents for a better comparability of the different markets:

	150 agents	300 agents
DEMUSD	7%	4%
YENUSD	-3%	-10%
Treasury Note	-9%	1%
S&P 500	-11%	2%
Bund	-7%	0%
DAX	6%	1%
JGB	-1%	2%
Nikkei 225	-4%	-18%

The results are not significantly different from the results achieved by chance. A static simulation environment seems not to be realistic enough for meaningful experiments, if one takes into account the hard constraint of at least qualitative adherence to real financial markets behaviour. However, an improvement of the results is obtained through an increasing number of agents, but the absence of adaptive agents is not realistic and does not allow to model real financial markets.

10.3 Dynamic simulation

The lessons learned from the previous static simulations lead to the implementation of more agent and adaptive agents which is realised in this simulation.

10.3.1 Configuration

The artificial simulation environment contains 2500 agents. There are only three types of agents below the description of the previous chapter:

- Technical traders
- Evolutionary multivariate traders
- Market markers

with different adaptive genes. The genes are initialised randomly.

10.3.2 Results

The *d* statistics for the different markets are:

DEMUSD	53%
YENUSD	52%
Treasury Note	54%
S&P 500	50%
Bund	55%
DAX	52%
JGB	51%
Nikkei 225	57%

The returns are:

DEMUSD	6%
YENUSD	10%
Treasury Note	5%
S&P 500	-1%
Bund	4%

DAX	7%
JGB	-2%
Nikkei 225	6%

The results are significantly better than the previous results of an artificial market simulation with non-adaptive agents, especially as far as the trend prediction is concerned, which is important in practice. Indeed, trend prediction is now in excess of 50% in every market and the agents make profit in six out of the eight markets. The results are now significantly different from results achieved by chance.

The adaptive, evolving traders are the major change in respect to the previous simulation environment. This adaptation of the agents with evolutionary algorithms introduces a second source of non-linearity in the artificial markets. The first source of non-linearity is the traded lot size of the agents based on the accumulated profit. Both sources improve the importance of successful agents, and allows the artificial market environment to adapt and to improve the similarity to the changing reality. The evolution number or the evolution of the parameters do not converge to an equilibrium. They are changing over time depending of the market conditions. There exist not an optimal agent type with optimal parameters overall. The agent types and parameters are changing dynamically over time without any tendencies to clustering of the agents.

The adaptive capabilities of the agents and the higher number of agents seems to be responsible for the improvement of the results. This allows to model the non-linearity of the financial markets for good results with very few assumption at the beginning. There are only a general well known market structure and only three types of agents with a random initialisation of the parameters.

The often cited drawback of genetic algorithms that they are not appropriate tools for local fine tuning, is not severe in the presented application because we are using integer representation of the genes and we are not using very high precision [Michalewicz96]. The drift component of the financial markets forces an endless adaptation without a perfect global optimisation. Further, the noise component of the financial markets makes it more desirable to find an optimal region than a perfect point with a high precision in the search space.

10.4 Complex agent simulation

Complex agents based on artificial neural networks introduces a third source of non-linearity in the previous dynamic simulation environment. The neural network agents are able to autonomously modify their trading strategy as a function of their self-observed performance and of the generally available market information. This increases the diversity of the agents and is therefore closer to reality where many different trading strategies simultaneously coexist and adapt reciprocally in a endless dynamic. The complex agent simulation combines Darwanian dynamics and the neural learning theory. Darwin dynamics optimise globally the fitness of the agents through the genetic algorithms which change the genotypes. The neural learning theory optimises locally the neural network weights of the agents what expresses a change of the phenotype. The neural network is a non-linear mapping of the genotype to the phenotype which represents a second structure [Forst97]. The complex agent simulation integrates also the ideas of the microeconomic market prediction of the first part of this work.

10.4.1 Configuration

The artificial simulation environment contains 10000 agents. There are four different types of agents:

- Technical traders
- Evolutionary multivariate traders
- Artificial neural network
- Market markers

with different adaptive genes. The genes are also initialised randomly.

10.4.2 Results

The *d* statistics for the different markets are:

DEMUSD	53%
YENUSD	57%
Treasury Note	57%
S&P 500	52%
Bund	56%
DAX	54%
JGB	57%
Nikkei 225	59%

The returns are:

DEMUSD	9%
YENUSD	30%
Treasury Note	21%
S&P 500	1%
Bund	14%
DAX	2%
JGB	11%
Nikkei 225	33%

The introduction of artificial neural network traders improves the results of the previous dynamic simulation. Every market is now profitably traded. This simulation environment seems to be realistic to some extent in respect of the very good performance measurements.

The process of adaptation in the multi-agent system consist of two complementary phases. Firstly, the adaptation occurs with learning within each agent's individual lifetime. Secondly, the adaptation occurs with evolution with successive generations of the population. Between the two classic adaptation of the science of artificial life is the adaptation of the influence of an agent based on his past profit over the changing of the lot size which is practically motivated. The individual learning is more effective in static environments; while the Darwinian evolution is more stable and performs better in dynamic environments [Sasaki97]. The integration of the different theoretical and practical approaches allows a good explanation of the evolution of the real financial markets with a artificial market simulation.

10.5 Conclusion

The performance indicators, *d* statistics and returns, show a significant similarity of the artificial market with the real financial markets. The result improves with the introduction of more sources of adaptation and non-linearity. The summary of the average *d* statistics and returns of the different simulations shows very clearly this improvement:

	d statistics	return
static simulation with 150 agents	47%	-3%
static simulation with 300 agents	52%	-2%
dynamic simulation	53%	4%
complex agent simulation	56%	15%

The adaptation happens on the macro level of the financial market with the changing of the traded lot size of the agents based on the success. This adaptation process is motivated from practical experiences and observations of financial markets. The other two adaptations processes indeed are motivated from the artificial life and artificial intelligence theory. The evolution of the trading strategies is also based on the success and happens with the evolutionary algorithm. An adaptation on the micro level with the introduction of neural network trading strategies improves the results further. The improvement of the results with additional adaptation possibilities supports the presence of the drift component in the financial markets. Other studies also observe the drift component by modelling the financial markets [Burgess97, Mason97]. The combination of individual learning on micro level and collective learning on the macro level is also successfully applied to an ecological systems simulation. In this ecological study a deep connection between the intelligence of the agent behaviour and the complexity of the artificial worlds is observed. In worlds where it was very easy to survive, only dumb agents appeared. But if worlds were difficult like the financial markets, intelligent beings, like complex agents are very useful [Gracias97]. Another source of improvement is the addition of agents. A too small population size lets converge the genetic algorithm too fast [Michalewicz96] which reduces the important diversity in the artificial market. The use of cellular automata for simulation of the introduction of a new product supports the importance of using heterogeneous diversified agents and this diversification allows to model the S-shaped curve of the selling quantities of new products [Oda97].

We conclude that distributed, decentralised, interacting assemblies of adaptive agents are a valuable tool for modelling financial markets with very few assumptions at the beginning. But these few assumptions of the market structure and agents should be reasonable in respect to real financial markets and to the behaviour of the different actors in the financial world. If this condition is fulfilled, non-linear aggregation of local rules of agents is able to explain part of the complex dynamic behaviour of financial markets.

11. Discussion and Conclusions

Many phenomena which take place in the economy are poorly understood. The research programme that gave rise to this thesis was concerned with a better quantitative understanding of the evolution of the financial markets as a core of the economy. A definitive answer involving the degree of understanding gives the ability of making correct forecasts for the future behaviour of the financial markets which is the ambitious goal of this thesis. The financial markets are analysed and modelled at two different levels. Firstly statistics and artificial neural networks are used for modelling the economic behaviour of a single agent. Secondly due to the consequences of having different evolving agents for different markets, the agents interact in artificial but realistic market environment for simulating the whole market behaviour.

The financial industries are very interested in models of behaviour of financial markets. They need different kind of applications of financial markets models depending on their role in the financial markets, like investors, speculators, risk manager, regulators, and others. Investors and speculators are interested in applications of forecasting and trading which are proactive and which predict the direction of the future behaviour. There are also possible application for risk management. Risk management is reactive and more interested in the different statistics like volatility and not in the trend of the future behaviour. Another possible application especially for the artificial market simulation comes from the normative economics because the realistic simulation environment presented allows experiments which opens completely new possibilities for regulators of the financial markets. Possible experiments are the analysis of the influences of political interventions like changing taxes or other changes of the market structure. This should lead to an optimisation of the market structures and to an improvement of the social welfare. This thesis concentrates on an improvement of a better understanding of the future behaviour of the financial markets which is also manifested in the prediction driven proactive applications. The goal of the artificial market simulation is to model of a realistic environment in terms of forecasting which allows the regulators to do different experiments in a realistic model of the financial markets.

The academics have denied the possibility of beating the market for a long time. Findings of price studies of financial markets from traditional economics and finance professors are controversial and have angered many financial analysts. The markets seemed to be efficient and therefore it was impossible to gain any free lunch with the different traditional forecasting methods, but the field was still more an art than a science. The market is assumed to be efficient because it consists of a large number of rational, profit-seeking, risk-averting investors who compete freely with each other in estimating the future value of an asset. Financial time series are often cited as examples of randomness, but the statement of randomness is an empirical one. Behaviour that was previously believed random might become predictable with a better understanding of underlying dynamics, better measurements of market behaviour, or more advanced computational power to process a sufficient amount of information. The cause of unpredictability or randomness is often ignorance. If we do not know the forces that cause something to change, then we cannot predict its behaviour.

To extract useful dynamic information from time series, long time series of high quality are necessary, but the time series in economics and finance are short and noisy. Financial markets evade a quantitative description by pure deterministic model. A useful model for financial markets could also include noise or shocks and drift of the deterministic dynamics. Drift stands for a change over time in some parameters of the deterministic part of the dynamics. By introducing a number of ad hoc parameters (obtained from observation) one could get a more realistic and reliable model of financial markets than with deterministic techniques. For that reason a definition of financial markets as a dynamic deterministic system with shocks (noise), like politics and natural events for example, and drift seems to be most promising for the future work.

The working definition of financial markets as a dynamic deterministic system with stochastic shocks and drift is also supported by the results of the different forecasting tools of this thesis. The artificial neural networks in the micro-economic part are able to approximate every deterministic function. It shows its successful application with two case studies of the Swiss stock and bond market. The results of the two case studies are equal or better than similar studies [Rauscher97]. The linear and non-linear neural network models presented are stable and their performance is good. The feasibility analysis identifies the important indicators and allows to make models with good generalisation capabilities which do not seem to be over-fitted. The methodology is well suited for analysing and modelling financial markets under realistic assumptions, like very few data, of the financial markets. The only exception is the linear regression model of the SPI which has a bad out-of-sample performance, but over-optimisation is not common, because the model has few degrees of freedom, and the performance of the validation set is stable. Another possibility for the outcome of the bad out-of-sample performance is a change in the linear market structure. It seems to be a market drift; the market has evolved and is more efficient in a linear manner. A combination of the different models improves the results further, but the danger of over-fitting is very high, because the models are under-sampled and the number of degrees of freedom increases with a combination of the models.

However they are not able to catch the drift component and the noise from the external shocks. The macro-economic artificial live simulation allows to model the additional non-linearity of the drift component. This non-linearity is introduced by the trading lot size based on the success in terms of profit of the market participant and the evolution of the trading strategies with an evolutionary algorithm. An adaptation on the micro level with the introduction of neural network trading strategies improves the results further. The improvement of the results with additional adaptation possibilities supports the presence of the drift component in the financial markets. The artificial market is a valuable tool for modelling financial markets with very few assumptions at the beginning. These few assumptions of the market structure and agents should be reasonable in respect to real financial markets and to the behaviour of the different actors in the financial world. If this condition is fulfilled, non-linear aggregation of local rules of agents is able to explain a part of the complex dynamic behaviour of financial markets like the artificial markets presented show.

The neural network is a deterministic model and the artificial market is a deterministic model with drift of the real financial markets. The focus of the models lies on the explanation in term of trends which makes them to very useful tools for forecasting and trading. The different real life applications explain profitably the behaviour of the financial markets. They are also able to model rare events like crashes and other strange situations. Both of the models presented explain a lot of the evolution of the financial markets, but there is still a rest of unexplained noise. The unpredictable external shocks are a simple explanation of this noise. There could also be more regular patterns in the residual noise which are waiting for detection in the future by new yet unknown technologies.

12. Bibliography

[Ackley97]
Ackley, D.H.: ccr: A network of worlds for research. In Langton, C.G., Shimohara, K. (Ed.): Artificial Life V. The MIT Press, 1997.

[Akiro97]
Akiro, I.: How Do Selfish Agents Learn to Cooperate? In Langton, C.G., Shimohara, K. (Ed.): Artificial Life V. The MIT Press, 1997.

[Akiyama97]
Akiyama, E., Kaneko, K.: Evolution of Communication and Strategies in an Iterated Three-person Game. In Langton, C.G., Shimohara, K. (Ed.): Artificial Life V. The MIT Press, 1997.

[Allen94]
Allen, P.M.: Evolutionary Complex Systems: Models of Technology Change. In Leydesdorff, L., Van den Besselaar, P.: Evolutionary Economics and Chaos Theory. London: Pinter, 1994.

[Anders97]
Anders, U.: Statistische Neuronale Netzwerke. Munich: Vahlen, 1997.

[Ankenbrand95a]
Ankenbrand, T., Tomassini, M.: Multivariate time series modelling of financial markets with artificial neural networks. In Pearson, D.W., Steele, N.C., Albrecht, R.F.: Artificial Neural Nets and Genetic Algorithms. Springer, 1995.

[Ankenbrand95b]
Ankenbrand, T., Tomassini, M.: Neural Networks and Statistical Forecasting of Multivariate Economical and Financial Time Series. In Neural Network World. IDG, 1995.

[Ankenbrand96a]
Ankenbrand, T., Tomassini, M.: Predicting multivariate financial time series using neural networks: the Swiss bond case. In Conference on Computational Intelligence for Financial Engineering. New York: IEEE, 1996.

[Ankenbrand96b]
Ankenbrand, T., Tomassini, M.: Forecasting Financial Multivariate Time Series with Neural Networks. In International Symposium on Neuro-Fuzzy Systems. Lausanne, 1996.

[Ankenbrand97]
Ankenbrand, T., Tomassini, M.: Agent based simulation of multiple financial markets. In Neural Network World. IDG, 1997.

[Arthur95]
Arthur, B.: Complexity in economics and financial markets. Complexity, 1, 1995.

[Atiya97]
Atiya, A.: Design of Time-Variable Stop Losses and Profit Objectives using Neural Networks. In Weigend, A.S., Abu-Mostafa, Y., Refenes, A.-P.N.: Decision Technologies for Financial Engineering. World Scientific, 1997.

[Azoff94]
Azoff, E.M.: Neural network time series forecasting of financial markets. Wiley, 1994.

[Baldwin96]
Baldwin, J.: A new factor in evolution. The American Naturalist, 1896.

[Beaufils97]
Beaufils, B., Delahaya, J.-P., Mathieu, P.: Our Meeting With Gradual: A Good Strategy For The Iterated Prisoner's Dilemma. In Langton, C.G., Shimohara, K. (Ed.): Artificial Life V. The MIT Press, 1997.

[Bernstein92]
Bernstein, P.L.: Capital Ideas. New York: The Free Press, 1992.

[Beuret97]
Beuret, O., Tomassini, M.: Behaviour of Multiple Generalized Langton's Ants. In Langton, C.G., Shimohara, K. (Ed.): Artificial Life V. The MIT Press, 1997.

[Brock86]
Brock, W.A.: Distinguishing Random and Deterministic Systems: Abridged Version. In Journal of Economic Theory. Academic Press, 1986.

[Brock91]
Brock, W.A., Hsieh, D.A., LeBaron, B.: Non-linear Dynamics, Chaos, and Instability: Statistical Theory and Economic Evidence. Cambridge, London: MIT Press, 1991.

[Burgess97]
Burgess, A.N.: Asset Allocation across European Equity Indices using a Portfolio of Dynamic Cointegration Models. In Weigend, A.S., Abu-Mostafa, Y., Refenes, A.-P.N.: Decision Technologies for Financial Engineering. World Scientific, 1997.

[Chattfield89]
Chattfield, C.: The analysis of time series. London: Chapman & Hall, 1989.

[Choey97]
Choey, M., Weigend, A.S.: Nonlinear Trading Models through Sharpe Ratio Maximisation. In Weigend, A.S., Abu-Mostafa, Y., Refenes, A.-P.N.: Decision Technologies for Financial Engineering. World Scientific, 1997.

[Cohen87]
Cohen, J.B., Zinbarg, E.D., Zeikel, A.: Investment Analysis and Portfolio Management. Homewood: Irwin, 1987.

[Davis93]
Davis, D.D., Holt, C.A.: Experimental Economics. Princeton: Princeton University Press, 1993.

[Deboeck94]
Deboeck, G.J., Cader, M.: Pre- and Postprocessing of Financial Data. In Deboeck, G.J. (Ed.): Trading on the Edge. John Wiley & Sons, 1994.

[Groot92]
de Groot, C.: Non-linear time series analysis with connectionist networks. Zurich, 1992.

[Eakins98]
Eakins, S.G., Stansell, S.R., Buck, J.F.: Analyzing the Nature of Institutional Demand for Common Stocks. In Quarterly Journal of Business and Economics, 1998.

[Eckmann92]
Eckmann, J.-P., Ruelle, D.: Fundamental limitations for estimating dimensions and Lyapunov exponents in dynamical systems. Physica D, 1992.

[Efron93]
Efron, B., Tibshirani, R.J.: An introduction to the bootstrap. Chapman & Hall, 1993.

[Farmer88]
Farmer, J.D., Sidororwich, J.J.: Can New Approaches to Non-linear Modelling Improve Economic Forecast. In Anderson, P.W., Kenneth, J.A., Pines, D. (Ed.): The Economy a an evolving complex System. New York: Addison-Wesley, 1988.

[Ferber94]
Ferber, J.: Simulating with reactive agents. In Hillebrand, E., Stender, J. (Ed.): Many - agent simulation and artificial life. Amsterdam: IOS Press, 1994.

[Franke86]
Franke, J.: Grundzüge der Mikroökonomik. München: Oldenbourg, 1986.

[Friedman98]
Friedman, D.: On economic applications of evolutionary game theory. In Journal of Evolutionary Economics. Springer, 1998.

[Forst97]
Forst, C.V.: Molecular Evolution of Catalysis. In Husbands, P., Harvey, I. (Ed.): Fourth European Conference on Artificial Life. The MIT Press, 1997.

[Foss98]
Foss, N.J.: The competence-based approach: Veblenian ideas in the modern theory of the firm. In Cambridge Journal of Economics. Oxford University Press, 1998.

[Gracias97]
Gracias, N., Pereira, H., Lima, J.A., Rosa, A.: Gaia: An Artificial Life Environment for Ecological Systems Simulation. In Langton, C.G., Shimohara, K. (Ed.): Artificial Life V. The MIT Press, 1997.

[Godfrey98]
Godfrey, L.G.: Test of non-tested regression models, Some results on small sample behaviour and the bootstrap. In Journal of Econometrics. Elsevier, 1998.

[Goodwin91]
Goodwin, R.M.: Economic Evolution, Chaotic Dynamics and the Marx-Keynes-Schumpeter System. In Hodgson, G.M., Screpanti, E. (ed.): Rethinking Economics. Hants: Edward Elgar, 1991.

[Guillaume94]
Guillaume, D.M.: A low-dimensional fractal attractor in the foreign exchange markets? Working Paper, 1994.

[Hall92]
Hall, P.: The bootstrap and Edgeworth expansion. Springer, 1992.

[Hanappi94]
Hanappi, H.: Evolutionary Economics. Hants: Avebury, 1994.

[Hassoun95]
Hassoun M. H.: Fundamentals of Artificial Neural Networks. MIT Press, 1995.

[Heri96]
Heri, E.W.: Was Anleger auch noch wissen sollten .... Basel: Helbling & Lichtenhahn, 1996.

[HFDF95]
Goodhart, C.A.E., O'Hara, M.: High Frequency Data in Financial Markets: Issues and Applications. Introductory Lecture. Zurich, 1995.

[Hinton87]
Hinton, G.E., Nowlan, S.J.: How learning can guide evolution. Complex Systems, 1987.

[Hsieh91]
Hsieh, D. A.: Chaos and Non-linear Dynamics. In The Journal of Finance, 1991.

[Hodgson91]
Hodgson, G.M., Screpanti, E. (ed.): Rethinking Economics. Hants: Edward Elgar, 1991.

[Holland88]
Holland, J.H.: The Global Economy as an Adaptive Process. In Anderson, P.W., Kenneth, J.A., Pines, D. (Ed.): The Economy a an evolving complex System. New York: Addison-Wesley, 1988.

[Holland97]
Holland, O.E., Melhuish, C.R.: Getting the Most from the Least: Lessons for the Nanoscale form Minimal Mobil Agents. In Langton, C.G., Shimohara, K. (Ed.): Artificial Life V. The MIT Press, 1997.

[Kanas98]
Kanas, A.: Testing for a unit root in ERM exchange rates in the presence of structural breaks: evidence from the Bootstrap. In Applied Economics Letters. Routledge, 1998.

[Karandikar98]
Karandikar, R., Mookherjee, D., Ray, D., Vega-Redondo, F.: Evolving Aspirations and Cooperation. In Journal of Economic Theory. Academic Press, 1998.

[Kikuchi98]
Kikuchi, Y., Doi, H.: Simulating Evolution of the Vitality by the Biased Mutation Model. In Langton, C.G., Shimohara, K. (Ed.): Artificial Life V. The MIT Press, 1997.

[Klimasauskas94]
Klimasauskas, C.C.: Neural Network Techniques. In Deboeck, G.J. (Ed.): Trading on the Edge. John Wiley & Sons, 1994.

[Langton89]
Langton, C.G.: Artificial life. Adison-Wesley, 1989.

[Langton97]
Langton, C.G., Shimohara, K. (Ed.): Artificial Life V. The MIT Press, 1997.

[Lee93]
Lee, T.-H., White, H., Granger, W.J.: Testing for neglected nonlinearity in time series models. In Journal of Econometrics. Elsevier Science Publishers, 1993.

[Li98]
Li, Q., Wang, S.: A simple consistent bootstrap test for a parametric regression function. In Journal of Econometrics. Elsevier, 1998.

[Loomes98]
Loomes, G.: Probabilities vs Money: A test of some fundamental assumptions about rational decision making. In The Economic Journal. Blackwell Publishers, 1998.

[Lux98]
Lux, T.: The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fait tails of return distributions. In Journal of Economic Behaviour & Organization. Elsevier, 1998.

[Marquez92]
Marquez, L., Hill, T., Worthley, R., Remus, W.: Neural Network Models as an Alternative to Regression. In Trippi, R.R., Turban, E. (Ed.): Neural Networks in Finance and Investing. Chicago, Cambridge: Probus, 1992.

[Martens98]
Martens, M., Kofman, P.: The inefficiency of Reuters foreign exchange quotes. In Journal of Banking & Finance. Elsevier, 1998.

[Mason97]
Mason, J.M., Badii, A.: Neural Network Compatible Memory with Controlled Depth and Resolution for Financial Temporal Forecasting Data. In Weigend, A.S., Abu-Mostafa, Y., Refenes, A.-P.N.: Decision Technologies for Financial Engineering. World Scientific, 1997.

[Mayley97]
Mayley, G.: Guiding to Hiding: Explorations into the Effects of Learning on the Rate of Evolution. In Husbands, P., Harvey, I. (Ed.): Fourth European Conference on Artificial Life. The MIT Press, 1997.

[Maza97]
De la Maza, M.: A Computer Simulation of Currency Market Participants. In Weigend, A.S., Abu-Mostafa, Y., Refenes, A.-P.N.: Decision Technologies for Financial Engineering. World Scientific, 1997.

[McClelland89]
McClelland, J.L., Rumelhart, D.E.: Explorations in Parallel Distributed Processing. Cambridge, London: The MIT Press, 1989.

[Michalewicz96]
Michalewicz, Z.: Genetic Algorithms + Data Structures = Evolution Programs. Heidelberg: Springer, 1996.

[Miller94]
Miller, M.: Das Optimieren von Neuronalen Netzen für den Einsatz zur Prognosen in der Ökonomie. In Bol, G., Nakhaeizadeh, G., Vollmer, K.-H. (Ed.): Finanzmarktanwendungen neuronaler Netze und ökonometrischer Verfahren. Heidelberg: Physica, 1994.

[Minsky97]
Minsky, M.: How Computer Science will change our lives. In Langton, C.G., Shimohara, K. (Ed.): Artificial Life V. The MIT Press, 1997.

[Miura97]
Miura, H., Yasuda, T., Fujisawa, Y.K., Kuwana, Y., Takeuchi, S., Shimoyama, I.: Insect-Model Based Microrobot. In Langton, C.G., Shimohara, K. (Ed.): Artificial Life V. The MIT Press, 1997.

[Monostori97]
Monostori, L., Hornyak, J., Egresits, Cs.: ANN and Hybrid AI Approaches to Financial and Business Problems. In Neural Network World. IDG, 1997.

[Moss92]
Moss, S., Rae, J. (Ed.): Artificial Intelligence and Economic Analysis. Edward Elgar, 1992.

[Nakamura97]
Nakamura, M., Kurumatani, K.: Formation Mechanism of Pheromone Pattern and Control of Foraging Behaviour in an Ant Colony Model. In Langton, C.G., Shimohara, K. (Ed.): Artificial Life V. The MIT Press, 1997.

[Neely97]
Neely, C., Weller, P., Dittmar, R.: Is Technical Analysis in the Foreign Exchange Market Profitable? A Genetic Programming Approach. In Journal of Financial and Quantitative Analysis, 1997.

[Oda97]
Oda, S.H., Miura, K., Ueda, K.: The Application of Cellular Automata to Network Externalities in Consumer's Theory: A Generalisation of Life Game. In Langton, C.G., Shimohara, K. (Ed.): Artificial Life V. The MIT Press, 1997.

[Olsen92]
Olsen, R.B., Dacorogna, M.M., Müller, U.A., Pictet, O.V.: Going Back to the Basics - Rethinking Market Efficiency. A discussion paper of the O&A Research Group, 1992.

[Palmer94]
Palmer, R.G., Arthur, W.B., Holland, J.H., Le Baron, B., Tayler, P.: Artificial economic life: a simple model of the stockmarket. Physica D, 75, 1994.

[Pardo92]
Pardo, R.: Design, Testing, and Optimization of Trading Systems. John Wiley & Sons, 1992.

[Peters91]
Peters, E.E.: Chaos and Order in the Capital Markets. John Wiley & Sons, 1991.

[Peters94]
Peters, E.E.: Fractal Market Analysis. John Wiley & Sons, 1994.

[Phang94]
Allen, P.M., Phang, H.K.: Managing Uncertainty in Complex Systems: Financial Markets. In Leydesdorff, L., Van den Besselaar, P.: Evolutionary Economics and Chaos Theory. London: Pinter, 1994.

[Pineda94]
Pineda, F.J., Sommerer, J.C.: Estimating Generalized Dimensions and Choosing Time Delays: A Fast Algorithm. In Weigend, A.S., Gershenfeld, N.A. (Ed.): Time Series Prediction. Addison Wesley, 1994.

[Poddig94]
Poddig, T., Rehkugler, H., Jandura, D.: Ein Weltmodell integrierter Finanzmärkte. In Rehkugler, H., Zimmermann, H.G. (Ed.): Neuronale Netze in der Ökonomie. München: Vahlen, 1994

[Pollack97]
Pollack, J.B., Blair, A.D., Land, M.: Coevolution of a Backgammon Player. In Langton, C.G., Shimohara, K. (Ed.): Artificial Life V. The MIT Press, 1997.

[Rauscher97]
Rauscher, F.A.: Exchange rate forecasting: A neural VEC approach. In Neural Network World. IDG, 1997.

[Refenes95]
Refenes A. P. (Ed.): Neural Networks in the Capital Markets. John Wiley & Sons, 1995.

[Refenes97]
Refenes, A.-P., N., Zapranis, A.D., Utans, J.: Neural Model Identification, Variables Selection and Model Adequacy. In Weigend, A.S., Abu-Mostafa, Y., Refenes, A.-P.N.: Decision Technologies for Financial Engineering. World Scientific, 1997.

[Rehkugler94]
Rehkugler, H., Poddig, T.: Kurzfristige Wechselkursprognosen mit Künstlichen Neuronalen Netzwerken. In Bol, G., Nakhaeizadeh, G., Vollmer, K.-H. (Ed.): Finanzmarktanwednungen neuronaler Netze und ökonometrischer Verfahren. Heidelberg: Physica, 1994.

[Ruelle88]
Ruelle, D.: Can Non-linear Dynamics Help Econimists? In Anderson, P.W., Kenneth, J.A., Pines, D. (Ed.): The Economy a an evolving complex System. New York: Addison-Wesley, 1988.

[Ruelle90]
Ruelle, D.: Deterministic chaos: the science and the fiction. In Proc. R. Soc. London, 1990.

[Ruelle94]
Ruelle, D.: Where can one hope to profitably apply the ideas of chaos. Physiscs Today, 1994.

[Russell95]
Russell, S., Norvig, P.: Artificial Intelligence. Prentice-Hall, 1995.

[Samuelson89]
Samuelson, P.A., Nordhaus, W.D.: Economics. New York: McGraw-Hill, 1989.

[Sasaki97]
Sasaki, T., Tokoro, M.: Adaptation toward Changing Environments: Why Darwinian in Nature? In Husbands, P., Harvey, I. (Ed.): Fourth European Conference on Artificial Life. The MIT Press, 1997.

[Scheinkman89]
Scheinkman, J.A., Le Baron, B.: Non-linear dynamics and stock returns. Journal of Business, 1989.

[Schlittgen87]
Schlittgen, R., Streitberg, B.H.J.: Zeitreihenanalyse. München, Wien: R. Oldenburg, 1987.

[Schnidrig97]
Schnidrig, R., Würtz, D., Hanf, M.: A Realtime Computer Simulation of a Foreign Exchange Trading Room. In Neural Network World. IDG, 1997.

[Selten91]
Selten, R.: Evolution, learning, and economic behaviour. Games and economic behaviour, 1991.

[Sen95]
Sen, T.K., Oliver, R., Sen, N.: Predicting Corporate Mergers. In Refenes, A.P. (Ed.): Neural networks in the Capital Markets. John Wiley & Sons, 1995.

[Shapiro92]
Shapiro, S.C.: Encyclopedia of artificial intelligence. John Wiley & Sons, 1992.

[Sharda92]
Sharda, R., Patil, R.B.: A Connectionist Approach to Time Series Predicition: An Empirical Test. In Trippi, R.R., Turban, E. (Ed.): Neural Networks in Finance and Investing. Chicago, Cambridge: Probus, 1992.

[Shin98]
Shin, H.-H., Stulz, R.M.: Are International Capital Markets Efficient? In The Quarterly Journal of Economics. The MIT Press, 1998.

[Shinjoh97]
Shinjoh, A., Takeyama, M.: Spatial Analysis of Artificial World. In Langton, C.G., Shimohara, K. (Ed.): Artificial Life V. The MIT Press, 1997.

[Simon82]
Simon, H.: Models of Bounded Rationality. Cambridge: MIT Press, 1982.

[Steurer97]
Steurer, E.: Ökonometrische Methoden und maschinelle Lernverfahren zur Wechselkursprognose - theoretische Analyse und empirischer Vergleich. Heidelberg: Physica, 1997.

[SwissIndex]
SwissIndex: Die Indexfamilie der Schweizer Börsen. Association Tripartite Bourses, Zürich.

[Takens81]
Takens, F.: Detecting Strange Attractors in Turbulence. In Rand, D.A., Young, L.-S. (Ed.): Dynamical Systems and Turbulence. Berlin: Springer, 1981.

[Tesfatsion97]
Tesfatsion, L.: How economist can get alive. To appear in Arthur, B., Durlauf, S., Lane, D.: The economy as an evolving complex system II. Adison-Wesley, 1997.

[Tordjman97]
Tordjman, H.: Spéculation, hétérogénéité des agents et apprentissage. In Revue Economique. Presse de Sciences Po, 1997.

[Türker96]
Türker, C., Conrad, S., Saake, G.: Dynamically Changing Behaviour: An Agent-Oriented View to Modeling Intelligent Information Systems. In Ras, Z.W., Michalewicz, M. (Ed.): Foundations of Intelligent Systems. Springer, 1996.

[Veblen98]
Veblen, T.B.: Why is economics not an evolutionary science? In Quarterly Journal of Economics, 1898.

[Wakeley98]
Wakeley, T.M.: Schumpeterian process competition, welfare and laissez-faire: An experiment in artificial economic world. In Journal of Evolutionary Economics. Springer, 1998.

[Walter51]
Walter, W.G.: A Machine that Learns. Scientific American, 1951.

[Weller97]
Weller, K.: Optimization Criteria for Design of Serial Transfer Evolution Experiments. In Husbands, P., Harvey, I. (Ed.): Fourth European Conference on Artificial Life. The MIT Press, 1997.

[Weigend94]
Weigend, A.S., Gershenfeld, N.A. (Ed.): Time Series Prediction. Addison Wesley, 1994.

[Werbos94]
Werbos, P.J.: The roots of backpropagation. John Wiley & Sons, 1994.

[Williams98]
Williams, G., Parikh, A., Bailey, D.: Are exchange rates determined by macroeconomic factors. In Applied Economics. Routledge, 1998.

[Yokoi97]
Yokoi, H., Mizuno, T., Takita, M., Hakura, J., Kakazu, Y.: Amoeba Like Self-Organization Model using Vibrating Potential Field. In Langton, C.G., Shimohara, K. (Ed.): Artificial Life V. The MIT Press, 1997.

[Zimmermann94]
Zimmermann, H.G.: Neuronale Netze als Entscheidungskalkül. In Rehkugler, H., Zimmermann, H.G. (Ed.): Neuronale Netze in der Ökonomie. München: Vahlen, 1994.

[Zimmermann96]
Gehrig, B., Zimmermann, H.: Fit for Finance. Zürich: NZZ, 1996.

13. Index