The Efficient Market Hypothesis through the Eyes of an Artificial Technical Analyst: An Application of a New Chartist Methodology to High-Frequency Stock Market Data

Timur Yusupov*
Department of Economics,
University of Kiel,
Olshausen Str. 40, Kiel,
24118, Germany,
yusupov@bwl.uni-kiel.de

Thomas Lux
Department of Economics,
University of Kiel,
Olshausen Str. 40, Kiel,
24118, Germany

Abstract

The academic literature has been reluctant to accept technical analysis as a rational strategy of traders in financial markets. In practice traders and analysts heavily use technical analysis to make investment decisions. To resolve this incongruence the aim of this study is to translate technical analysis into a rigorous formal framework and to investigate its potential failure or success. To avoid subjectivism we design an Artificial Technical Analyst. The empirical study presents the evidence of past market inefficiencies observed on the Tokyo Stock Exchange. The market can be perceived as inefficient if the technical analyst's transaction costs are below the break-even level derived from technical analysis.

1. Introduction

For long time the efficient market hypothesis (EMH) has been the dominant paradigm finance. Its weak form postulates that in a competitive market it should not be profitable to base investment decisions on information obtained from past prices or returns of publicly traded securities. Numerous empirical studies, however, show that technical analysis, which directly contradicts the weak form of the EMH, could exploit to some extent hidden patterns in past prices. To avoid the joint hypothesis problem of direct tests of the EMH, an artificial technical analyst is created to conduct the test. This approach has two advantages. First, it is free of equilibrium model limitations, and second, technical analysis can be tested in a robust way, which should validate its existence.

The EMH is the cornerstone of modern financial economics. The paradigm was coined in the 1960-70s by Harry Roberts [133] and formalized by Eugene Fama [50]. They identified three forms of market efficiency distinguished by which information prices of securities should correctly incorporate. The weak form of market efficiency postulated that past prices or returns should have no information, which can be used to predict next period values. This form was linked to the random walk hypothesis, which constituted the majority of tests performed that time. The weak form EMH was supported by empirical studies conducted before and shortly after the 1970s. Its association to the random walk allowed the development of many important analytical tools in financial theory. The most famous example is the application of the random walk hypothesis by Myron Black and Fischer Scholes [21] to derive their seminal option pricing formula, which has caused a boom of derivative markets and further developments in financial theory.

With the development of statistical techniques more and more deviations from the random walk hypothesis were observed in time series of prices. Finally, Andrew Lo and Craig MacKinlay [100] used a simple specification test to reject the random walk hypothesis for stock prices and returns. The test was based on the variance properties of random walk time series. It is robust to different heteroskedasticities and non-normality of data. In reaction to this the assumptions of the random walk hypothesis were relaxed, first, to allow only independence of increments, and, later, to require only zero correlation of increments.

In the 1980s technical analysis appeared as a new aspect in the empirical literature on testing the EMH. This approach, reported to be wide-spread among professional financial practitioners, attempts to exploit predictability of prices for profit, and thus is in direct contradiction to the weak form of market efficiency. For example, in 1989 He-

---

*Corresponding author.

len Allen and Mark Taylor [5] surveyed that at the shortest horizons, intraday to one week, approximately 90% of respondents [professional traders and analysts at foreign exchange markets] used some chartist input in forming their exchange rate expectations, with 60% judging charts to be at least as important as fundamentals. Moreover, there appeared to be a persistent 2% of presumably "pure" chartists, who only used technical analysis. Numerous empirical studies reported either profitability of technical analysis or its positive added value to investment decision making.

Along with the empirical studies financial theory turned its attention to technical analysis. This paradigm was incorporated in behavioral asset pricing models to capture observed empirical "anomalies", like volatility clustering, high transaction volumes and erratic behavior of prices [17]. It was persuasively shown that technical analysis could bring additional information, and as such should be used by rational investors [22, 27]. Moreover, the application of technical analysis for speculation can produce stable equilibria in an economy [113]. In the extreme case, the dominance of an "irrational" technique, which technical analysis was referred to be, in the market can create its own space, where application of other "rational" techniques is suboptimal [45]. At present technical analysis and market efficiency are studied from a behavioral prospective, which promises to resolve their incongruence.

Since the original formulation of the weak form of the efficient market hypothesis became outdated, the empirical implementation of the EMH experienced several transformations. First, unpredictability of prices was replaced by inability to outperform passive benchmarks, such as the buy-and-hold strategy, and later, by adding the aspect of profitability. In this formulation a financial market is weakly efficient if an outcome of market interactions does not contain any information, which can be persistently and profitably exploited by predicting the next period prices or returns. The notion of profitability is taken in a strict sense, that is all transaction costs should be accounted for.

With the first formulation of the EMH, it became obvious that its tests might be sensitive to the joint hypothesis problem. Evidence against the EMH could be either due to a wrong equilibrium model or due to true market inefficiency. One way to avoid this problem is to construct tests that do not assume any underlying model.

Technical analysis stands in direct contradiction to the weak form of the EMH, and as such can be directly used for its testing. The design of the test is as follows. One should process past stock prices with technical analysis to obtain next period price predictions. Predicted values should be used in hypothetical investment decisions. Associated returns, adjusted for transaction costs, should be aggregated and can be used as a measure of financial market inefficiency if they are in excess of a passive benchmark.

Normally the application of technical analysis suffers from subjectivism. Technical analysis is taken rather as an art than a precise science. The application of principles of technical analysis in an autonomous or artificial decision-making system should eliminate the subjective factor. Additionally, the system should be relatively simple to insure its robustness, and transparent to provide its understandability. Spyros Skouras [143] proposed an artificial technical analyst as a quantifiable measure of market efficiency. Expanding his idea, an artificial technical analyst (ATA) will also be used for our test.

Our idea of creating and employing the ATA originates from the seminal paper by Arthur et al. [9]. The authors have created the Artificial Stock Market (ASM), known as the Santa Fe ASM. The market is populated by artificial agents engaged in stock trading. The agents use technical analysis to screen the market and an implementation of artificial intelligence, Learning Classifier Systems (LCS), for the optimal application of technical trading rules.

The success of the Santa Fe ASM inspired us to create the ATA. Unfortunately, the Santa Fe ASM had some shortcomings: computational and algorithmic limitations of that time, and a short and fixed list of technical trading rules. Additionally, the mechanism of forming the traders' expectations was based on early concepts of LCS that have undergone important modifications in the meantime. Our implementation takes into account the shortcomings of the Santa Fe ASM. It incorporates technical analysis in an adaptive way, where the core of the ATA is driven by a new implementation of LCS.

This new implementation of the ATA incorporates three main components: data-preprocessing, pattern-recognition and decision-making under transaction costs. The first component insures that raw input data are homogenized in a way that maximizes the informational content. Homogenization itself allows to reduce the complexity of pattern-recognition, by focusing attention only on levels in time series. The pattern-recognition is driven by an implementation of artificial intelligence, which allows for a transparent structure of the results. LCS was selected as a possible candidate. To match the required level of performance a new algorithm is used - denoted as True Classifier System (TiCS). The decision-making under transaction costs insures optimality of investment decisions.

In the empirical part the ATA is applied to historical time series of security prices. Since the part of economically relevant transaction costs is the impact on the price, in a study of hypothetical investments with historical data the aspect of profitability is inverted. Instead of measuring the aggregate return adjusted for specified transaction costs, the value of *revealed* transaction costs, which insure positive aggregate return, is used. This allows to derive breakeven transaction costs, which could be later used for benchmarking

markets, segments, and individual stocks.

This work differs from others in the extent to which the design of the market efficiency test is free of equilibrium model limitations. By applying technical analysis one can see whether price time series have predictable patterns. In case persistent patterns are detected, their profitability can be studied under the schedule of revealed transaction costs. In this way an absence of patterns is a clear indication of the weak form of market efficiency, which was advocated in the theoretical literature of the 1970s. Otherwise, the profitability of detected patterns is accessed through revealed transaction costs. When the level of breakeven transaction costs is below market transaction costs, a market can be perceived as efficient, according to the pertinent literature. At the same time a comparison to market transaction costs leaves a speculative space for those market participants, whose correctly accounted transaction costs are below the breakeven value.

The novelty of this work comes, first, from using an adaptive method of technical analysis, which allows to extract uncharted patterns; second, from developing a new type of pattern-recognition engine, which makes detected patterns accessible afterwards; and, third, from testing market efficiency under a wide schedule of transaction costs, which allows to identify boundaries of market efficiency.

This paper is organized as follows. Section 2 presents an overview of the efficient market hypothesis. It is followed by Section 3, which discusses the impact of transaction costs on the perception of market efficiency. Section 4 outlines the counterpart of the weak form of market efficiency - technical analysis. Section 5 introduces the design and key components of the ATA. Data used in empirical studies are presented in Section 6 along with a short description of the Tokyo Stock Exchange microstructure. Results of hypothetical empirical application are presented in Section 8. Conclusion finalizes the paper.

2. Efficient Market Hypothesis

The origin of the EMH dates back to 1900, when Louis Bachelier [12] first introduced the idea that the stock market fluctuations follow a stochastic process for which the future does not depend on the past except through the present and the best prediction of the subsequent price is the value of the current price.[1] That is if all relevant information is already contained in the quoted prices, the only cause of new variations could be elements of information that are not predictable. Bachelier already stresses the importance of the information concept, on which the efficiency concept

is based.[2]

The notion of "efficient markets" was coined by Harry Roberts and popularized by Eugene Fama. In his Ph.D. dissertation Fama convincingly made the argument that in an active market that includes many well-informed and intelligent investors, securities will be appropriately priced and will reflect all available information. If a market is efficient, no information or analysis can be expected to result in outperformance of an appropriate benchmark.

It has been customary since Harry Roberts [133] to distinguish three forms of market efficiency by considering three different types of information sets:

- *The weak form* of the EMH asserts that prices fully reflect the information contained in the historical sequence of prices. Thus, investors cannot devise an investment strategy to yield abnormal profits on the basis of an analysis of past price patterns.[3]

- *The semi-strong form* of the EMH asserts that current stock prices reflect not only historical price information but also all publicly available information relevant for company securities. If markets are efficient in this sense, then an analysis of balance sheets, income statements, announcements of dividend changes or stock splits or any other public information about a company will not yield abnormal economic profits.

- *The strong form* of the EMH asserts that all information that is known to any market participant about a company is fully reflected in market prices. Hence, not even those with privileged information can make use of it to secure superior investment results. There is perfect revelation of all private information in market prices [105].

These definitions of the variants of the EMH allows to construct tests focused on their specific aspects. Acceptance or rejection of some form of market efficiency allows to identify the level of financial market development. Obtained results can be used to generate recommendations and policies to improve efficiency of financial markets.

Theoretical arguments for the efficient market hypothesis are based on three assumptions. These are the rationality of investors, the irrationality of investors and the randomness of trades and, finally, the presence of rational arbitrageurs. Each assumption is progressively weaker but their

---

[1]This notion was later branded by Karl Pearson as the random walk hypothesis [125, 126].

[2]However, as Paul Cootner [36] comments - Bachelier's work received little attention from academicians and was forgotten for almost fifty five years. His work was rediscovered in 1955 by Leonard Jimmie Savage.

[3]The definition of Mark Rubinstein in [137] and William Beaver in [16] requires that publishing the information does not change equilibrium prices. Rubinstein's notion allows one to ask only if the market is efficient with respect to all information. William Beaver provides a definition that covers the information contained in historical prices.

combination allows to justify market efficiency in the most of market situations.

Rationality of investors. Investors are assumed to be rational and hence to value securities rationally. Investors value each security for its fundamental value. When investors learn something new about the fundamental values of securities, they quickly respond to the new information by bidding up prices (buying securities) when the news is good and bidding them down (selling) when the news is bad. As a result, security prices incorporate all the available information almost immediately and adjust to new levels corresponding to the new net present values of expected cash flows.

Irrationality of investors and randomness of trades. It is often admitted by proponents of the EMH that some investors are not rational, and they are trading randomly. When there is a large number of such investors and when their trading strategies are uncorrelated, their trades are likely to neutralize each other.[4]

Presence of rational arbitrageurs. Although some investors might be irrational in a similar way there are some rational arbitrageurs in the market. Arbitrage[5] is one of the most intuitively appealing and plausible arguments in all of economics. The main condition for arbitrage is the existence of an over- or undervalued security with a close substitute.[6] This could be the case when the trade involves irrational investors. Noting the overpricing the arbitrageur would sell (or even short sell) the overpriced security and simultaneously purchase another, "essentially similar" but truly valued, security to hedge any risk. The effect of this arbitrage is to bring the price of the overpriced security down to its fundamental value.

The arbitrage argument allows to cover the most complex case - the existence of irrational investors. To the extent that the securities that irrational investors are buying are overpriced and the securities they are getting rid of are undervalued, such investors earn lower returns than either passive investors or arbitrageurs. Relative to their peers, irrational investors lose money and in a long run should leave the market. Thus, not only investor rationality, but also the market forces bring about the efficiency of financial markets [142].

Since this work is presenting the efficient market hypothesis from the point of view of technical analysts this study focuses on the weak form of efficiency and its tests.

The earliest tests were concerned with short horizon returns. These tests typically assumed that in an efficient market the expected rate of return was constant through time and the realized returns should not be serially correlated. Eugene Fama [49] finds that the first-order autocorrelation of daily returns is positive. In [106] and [49] it is also recognized that returns are characterized by volatility clustering and leptokurtic unconditional distributions. Lawrence Fisher [60] suggests that autocorrelations of monthly returns of a diversified portfolio are bigger than those of individual stocks. However, the evidence often lacked statistical power and the EMH was not rejected.[7]

Later research used daily and weekly NYSE or AMEX data. Andrew Lo and Craig MacKinlay [100] find that weekly returns on portfolios of NYSE stocks show reliable positive autocorrelation, which is stronger for portfolios of small stocks. This can be due to their smaller liquidity and the non-synchronous trading effect discussed already in [60]. In [35] Jennifer Conrad and Gautam Kaul mitigated this problem, examining the autocorrelations of Wednesday-to-Wednesday returns for size-grouped portfolios. They also found positive autocorrelation especially in portfolios of small stocks.

However, in [79] Shmuel Kandel and Robert Stambaugh show that stock return predictability, which seems weak when evaluated by classical statistical criteria, may nevertheless be economically important in the sense that a rational Bayesian investors would substantially alter portfolio holdings in response to the current values of predictive variables.

As noted by Francis Diebold [46] and Robert Cumby and John Huizinga [38], the presence of conditional heteroskedasticity or excess kurtosis biases the test towards rejection of the null hypothesis of uncorrelated returns. In contrast to the weak evidence for autocorrelation in returns, Tim Bollerslev and Robert Hodrick [24] stress the importance of conditional heteroskedasticity. One should, however, say that finding a strong dependence in the even ordered moments does not necessarily imply market inefficiency, which is consistent with a martingale hypothesis for stock prices.

Overall, many researches showed that daily and weekly returns are predictable from past returns. Thus, these empirical findings reject the weak form of the EMH. At the same time the estimated autocorrelations are typically found to be very small and the variation of these returns is a small part of the overall return variance.

In recent times the return predictability research gradually moved in the direction of higher frequency time series. In [80, 81] Ludwig Kanzler developed a new version of the BDS test to verify the EMH on ultra-high frequency foreign exchange market data. He found that the EMH holds only in some periods, in particular, when a release of important news takes place.

---

[4]In such a market, there will be a substantial trading volume as the irrational investors exchange securities with each other, but the prices are nonetheless close to their fundamental values.

[5]That is the simultaneous purchase and sale of the same, or essentially similar, security in two different markets at advantageous different prices.

[6]In some cases the access to alternative markets can be taken as a form of substitute.

[7]Reported R^2 was often less than 0.01 for individual stocks.

The more striking evidence on the predictability of returns from past returns comes from tests on predictability of long-horizon returns. Robert Shiller [141] and Lawrence Summers [145] challenge the argument of EMH validity, based on very low autocorrelations of short-horizon returns. They provide evidence of stock market inefficiency by demonstrating that stock prices undergo large slowly decaying swings, even though the short-term returns have little autocorrelation.

In [42, 43] Werner DeBondt and Richard Thaler attacked market efficiency in a similar manner, trying to unmask irrational bubbles. They find that the NYSE stocks identified as the most extreme losers over a 3- to 5-year period tend to have strong returns relative to the market during the following years. The stocks identified as extreme winners tend to have, on the contrary, weak returns relative to the market. They attribute these results to market overreaction to extreme news.

In [77] Narasimhan Jegadeesh and Sheridan Titman observed that past winners realized consistently higher returns around their earnings announcements in the first 7 months following the portfolio formation date than past losers. They argue that to attribute the results to underreaction is overly simplistic. Buying past winners and selling past losers, consistent with positive feedback trading, moves prices further from their long-run values and thereby causes price to overreact. The interpretation is consistent with DeLong et al. [44] who explore the implications of positive feedback trading on market prices. Louis Chan [32] and Ray Ball and S. P. Kothari [13] argue that these results are due to a failure to risk-adjust returns.

James Poterba and Lawrence Summers [127] and Eugene Fama and Kenneth French [54, 55] also realized that the negative serial correlation in returns would manifest more transparently at longer horizons. Evidence in [127] and [54], using multi-period regressions and variance ratio statistics, suggests that for longer return horizons a large proportion of returns is explainable from the history of past returns alone.[8] James Poterba and Lawrence Summers [127] argue that asset prices are characterized by speculative fads, in which market prices experience long systematic swings away from their rational fundamental values.

However, whether the longer-horizon mean reversion really exists is controversial. For example, Narasimhan Jegadeesh [76], Kim et al. [84], Mankiw et al. [107], Richardson et al. [130], all argue that the case for predictability of long-horizon stock returns is weak when one corrects for the small sample biases in test statistics. In addition, [54] offers the counterargument that irrational bubbles and

swings in stock prices are indistinguishable from rational time-varying expected returns.

The subsequent work showed, that the apparent predictability of long-horizon returns should be interpreted very carefully. As [24] point out, the overlapping nature of the data in the multi-year return regressions gives rise to a non-standard small sample distribution of test statistics, which appear to be better approximated by the alternative asymptotic distribution derived by Richardson et al. [130].

In [24] Tim Bollerslev and Robert Hodrick developed tests based on the iterated version of the null hypothesis using Hansen's GMM [68] and found some improvement in the small sample performance of the test statistics. However, they conclude that there is still little evidence for predictability of returns.

To sum up, the degree of predictability is generally small compared to the high variability of returns. In [52] Eugene Fama supports this argument saying that market anomalies are chance events, i.e. they split randomly between overreaction and underreaction to news (see for example [2, 10, 111, 136]) and they tend to depend on the methodology used.

In the 1970s researchers interested in the efficiency of asset markets shifted their focus from the predictability of returns to the volatility of prices. The main reason was that price fluctuations seemed to be too large to be justified by the subsequent variation in dividend payments. The EMH could not be tested directly but only as a part of a joint hypothesis. Researchers were still required to specify a particular model of expected returns. In addition, the predictions of price volatility depended on the assumed time series properties of the dividend process and the information sets of economic agents [24].

Stephen LeRoy and Richard Porter [98] and Robert Shiller [139, 140] introduced another important class of tests for market efficiency: the volatility or variance bounds tests. They assumed a constant expected rate of return model and reported overwhelming rejections of market efficiency since excess price volatility was supposed to imply market inefficiency.

In the first generation of volatility tests the null hypothesis was taken to be the standard present value model with a constant discount rate. The vast majority of these tests resulted in clear rejections of market efficiency, with actual asset prices being excessively volatile compared to the implied price series calculated from the discounted value of the expected or actual future fundamentals. One possible explanations was the idea that asset prices may be characterized by self-fulfilling speculative bubbles that earn the fair rate of return but cause prices to differ from their rational fundamentals.

However, as Charles Nelson [120] and Eugene Fama and William Schwert [56] showed in their works, the assump-

---

[8]The estimated in [54] for monthly U.S. stock returns imply that for 3- to 5-year returns up to 40% of variability is predictable. However, this does not necessarily imply market inefficiency, since the variation could be due to a time-varying risk premium.

tion of constant expected return was unjustified. In response to this problem, subsequent research, in particular by Marjorie Flavin [61], Allan Kleidon [85, 86] and Eugene Fama [51] questioned the small sample statistical properties of these analyses.[9]

The volatility tests thus clearly show that expected returns vary through time, but give no help on the central issue of whether the variation in returns is rational.[10]

The introduction of the option pricing theory by Myron Black and Fischer Scholes [21] in the 1970s and 1980s turned the attention of the financial community to derivative markets. A. L. Tucker [156] studied the currency option markets. After accounting for transaction costs and bid-ask spreads he reported no possibility to earn riskless arbitrage profits on the currency options market. James Bodurtha and Georges Courtadon [23] achieved similar results. The currency option market was inefficient only before adjusting for transaction costs. However, the violations of the EMH disappear when transaction costs are taken into account.

Y. P. Chung [33] investigated the efficiency of the market for stock index futures and the profitability of index arbitrage. The results indicate that the size and frequency of boundary violations decreased significantly over the sample years for all levels of transaction costs, which indicates maturing of the futures market in which arbitrage trading has tended to correct mispricing.

In the process of testing the EMH a fundamental problem became obvious. It is known as the joint hypothesis test of the EMH. The problem comes from underlying probabilistic assumptions. One cannot speak of efficiency by itself, except through a model that defines the generation of prices with a representative probability system. Eugene Fama [50, 51] stressed that the market efficiency per se is not testable. One can test whether information is properly reflected in prices in the context of a pricing model. It means that when one finds anomalous evidence on the behavior of prices or returns it is ambiguous if this is caused by market inefficiency or/and a bad model of market equilibrium. This leads to the conclusion that the efficiency tests are always joint tests on the market efficiency and the pricing model and its probabilistic assumptions.

In [157] Christian Walter considers this overlap as a common cause of misinterpretations and errors, leading to rejection of efficiency when there is only a misspecification of the stochastic process.

3. Transaction Costs

At the end of the 1970's and beginning of the 1980's an increasing number of studies reported the violation of the weak form of market efficiency. In-depth analyses showed that the EMH does not rule out small abnormal returns before accounting for transaction costs. Given that collecting and processing information is a costly process, prices are expected to reflect information to the point where the marginal benefits of acting on information do not exceed the marginal costs [67, 78].

Steven Thorley [154] defines four components of transaction costs: *brokerage commissions*, *bid-ask spreads*, *taxes*, *suboptimal diversification and research*.[11] Based on the performed simulations he states that 67% of portfolios are underperforming when transaction costs are correctly taken into account. Thus, bravado reports of practitioners beating the market could be just a result of myopic accounting, but not of market inefficiency.

John Hussman [75] argues that transaction costs create a region in which the market may be inefficient while still excluding the possibility of abnormal risk-adjusted returns. If sufficiently high trading costs reduce long-term returns below those of a passive approach, an active approach may still be optimal from the standpoint of utility maximization for "myopic" investors whose utility is defined over the sequence of returns during individual holding periods, instead of terminal wealth.

Analysts could therefore still have an incentive to obtain and act on valuable information. As Elroy Dimson and Massoud Mussavian [47] suggest, time-varying expected returns could also explain these patterns.

To sum up, transaction costs, in particular bid-ask spreads, are one of the main reasons for rejecting the EMH. That is the stock market is efficient when transaction costs are considered [11]. Correct accounting for transaction costs can remove perception of market inefficiency. Thus, the magnitude of transaction costs is crucial for measuring market efficiency.

4. Technical Analysis

There are two competing ways to forecast the price development of financial instruments: fundamental and technical analysis. The fundamental analysis relies on the fundamental attributes of the instrument, such as price/earning ratio, return on investment and associated economic statistics. The aggregation of these measures provides an intrinsic value of the instrument, which in an efficient financial market should be equal to the trading price of the instrument. Unfortunately, this is not the case observed in reality.

The existence of a human factor brings distortions causing deviations of the trading price from its intrinsic value. Technical analysis is aimed at detecting a psychological component of financial trading and consequently converting findings into profit.

Technical analysis is the practice of identifying recurring patterns in historical prices in order to forecast future price trends.[12] The technique relies on the idea that, as Martin Pring [129] puts it - prices move in trends which are determined by the changing attitudes of investors toward a variety of economic, monetary, political and psychological forces. Detection of trends is performed through indicators or technical rules which are aimed to capture underlying dependencies.

The Japanese were the first to use technical analysis to trade rice on the Dojima Rice Exchange in Osaka as early as the 1600s. A Japanese man called Munehisa Homma who traded in the futures markets in the 1700s discovered that although there was a link between supply and demand of rice, the markets were also strongly influenced by emotions of traders. As a result there could be a vast difference between the value and price of rice. Homma realized that he could benefit from understanding the emotions to help predict the future prices. He formulated his trading principles in two books, *Sakata Senho* and *Soba Sani No Den*, which were said to have been written in the 1700s. His work, as applied to the rice markets, evolved into the candlestick methodology which is still popular among chartists in Japan. Unfortunately, the results of four hundred years old studies were isolated by cultural and language barriers from the western world up to a moment when they have been rediscovered in the second half of 20th century [121, 161].

In the western world technical analysis starts in the early 20th century with the Dow theory. The theory was developed by Charles Dow based on his analysis of market price action in the late 19th century. Charles Dow never wrote a book or scholarly article on his theory. Instead, he put down his ideas of stock market behavior in a series of editorials that *The Wall Street Journal* published around the turn of the century. In 1903, the year after Dow's death, S. A. Nelson compiled these essays into a book entitled *The ABC of Stock Speculation*. In this work, Nelson first coined the term "Dow's Theory". In 1922, William P. Hamilton categorized and published Dow's tenets in a book titled *The Stock Market Barometer*. Robert Rhea developed the theory even further in the *Dow Theory* (New York: Barron's), published in 1932 [115].

The Dow theory addresses the fundaments of technical analysis as well as general principles of financial markets, which are primarily applied to stock indexes. The theory assumes impossibility of manipulating the primary trend[13], while at short time intervals or with individual stocks the market could be prone to manipulation by large institutional investors, speculators, breaking news or rumors [131].

William Hamilton and Charles Dow openly admitted that the Dow theory is not a sure-fire means of beating the market. It is looked upon as a set of guidelines and principles to assist investors with their own study of the market. The Dow theory was thought to provide a mechanism to help make decisions less ambiguous [131].

During the 1920s and 1930s, Richard W. Schabacker refined the subject of the Dow theory in a somewhat new direction. He realized that whatever significant action appeared in a stock index it must derive from similar action in constituent stocks. In his books, *Stock Market Theory and Practice*, *Technical Market Analysis* and *Stock Market Profits*, Schabacker showed how the principles of the Dow theory can be applied to the charts of individual stocks [48].

Further development of technical analysis was pretty straightforward. First, Richard Schabacker, the interpreter of the Dow theory, was joined by Robert D. Edwards. Then, in 1942 John Magee joined the study of technical analysis. With his participation the entire process of technical evaluation became more scientific. As a result of their research from 1942 to 1948, Edwards and Magee developed new technical methods of technical analysis. They put these methods to practical use in actual market operation. And eventually, in 1948, these findings were published in their definitive book, *Technical Analysis of Stock Trends* [48]. The 8th edition of this book was published in 2001. It demonstrates strong interest of investors in methods of technical analysis.

Technical analysts distinguish five points, which define the importance of technical analysis [121]:

1. While fundamental analysis may provide a gauge of the supply/demand situations, price/earnings ratios, economic statistics, and so forth, there is no psychological component involved in such analysis. Technical analysis provides the only mechanism to measure the "irrational" (emotional) components present in all markets.

2. The application of technical analysis allows investors to separate investment decisions from investors sentiments and to see the market without the prism of subjectivity.

3. Following technical analysis is important even if one does not fully believe in it. This is because, at times, technical analysts themselves are the major reason for

---

[12]John Murphy [114] defined the technical analysis, as a study of market action, primarily through the use of charts, for the purpose of forecasting future price trends.

[13]*Primary trend* is a long-running (up to five years) general movement in price data.

a market move. Since they are a market moving factor, they should be monitored.

4. People remember prices from one day to the next and act accordingly. Peoples' reaction affect prices, but prices also affect peoples' reactions. Thus, price itself is an important component in market analysis.

5. The price change is the most direct and easily accessible information of the combined effect of different factors.

All but the second point seem to be acceptable. The second point is unrealistic since it requires enormous self-control of a technical analyst. To make it valid one would need a machine with intelligence and expertise of a technical analyst and zero whatsoever emotions.

Following the classification by Christopher Neely [117] the methods of technical analysis attempt to identify trends and reversals of trends. To distinguish trends from shorter-run fluctuations, technical analysts employ two types of analysis: *charting* and *technical (mechanical) rules*. Charting, the older of the two, involves graphing the history of prices over some period - determined by a practitioner - to predict future patterns in the data from the existence of past patterns.[14] The second type of methods, technical rules, imposes consistency and discipline on technical analysts by requiring them to use rules based on mathematical functions of present and past prices.

To identify trends through the use of charts, technical analysts must first find *peaks* and *troughs* in the price series. A peak is the highest value of the price within a time interval under consideration (a local maximum), while a trough is the lowest value the price has taken on within the same time period (a local minimum). A series of peaks and troughs establishes downtrends and uptrends, respectively.

Detecting a trendline allows technical analysts to issue a short-term investment recommendation. Usually if an uptrend is detected the recommendation is a long position, alternatively, for the downtrend it is a short position.

Spotting the reversal of a trend is just as important as detecting trends. Peaks and troughs are important in identifying reversals too. Local peaks are called *resistance levels*, and local troughs are called *support levels*. If the price fails to break a resistance level (a local peak) during uptrend period, it may be an early indication that the trend may soon reverse.

Technical analysts identify several patterns that are said to foretell a shift from a trend in one direction to a trend in the opposite direction. The best known type of reversal formations called "head and shoulders". The head and shoulders reversal following an uptrend is characterized by three local peaks with the middle peak being the largest of the three. The line between the troughs of the shoulders is known as the "neckline". When the price penetrates the neckline of a heads and shoulders, technical analysts confirm a reversal of the previous uptrend and issue a recommendation to take a short position.[15]

Another method of charting is the *candlestick technique* developed in Japan more than four centuries ago for rice and its futures market. The technique is based on the recognition of visual patterns that take the shape of candlesticks. Every candle includes information on the high, low, opening and closing prices of a particular time interval. The figurative "body" shows the difference between opening and closing prices, and its length depends on this difference. If the closing price is higher than the opening price, the body is white, which signals rising prices. If the opening price is higher than the closing price, the body is black, which signals falling prices. Above and below the candle's body are the "shadows", called upper shadow and lower shadow. They depict the high and the low of the trading interval [62]. In general, the candlestick technique consists of a set of patterns, defined by candlesticks, and respective expectations of market reaction.

The advantage of candlestick technique is that, first, it allows to express several relative to each other values within one graphical symbol.[16] Second, this technique can be easily combined with other charting methods or with technical rules.

In general, the problem with charting is that it is very dependent on the interpretation of a technical analyst who is drawing the charts and interpreting the patterns. Subjectivity can permit emotions like fear or greed to affect the trading strategy. Technical rules make the analysis more consistent and disciplined and thus allow to avoid the problem of subjective analysts' judgment [117].

There are many types of technical rules. In general, they aim at identifying the initiation of new trends. The best known technical rules are the following:

1. *Filter rules* - buy when the price rises by a given proportion above a recent through.

2. *Trading Range Break* or *Channel rules* - buy when the price rises by a given proportion above a recently established trading range.

3. *Moving Average rules* - buy when the current price level is above the moving average.

4. *Moving Average Intersection rules* - buy when a shorter moving average penetrates a longer moving average from below. They can have a form of Variable

---

Length Moving Average or Fixed Length moving average, which differs in the number of days during which the buy or sell signal is assumed to be issued.

5. *Oscillator rules* - buy (sell) when the oscillator index takes an extremely low (high) value. A simple type of oscillator index is a difference between two moving averages with short and long horizons [94, 117].

6. *Statistical rules* are based upon $ARMA$-family models for rescaled returns. The rules rely on a standardized forecast, given by the one-period-ahead forecast divided by an estimate of its standard error. For example, for $ARMA(1,1)$ an upward trend is predicted when the value of the standardized forecast is positive [153].

7. *Other rules*. Many technical analysts assign a special role to round numbers in support or resistance levels, and to historical record prices.[17] Other prominent types of technical analysis use exotic mathematical concepts such as Elliot wave theory and/or Fibonacci numbers.[18] Finally, technical analysts sometimes use technical analysis of one market's price history to take positions in another market, a practice called intermarket technical analysis [117].

Each rule has a mirror equivalent, which suggests short position. In each case a technical analyst has to choose the time horizon over which troughs and peaks are identified and moving averages calculated as well as the threshold before a decision is made.

Since the introduction of technical analysis there is a growing evidence that many contemporary professional investors use it. The profitability of technical analysis would be in contradiction to the EMH, which postulates that in efficient markets it is impossible to profit by predicting price development based on its past performance.

Alfred Cowles [37] was one of the first scholars who analyzed the profitability of technical analysis. In 1933 he reported results of a hypothetical investment strategy based on market forecasts of William Hamilton in his editorials to *The Wall Street Journal*. The hypothetical application of published forecasts of the stock market based on the Dow theory over a period of 26 years, from 1904 to 1929, achieved a result better than what would be ordinarily regarded as a normal investment return, but poorer than the result of a continuous outright investment in representative common stock for this period.

The study of Harry Roberts [132] conducted in 1959 on American data, for both indexes and individual stocks,

questioned the applicability of technical analysis since time series of prices seemed to follow an extremely simple chance model. He referred to Maurice G. Kendall [82], who obtained the same results for British stock indexes and American commodity prices in 1953. Moreover, Roberts found that even in 1934 Holbrook Working [162] achieved the same conclusion: that [financial] time series commonly possess in many respects the characteristics of series of cumulated random numbers.

In contrast to these results Hendrik S. Houthakker [73] in 1961 found elements of non-randomness in speculative price movements. He presented evidence that stop orders gave rise to a non-random profit. Sidney S. Alexander [3] by using 5-percent filtering of noise showed that after filtering large changes are more likely to continue than to reverse: in speculative markets price changes appear to follow a random walk over time, but a move, once initiated, tends to persist.

Robert Weintraub [158] analyzed the pertinent literature of that time on testing technical analysis. He found that the studies up to 1963 were using too restrictive assumptions about the behavior and abilities of technical analysts. For example, Weintraub argued that Kendall's assumption of fixed interval between trades did not reflect the reality and did reduce potential profit opportunities. By using a more realistic varying waiting time Weintraub obtained results which spoke more in favor of technical analysis than the random walk hypothesis. He concluded that the lack of serial correlation between first differences of closing prices simple meant that speculators [technical analysts] who were supposed to smooth out price movements over time were doing their job well.

In 1964 Sidney S. Alexander [4] tested a number of filter rules. Although they appeared to yield returns above the buy-and-hold strategy for the DJIA and S&P stock indexes, he concluded that adjusted for transaction costs, the filter rules were not profitable. Eugene Fama [49] came to an even more restrictive conclusion: the data seem to present consistent and strong support for the [random walk] model. This implies, of course, that chart reading [technical analysis], though perhaps an interesting pastime, is of no real value to the stock market investor. In [53] Eugene Fama and Marshall Blume achieved similar conclusions,[19] which in 1970 led Eugene Fama [50] to dismiss technical analysis as a futile activity.

In 1967, in spite of the tendency to reject technical analysis, M. F. M. Osborne [123] found that applicability of the random walk theory and technical analysis can be dependent on the underlying time frequency of prices. He con-

---

[17]One can argue that this rule captures the effect of the "psychological" barrier, which market has to overcome.

[18]To get more details on Elliot wave theory and Fibonacci numbers see [114].

[19]They conducted a study of the thirty Dow Jones Industrial Stocks using 24 different filter rules, ranging in size from 0.5% to 50% for the time period from 1957 to 1962. They concluded that the filter rules were not profitable when the effect of interim dividends and brokerage commissions were considered [53].

cluded that in general shorter intervals (daily, weekly) tend to show more "non-random walk" properties than longer intervals (monthly).

After Eugene Fama silenced empirical studies of technical analysis for almost twenty years, in the second half of the 1980s the interest of academic community returned to the topic. New empirical studies either found evidence in favor of technical analysis[20] or defined segments and markets, where the weak form of market efficiency prevails and technical analysis brings a small added value[21].

The return and price predictability is of interest not only in stock markets, but also in the currency markets dominated by professional investors [5, 151].[22] Blake LeBaron [93] showed that simple rules used by technical traders have some predictive value for the future movement of foreign exchange rates. He explained that the reason can be in the nature of foreign exchange markets, where there are several major players whose objectives may differ greatly from those of maximizing economic agents. The results of his study showed that this predictability was greatly reduced, if not eliminated, on the days in which the Federal Reserve was actively intervening were removed.[23]

Along with empirical support technical analysis also received more interest by the theory of financial markets. A typical example is a study of Avraham Beja and Barry Goldman [17], where they showed that incorporating technical analysis could help to explain empirical properties of financial markets. David Brown and Robert Jennings [27] constructed a two-period dynamic equilibrium model to demonstrate that rational investors should use historical prices in forming their demands.

Arthur et al. [9] created a model, populated with artificial technical analysts, known as Santa Fe artificial stock market (ASM). The model allowed the authors to explain the contradiction between the theoretical literature and practitioners' view on market efficiency. The simulation results showed that both views were correct, but within different regimes of the market. The market settled into the rational-expectations equilibrium of the efficient market literature when the ASM agents had enough time to accumulate and process information. Otherwise, if in a hurry to accumulate and process market information investors did place more weight on technical trading. As a result, the market dominated by technical trading experienced temporary bubbles and crashes.

William Clyde and Carol Osler [34] provided theoretical foundations for technical analysis as a method for doing nonlinear forecasting in high dimension systems. They argued that traditional graphical technical modeling methods might be viewed as an equivalent to nonlinear methods that use the Taken method of phase space reconstruction combined with local polynomial mapping techniques for nonlinear forecasting. The study presented evidence in support of this hypothesis in the form of an application of the head-and-shoulders formation identification algorithm to high-dimension nonlinear data, resulting in successful pattern identification and prediction.

Alexandra Ilinskaia and Kiril Ilinski [90] used the framework of Gauge Theory of Arbitrage to show that technical analysis and market efficiency corresponded to different time regimens. From their point of view, technical predictions exist due to an internal deterministic dynamics, which brings the market to an equilibrium. They showed that technical analysis indicators and their predictions existed for short time horizons while for long time horizons the model produced an EMH state with realistic statistical behavior.[24]

Alan Morrison and Nir Vulkan [113] studied a version of the standard Kyle [89] model with endogenous information acquisition. They found that there was a robust equilibrium, which allowed free entry and in which speculators attained positive profits.

Based on this, one might argue that individual investors are aware of short-term financial market returns and price predictability and try to exploit the trading opportunities by using technical analysis.

5. Artificial Technical Analyst

Previous sections have emphasized two competing paradigms in finance: the efficient market hypothesis and the legacy of technical analysis. The overview of empirical tests did not reveal a clear dominance of one paradigm or another. The main reason of this ambiguity is the joint hypothesis problem, which constrains tests of market efficiency. To overcome this limitation one needs to design a model-free test. Spyros Skouras [143] proposed a solution, which allows to test market efficiency independently of an equilibrium model, - his artificial technical analysts. His concept of the ATA will be used to design our artificial technical analyst, which will incorporate latest advances in artificial intelligence, pattern recognition and data preprocessing. This section outlines key components of a new implementation of the ATA.

A blueprint of the ATA has three major components: data preprocessing, pattern recognition and decision-making mechanisms.

The first part of data preprocessing is a homogenization mechanism. Since the idea of technical analysis is the interpretation of past financial data, its quality and consistency

---

[20]For examples see [1, 6, 8, 25, 27, 41, 58, 74, 87, 91, 101, 102, 108, 110, 112, 119, 128, 134, 138, 144, 146, 147, 152].

[21]For examples see [14, 40, 57, 59, 104, 112, 122, 128, 134, 138].

[22]For examples see [19, 63, 99, 103, 116, 122, 124, 148, 149, 150].

[23]A similar topic is discussed by Christopher Neely [118].

[24]The investors aware of this fact might increase the trading frequency to exploit these opportunities.

crucially depends on the quality of data. Two aspects of data preprocessing will be taken into account. The first one is a reduction of problem complexity. And the second one is a maximization of the informational content of the data. Both aspects can be resolved through the mechanism of data homogenization.

The second part of data preprocessing is a piecewise linear approximation (PLA) mechanism. This mechanism serves to identify important peaks and troughs in time series of prices. Since the ATA should have good proficiency in detecting peaks and troughs the quality of piece wise approximation has a crucial role.

The second component in the ATA is responsible for pattern recognition. It will use peaks and troughs, identified by the PLA, to learn patterns in the data. The pattern recognition of the ATA is similar in a sense to the candlestick technique. Like with the candlestick patterns actual values of prices have no importance, but their relative values. In addition, the ATA will incorporate relative waiting time of appearance of each price observation. This approach will allow to see which patterns can predict price behavior.

The cognitive mechanism of the ATA is based on artificial intelligence. Since learning classifier systems (LCS) are one of few implementations of artificial intelligence that have a transparent structure of solutions they will be used for building the cognitive mechanism. There are many implementations of LCS, but all of them are dominated by the XCS, developed by Stewart Wilson [159, 160]. Unfortunately, the XCS does not have the power to solve our problem within reasonable time. To overcome this limitation a new implementation of LCS will be used - the true classifier system (TiCS), - developed by Timur Yusupov.

Finally, using the concept of revealed transaction costs a decision-making function will be introduced.

This section is organized as follows. First, Subsection 5.1 presents the data preprocessing component. It is followed by Subsection 5.2, which outlines the implementation of artificial intelligence, the TiCS. Subsection 5.3 introduces the pattern encoding mechanism. In Subsection 5.4 revealed transaction costs and decision making are presented. A discussion concludes the section.

5.1. Data Preprocessing

This subsection deals with the preprocessing of the data.

With the data coming in raw, tick-by-tick form one has to apply some homogenization procedure to reduce the data and the problem complexity. There are two methods of homogenization. That is previous- and linear-tick interpolation. Empirical literature does not indicate any pronounced difference in the results produced by either method [39]. Due to simplicity and fast computability the previous-tick interpolation is selected.

The drawback of homogenization is inefficient use of data. High-frequency data has much more observations than is required for some fixed frequency of homogenization. At the same time the ATA needs as much data as possible for fast learning. One way to "feed" the ATA and to improve the efficiency of data processing is to use sequential root time shift in the homogenization procedure. This allows to achieve almost 100 percent efficiency of data utilization.

One important aspect of homogenization is the necessity to exogenously specify the optimal frequency. The criterion of optimality is profitability. To get the sign and the magnitude of the profit one has to consider transaction costs (TC). With no knowledge of TC the notion of revealed transaction costs can shed light on hypothetical profits.[25] Aggregate returns adjusted for a schedule of transaction costs, as a direct indicator of profitability, should be computed for a grid of frequencies of homogenization. The maximum value of breakeven revealed transaction costs will indicate the optimal frequency of homogenization.

Another related aspect of data preprocessing is selection of the optimal subsample size. For consistent application of technical trading rules the ATA should have a well defined and fixed subsample of recent observations. Like in the previous case the grid analysis and revealed TC should be used for the search.

To sum up, without any exogenous information on the frequency of homogenization, the value of transaction costs or optimal subsample size one needs to perform parallel search along all three variables, which results in a three-dimensional search problem. The optimal value of the subsample size, or the frequency of homogenization can vary across stocks, markets and through the time. As a result, wrong selection of these parameters can destine the ATA performance to be very low from the very beginning.

The PLA is an important ingredient of the ATA.[26] Without this mechanism the system is insensitive to turning points in time series. The segmentation algorithm is an effective way to perform identify these points. For this the PLA constructs a representation of the original time series by several linear segments. Assuming the ends of linear segments are connected, one can use those joints as identification points of the underlying time series. These identification points will be used for similarity search and subsequent forecasting of the next observation in the time series.

Depending on speed or accuracy needed one can pick the most appropriate algorithm out of possible alternatives.[27]

---

[25]For details on revealed transaction costs see Subsection 5.4.

[26]The formal definition of the *piecewise linear approximation* is an approximation of a time series of the length n by k linear segments. Normally, one chooses k to be much smaller than n, which makes the storage, transmission and computation of the data more efficient [83]. The abbreviation of the piecewise linear approximation is PLA. One can refer to it as *segmentation algorithm* or *approach*.

[27]For an overview of PLAs see [83].

Since the Top-Down algorithm[28] takes a constant time, and has acceptable degree of accuracy it is chosen for application in the ATA.

5.2. Learning Classifier Systems: TiCS

The artificial technical analyst is designed to replicate the professional expertise of a technical analyst. The quality of an analysis crucially depends on the cognitive ability of the ATA. Imitation of human-like cognition is a very cumbersome task, since one needs to match the abilities of the human brain. The sophistication of the human brain allows it to instantaneously process incoming information, to associate it with past experience, and to make sensible prognoses. As a result we are able to learn from our experience and to generalize it to new, unseen situations. With the help of Artificial Intelligence (AI) the ATA should be able to replicate the process of human-like cognition.

In the subclass of AI - Machine Learning, learning classifier systems (LCS) are algorithms meant to imitate the human ability for classification, learning and generalization. The Encyclopedia Britannica defines these abilities as follows. *Classification* is the ability to systematically arrange in groups or categories according to established criteria. *Learning* is the ability to adapt to the environment and to alternate behavior as a result of individual experience. And *generalization* is the ability to respond in the same way to different but similar environmental conditions. To meet these criteria, LCS should possess the following characteristics.

(i) the ability for on-line classification and establishing the patterns of the different environmental situations.

(ii) the ability to distinguish and preserve the most persistent patterns.

(iii) the ability to ignore any irrelevant or noise information.

The LCS were first introduced by J. Holland [70, 71].[29] LCS were designed to read the current environment at state in terms of a fixed number of predetermined conditions and to provide the most adequate mapping into the space of coming events. For this purpose LCS employ Genetic Algorithms (GA) during the training period to identify a correct mapping from the combination of predetermined conditions to the most probable event. In the process, LCS should identify irrelevant conditions from noise and distinguish persistent combinations of conditions.

In the seminal work of J. Holland and in the related literature of the following decades the adequacy of each classifier was measured by the predetermined criteria, known as strength or fitness. This parameter was serving both as a predictor of future payoff and as the classifier's fitness for the genetic reproduction. Unfortunately, this primitive aggregation resulted in low performance of LCS. As a result, the considerable enthusiasm of the 1980's declined in the early 1990's. LCS seemed too complicated to be studied, with only few successful applications reported. In the mid 1990's the field appeared almost at dead end [72].

In response to this situation Wilson [159] introduced the XCS. The primary distinguishing feature of the XCS is that classifier fitness is based on the accuracy of classifier payoff prediction rather than on payoff prediction (strength) itself.

Although the XCS is currently a favorite, it has some disadvantages. Its overcomplicated structure slows down the algorithm in finding solutions. The XCS originates from a "zeros level" classifier system, which was intended to simplify Holland's canonical framework while retaining the essence of the classifier system idea [159]. At some step the intention to simplify turned out to overcomplicate. To illustrate this point [29, 28] list at least 28 parameters and switches, which need to be specified. Their numerosity and vague explanation in the source literature[30] makes tuning of XCS an art rather than a precise science. These parameters are claimed to change the XCS's behavior, adjust it to the current problem and specify output characteristics [29]. Unfortunately, low transparency of the algorithm and a lack of theoretical studies make re-specification of the parameters of the XCS impossible.[31] As a result the XCS shows high inertia in learning relatively simple problems. To overcome those limitations a new classifier system was introduced - the True Classifier System (TiCS).

The key element of the TiCS is the Micro-Genetic Algorithm (μGA)[32]. Unlike conventional algorithms the μGA requires specification of only 2 parameters. The same 2 parameters are used for the TiCS activation and run. The first parameter N_g defines the global population size of classifiers. The second parameter N_s instructs the algorithm on how many classifiers should match the input signal to form the sub-population.

---

[28]In the Top-Down algorithm a time series is recursively partitioned until some stopping criteria. The first approximation is one line, which connects the first and the last point in the original time series. To find a better approximation one evaluates every possible partitioning of the previous approximation. The partitioning which provides the best goodness of fit locates next split. Newly segmented approximation is reexamined for a new partitioning and the process repeats. The algorithm runs until either a benchmark goodness of fit is reached, or an approximation gets enough linear segments.

[29]For the introduction to LCS see [65, 109], for a recent surveys on its applications and development see [92].

[30]There are only few articles were some of this parameters are studied in detail. The best description is provided by [29, 28].

[31]For this reason the comparative study uses the default parameters of the XCS.

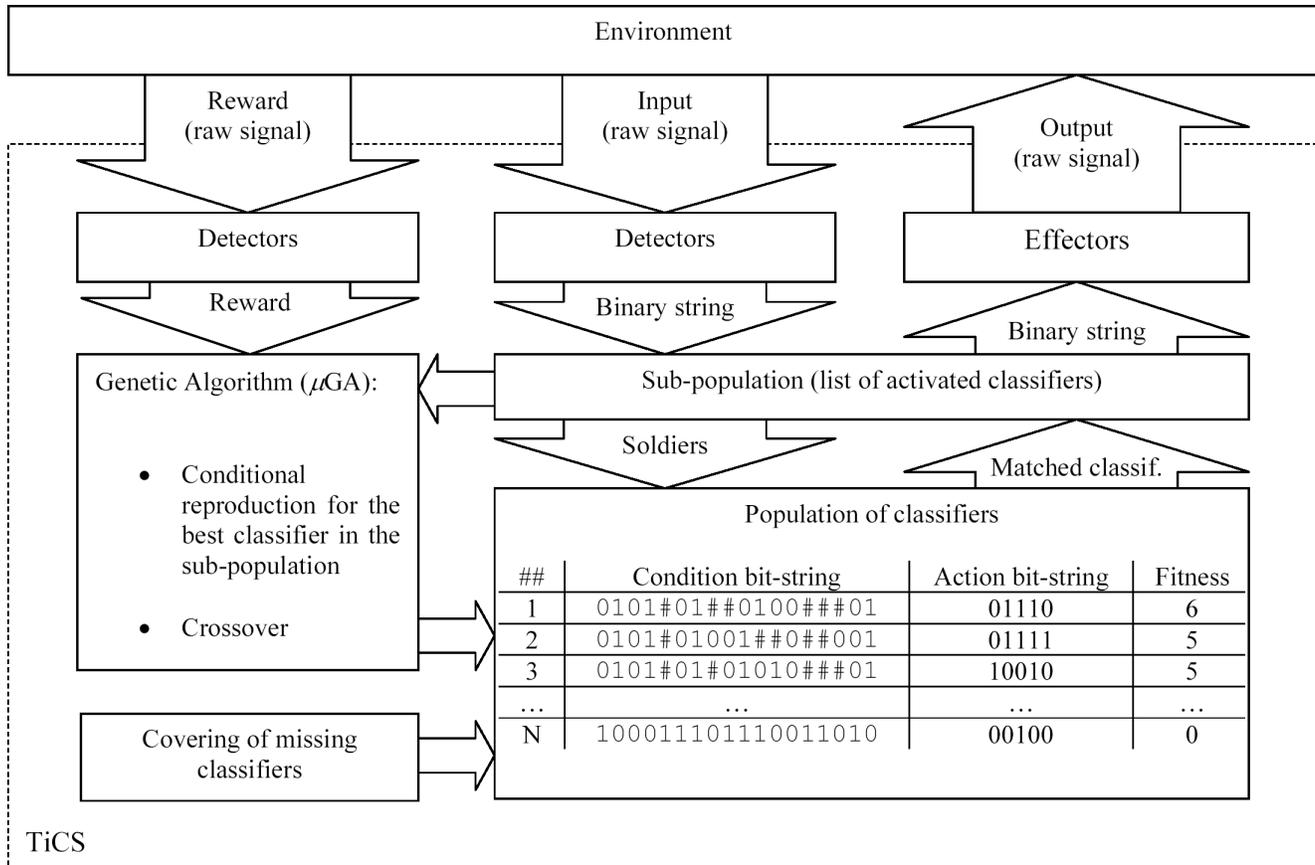[32]For details on the μGA see [88].

Figure 1. TiCS and its environment

Figure 1 outlines the operation of the TiCS. It follows the original idea of Holland's LCS. From a dynamic environment the TiCS extracts static sub-problems to provide an adequate response. For this purpose an input signal is matched against the condition part of each classifier. The matched classifiers form a current sub-population. Each classifier in the population should be as general as possible, i.e. it should correctly respond to the maximum number of states. A matched classifier is assumed to have the correct response. Given the goal of generalization, the measure of fitness is the quantity of "don't care" symbols in the condition part. This allows to rank the classifiers. If the number of matched classifiers is more than N_s, the classifiers with the lowest fitness are excluded from this sub-population.

The classifiers in the sub-population are referred to as **S**oldiers, since they are too **S**pecialized. The best in fitness classifier is referred to as a **G**eneral, since it achieved the most **G**eneral level without failing. The General forms the TiCS response to the input signal. Meanwhile, the classifiers in the sub-population share the "experience" through the μGA. The probability of being selected is proportional to fitness. Selected classifiers crossover their condition and action parts to form a new sub-population. The General enters the new sub-population only if his response was adequate. The new sub-population then replaces the current sub-population in the global population.

In case there is not a sufficient number of classifiers in the sub-population a covering mechanism fills missing positions. Covering creates classifiers that match a current input signal by copying it to the condition part of new classifiers and replacing some condition bits with "don't care" symbol #. The corresponding actions are randomly generated.

5.3. Pattern Encoding Mechanism

Since the TiCS is intended to be used for pattern recognition in time series it requires a particular encoding mechanism. This mechanism allows to encode time series into strings of conditional bits.

At the first step a subsample of the time series is demeaned and normalized to have unit variance. This allows to recognize similar patterns irrespective of the current level or the variance. Next, the PLA algorithm identifies positions of key points describing the subsample of time series.

Each key point connects linear approximations found by PLA. With n predefined linear segments there are $n + 1$ key points.

Application of demeaning and normalization allows to describe the position of each key point within the same two-dimensional coordinate space. An ordinate dimension corresponds to the level in time series and is constrained to the interval $[-2, 2]$. A coordinate dimension corresponds to time stamps of time series and is constrained by the interval $[1, L]$, where L is the length of the subsample.

At the second step the encoding takes place. First, two equal segments of a coordinate and ordinate space are defined. In the beginning they are $[-2, 0)$ and $[0, 2]$ for the coordinate axe, and $[1, L/2)$ and $[L/2, L]$ for the ordinate axe. Next a coordinate (ordinate) of a key point is assigned to one of the two segments. If it belongs to the below (left) segment of the coordinate (ordinate) then the first condition bit is assumed to be 1, otherwise it is assumed to be 0. Next the segment, where the key point is located is divided into two new equal size segments. And again, the key point is assigned to one of the two segments and the corresponding bit is added to the condition bit-string. After a predefined number of iterations the coordinate (ordinate) position of the key point is encoded by a sequence of zeros and ones.

Example 1: One needs to encode a position of a key point from a demeaned and normalized subsample of 50 observations. The level of the key point is 1.17, and the time stamp is 11. Table 1 provides the results for encoding the first 5 bits. After performing iterative process the level is encoded as 11000, and the time stamp is encoded as 00111.

The encoding mechanism allows decoding of a condition bit-string to get the coordinate and ordinate limits. Since the TiCS has in the alphabet of condition bits the "don't care" symbol #, during decoding the process stops either when all bits are processed or when the first symbol # is found. The latter case allows the TiCS to have varying interval within which key points should be.

Example 2: Using the same settings as in the previous example one needs to decode bit-strings of coordinates and ordinates: 010## and 11010, respectively. Table 2 provides the results for decoding these bit-strings. After decoding the level is found to be within $[0.500, 1.000)$, and the time stamp is decoded to be within $[07, 09)$.

5.4. Revealed Transaction Costs and Decision Making

Long and short selling are the most basic transactions in financial markets. Any complex arbitrage strategy is just their combination. The applicability of either transaction normally depends on the expectation about the future price movement and transaction costs. Wrong assessment of either of them can cause a financial loss. With endogenous

Step ##	Ordinate segment		Ordinate bits
	Left	Right	
1.	$[01, 25)*$	$[25, 50]$	1____
2.	$[01, 12)*$	$[12, 25)$	_1___
3.	$[01, 06)$	$[06, 12)*$	__0__
4.	$[06, 09)$	$[09, 12)*$	___0_
5.	$[09, 10)$	$[10, 12)*$	____0
	Final ordinate bit-string:		11000

Step ##	Coordinate segment		Coordinate bits
	Low	Up	
6.	$[-2.00, 0.000)$	$[0.000, 2.000]*$	0____
7.	$[0.000, 1.000)$	$[1.000, 2.000]*$	_0___
8.	$[1.000, 1.500)*$	$[1.500, 2.000]$	__1__
9.	$[1.000, 1.250)*$	$[1.250, 1.500)$	___1_
10.	$[1.000, 1.125)*$	$[1.125, 1.250)$	____1
	Final coordinate bit-string:		00111

Table 1. Illustation of pattern-to-bits encoding mechanism

A subsample of time series has 50 observation (ordinate axe). Values are demeaned and normalized. Encoding of a key point with the level of 1.17 and the position of 11. Asterisk denotes that the position is within the marked segment.

price expectations one can evaluate maximum transactions costs at which arbitrage strategies generate positive profits.

Transaction costs include commissions paid per transaction, internal R&D expenses, a bid-ask spread and the impact on price. Every component complicates measuring actual transaction costs. If, for example, there is only a time series of past prices, then the bid-ask spread, past commissions and the impact on price have to be estimated from the data. Roll [135] introduced an implicit measure of effective bid-ask spread that can be estimated from the univariate time series. Lesmond et al. [96] used time series observations of zero returns to estimate transaction costs of a marginal investor.[33] Unfortunately, for the precise estimation of transaction costs one needs a complete data set to capture all the components, which are usually not accessible in time or have confidential nature.

Known transaction costs enable one to reveal the return sufficient to insure a positive profit. If, for some reason, the exact value of transaction costs is not known, a reverse solution can be found, i.e. one can find transaction costs that insure non-negative profits. Those costs can be derived from the arbitrage strategies mentioned above and are referred to as *revealed transaction costs*. Bessembinder et al. [20] used a similar concept, known as breakeven costs. In relation to our definition the breakeven transaction costs are maximum

---
[33]For the latest review of literature on transaction costs see [97].

Step ##	Ordinate bit-string	Ordinate segment	
		Left	Right
1.	1____	[01, 25)*	[25, 50]
2.	_1___	[01, 12)*	[12, 25)
3.	__0__	[01, 06)	[06, 12)*
4.	___1_	[06, 09)*	[09, 12)
5.	____0	[06, 07)	[07, 09)*
	Final ordinate interval:	[07, 09)	

Step ##	Coordinate bit-string	Coordinate segment	
		Low	Up
6.	0____	[−2.00, 0.000)	[0.000, 2.000]*
7.	_1___	[0.000, 1.000)*	[1.000, 2.000]
8.	__0__	[0.000, 0.500)	[0.500, 1.000)*
9.	___#_		
10.	____#		
	Final coordinate interval:	[0.500, 1.000)	

Table 2. Illustation of bits-to-pattern decoding mechanism

A subsample of time series has 50 observation (ordinate axe). Values are demeaned and normalized. Decoding of coordinate bit-string 010##, and ordinate bit-string 11010. Asterisk denotes selected segment.

costs that insure zero profit.

Following the differentiation of strategies, one can derive revealed transaction costs for the case of expected increase and decrease in the asset's price. Given the high-frequency nature of data and transactions, one can neglect interest rates as well as the influence of dividends and splits.

For the case of expected price increase the value of revealed transaction costs, $\widehat{c}_t^L$, is derived from long position, and should satisfy the following inequality to insure a positive profit:

$$c < \frac{E\left[p_{t+1}\right] - p_t}{E\left[p_{t+1}\right] + p_t} \equiv \widehat{c}_t^L, \qquad (1)$$

where c is actual transaction costs, p_t is the current price level, and $E\left[p_{t+1}\right]$ is expected next period price. For the short position the value of revealed transaction costs, $\widehat{c}_t^S$, should satisfy the following inequality to insure positive profit:

$$c < \frac{p_t - E\left[p_{t+1}\right]}{E\left[p_{t+1}\right] + p_t} \equiv \widehat{c}_t^S. \qquad (2)$$

A combination of both cases allows to formulate a decision-making function, which takes as arguments expected and current price, and the value of anticipated trans-action costs:

$$D_t\left(c, p_t, E\left[p_{t+1}\right]\right) = \begin{cases} B & \text{if} \quad c < \dfrac{E\left[p_{t+1}\right] - p_t}{E\left[p_{t+1}\right] + p_t}, \\ S & \text{if} \quad c < \dfrac{p_t - E\left[p_{t+1}\right]}{E\left[p_{t+1}\right] + p_t}, \\ N & \text{otherwise}, \end{cases} \qquad (3)$$

where B denotes a decision to buy, S - a decision to short-sell and N - to do nothing.

6. Tokyo Stock Exchange: Market Microstructure and Data Description

This section describes the data used in this paper. The data cover historical stock prices collected on the Tokyo Stock Exchange (TSE).

The TSE is a classical example of an order-driven market. It operates as a continuos auction, where buy and sell orders interact directly with one another. The study of the order-driven activity of the TSE did not show significant differences from the New York Stock Exchange or from stocks markets with designated market-makers [7, 15, 95]. At the same time the TSE market has some specific features, which can have an impact on empirical results. Since the results of empirical analysis can be sensitive to the method of data selection, the section describes our method of data selection, which minimizes the impact of selection bias on empirical results.

The TSE has no market-makers. All orders, whether limit or market orders, are placed by broker/dealer trading participants and matched in accordance with *price priority* and *time priority* rules. Under the price priority rule, a sell (buy) order with the lowest (highest) price takes precedence. Under the time priority rule, an earlier order takes precedence over others at the same price. Thus, when the lowest sell and highest buy orders match in price, the transaction is executed at the price.[34]

At the TSE there are two transaction methods: the *itayose* and *zaraba*. The itayose method is used mainly to determine opening and closing prices. Under the itayose method, the time priority rule is not applied and numerous orders placed before price setting are matched in aggregate. In contrast, under the zaraba method, both the price priority and time priority rules are applied, and pairs of buy and sell orders are matched continuously.

The TSE adopts several measures to prevent wild short-term fluctuations in prices: special bid and ask quotes, daily price limits, trading units and margin transactions.[35] These measures do not only help ensure price continuity, but also in effect work as "circuit breakers" in an emergency. In

---
[34] For more details and specific parameters see the source [155].
[35] For their details see [155].

addition the TSE uses *off-auction trading* to handle large block orders. The off-auction trading system allows to accommodate large block orders and basket order transactions, the trading of which is difficult under the competitive trading scheme of auctions. To eliminates the risk of trade counterparty default the TSE utilizes the central counterparty system and the clearing participant system for the settlement of trading on its exchange. Settlement for normal domestic equity transactions in the exchange market is made on the fourth business day starting from the transaction date ($T + 3$).

The TSE has three trading sessions per day: from 8:20 till 9:00, from 11:00 till 12:30 and from 15:00 till 16:30. For half-trading days only the first two sessions take place. The total trading time sums up to 240 minutes over the 490 minutes between the opening and closing time.

The empirical part of the paper studies the data collected on the TSE. Data comes in a raw, tick-by-tick form, and includes mid-prices and volumes of each transaction. The precision of time stamps is one minute. The dataset covers 2273 stocks traded at the TSE during a period from 11/03/1996 till 22/06/1998, i.e. 833 trading days.

Due to computational limitations three subsets of time series were selected for the empirical study. Under the assumption that the return (or price) predictability and, correspondingly, investors' attitude might be different for different levels of liquidity the dataset was divided into three groups based on the liquidity level of the underlying stocks.[36] Each subset was randomly (with withdrawing) populated by 9 time series, which belong to stocks with the same level of liquidity. For each group of liquidity the median stocks provided the tenth time series, thus creating a high-, medium and low-liquidity pool of time series.

Table 3 lists selected stocks and provides statistical properties of their time series. Statistical properties are obtained for daily data over the time period covered in the study. A short description of the occupational activity of underlying companies accompanies each record. The stock that represents the median in liquidity is marked by the asterisk next to a company name. Statistical properties cover mean (μ), sample variance (σ) and sample slope (β).

Figure 2 presents equally weighted indexes for each liquidity pool of TSE stocks (thick line). Prior to aggregation each constituent time series was rebased by setting its first observation in the series equal to unity (thin lines).

7. Application of the ATA to The Tokyo Stock Exchange Data

Technical analysis, i.e. detection of patterns in time series of prices, is a direct contradiction to the weak form of

---

[36]The criterion of liquidity is the number of trades over the covered time period.

Company name (sektor code)	Statistics of price series		
	μ	σ	β
1. High-liquidity pool			
Sumitomo Marine & Fire Insurance (1)	814.98	90.26	−0.27
Nisshin Steel (2)	288.24	97.34	−0.54
Kenwood (3)	524.83	122.22	−0.59
Daihatsu Motor (4)	613.29	71.99	−0.32
Melco (5)	3053.31	917.88	−4.39
Ricoh (5)	1386.98	211.53	0.82
Hitachi (6)	1056.57	107.76	−0.13
Kirin Brewery (7)	1140.67	119.82	−0.38
Matsushita Elec. Industrial (3)	1992.27	178.00	0.57
Adeka* (8)	772.51	164.61	−0.87
2. Medium-liquidity pool			
Hoko Fishing (9)	245.24	112.93	−0.61
Chinon Industries (10)	857.62	480.91	−2.62
Tokyo Style (11)	1517.73	239.48	−1.23
Miraca Holdings (12)	710.59	212.03	−1.10
CO-OP Chem. (8)	256.82	123.70	−0.70
Canon Electr. (5)	818.12	81.70	0.00
Tokyo Cosmos Electric (6)	326.55	110.50	−0.59
Chunkyo Coca-Cola (7)	1054.55	114.99	−0.46
Hitachi Medical (13)	1569.24	196.23	−0.86
Aichi Tokei Denki* (6)	458.38	171.07	−0.96
3. Low-liquidity pool			
Tonichi Carlife GP (14)	503.26	164.28	−0.86
Sonton Food Ind. (9)	1296.90	146.96	−0.72
Tensho Elec. Ind. (10)	352.64	120.43	−0.60
CFS (15)	1121.76	253.27	−1.42
Tokyo Auto Machine Works (10)	472.81	156.88	−0.86
Smbc Friend Sec. (16)	295.27	146.52	−0.80
Lohmeyer (9)	269.37	109.71	−0.60
Daiwa Heavy Ind. (10)	278.31	110.56	−0.61
JSP (8)	1021.06	207.80	−1.07
Unozawa-Gumi Iron Works* (10)	426.58	118.27	−0.56

Table 3. List and description of stocks used in the empirical study

The table uses the following sector codes:

(1) - Nonlife Insurance; (2) - Industrial Metals; (3) - Leisure Goods; (4) - Automobiles & Parts; (5) - Technological Hardware & Equipment; (6) - Electrical, Electronic Equipment; (7) - Beverages; (8) - Chemicals; (9) - Food Producers; (10) - Industrial Engineering; (11) - Personal Goods; (12) - Pharmacology, Biotechnology; (13) - Health Equipment & Services; (14) - General Retailers; (15) - Food & Drug Retailers; (16) - General Financial.
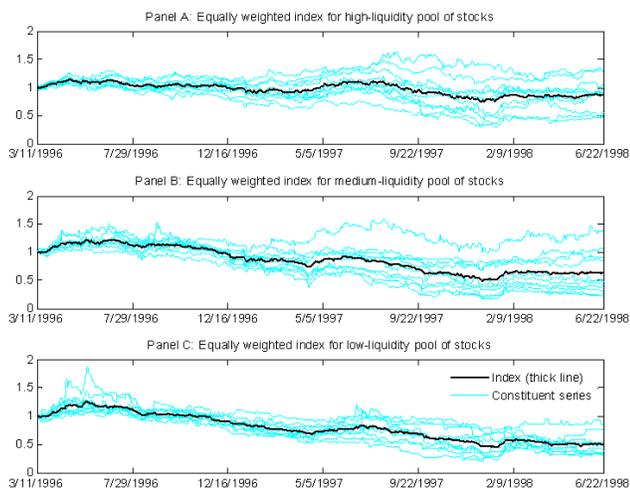
Figure 2. Equally weighted indexes for liquidity pools of TSE stocks

the EMH. This hypothesis is a corner stone of most theories of financial markets. The EMH postulates that in weakly efficient financial markets it is impossible to profit by predicting next period's price used on time series of past prices or returns. The contradiction turns technical analysis into an "outlaw" of financial theory separating the financial community into two camps: the camp of academics, who advocate the EMH; and the camp of practitioners, who keep using technical analysis.

The overview of practices in financial industry shows that many investors, including professional traders and analysts, are widely using technical analysis [5]. Previous investigations demonstrate that under some conditions technical analysis can provide substantial returns. Unfortunately, the findings are undermined by two limitations. Technical analysis is a complex method, which includes rigorous mathematical formulas as well as abstract visual patterns of time series. As a result, researches use only a fraction of the methodology, which dramatically constrains the potential of technical analysis. On the other hand, when the ability of an expert, who is knowledgeable of and fluent in all aspects of technical analysis, is studied one cannot distinguish whether the results are solely due to technical analysis or a subjective interference of an analyst.

The ATA overcomes these drawbacks. The ATA is a robust implementation of the technical analyst's expertise. It follows the aim of technical analysis to detect patterns in price time series. The cognitive mechanism of the ATA, the TiCS, uses the methodology of technical analysis to generate online forecasts of next period's price in a dynamic

environment. To insure the robustness and effectiveness of generated forecasts the ATA performs a sequence of data preprocessing steps. Investment decisions are generated to obtain aggregated returns adjusted for transaction costs.

In the process of calculation several assumptions are made. First, the impact of inflation, dividends and splits is disregarded, since at high frequency their impact is negligible. Second, the aggregate return is computed separately for every time series in the study. Third, trading strategies are applied at every period, i.e. at the beginning of each period the previous position should be closed. Correspondingly the aggregate return is a product of returns in all time periods. And finally, the decision-making process and calculation of returns take into account round-trip TC.

7.1. Decision-Making Function and Its Informational Base

This section presents the decision-making functions and its informational base.

The informational base of the decision-making process includes two components. The first component are expected transaction costs. Transaction costs have a direct influence on trading decisions, since the value of extracted return should be sufficient to cover round-trip transaction costs and with a stochastic return distribution the probability of trading is higher under small transaction costs than under high transaction costs. Due to its hypothetical nature this empirical study uses a set of discrete values of transaction costs taken in a grid from an interval $[0, 2]$ of percent of transaction value with an increment of 0.001 percent.

The second component of the informational base is a subsample of recent prices observations. Technical analysis requires this subsample for forecasting the next period's price. Specifically to the ATA, it requires specification of the subsample size. The ATA is able to identify relevant price observations and filter out the rest. But, in case subsamples are too small or too big, the ATA does not have enough price observations to compute forecasts or, respectively, it would take more time to detect the pattern than the duration of this pattern in the market. To cover this aspect the study uses a set of discrete subsample sizes taken in a grid from an interval $[50, 250]$ of price observations with an increment of 10 observations.

Along with the two explicit components of the informational base for a decision there is a third component - the frequency of decision making. It has an implicit influence on outcomes of decision making. A low frequency allows traders to avoid frequent expenses on transaction costs. At the same time, in a volatile market it constrains traders to take a speculative advantage from local trends. The study uses a set of discrete frequencies in a grid within the interval $[1, 10]$ of minutes with an increment of one minute.

In general, the process of decision-making is formalized through a function $D()$, which maps the arguments into the space of trading actions:

$$D\left(c, p_t, p_{t-1}, ..., p_{t-n}\right) = \{B, S, N\}, \qquad (4)$$

where c stands for transaction costs expressed in fractions of transaction value; $\{p_t, p_{t-1}, ..., p_{t-n}\}$ are current and lagged asset prices; B denotes the buy transaction; S denotes short-selling; and N means to remain inactive.

Normally, the decision-making process includes two steps: forecasting of next period's price and maximization of a next period return adjusted for transaction costs. Forecasting is specific for each strategy since in every case there will be specific forecast mechanism. In general forecasting provides an estimate of the next period asset price conditional on the informational base, $E\left[p_{t+1} | \{p_t, p_{t-1}, ..., p_{t-n}\}\right]$. The second step, return maximization, performs a search through the outcomes of possible trading actions to find the highest return. This step is based on the notion of the revealed transaction cost. Thus, the general form of the decision-making functions is equation (3) with expectations $E\left[p_{t+1}\right]$ been conditioned on past prices $p_t, p_{t-1}, ...$ that enter the TiCS's pattern recognition algorithm.

Unlike the conventional forecast methods the ATA performs interval price forecasts. This allows to convey the forecast of price developments and the ATA confidence in those forecasts.[37] The confidence is indicated by the difference between the upper and lower limits of the forecast. A high confidence is reflected in a small difference, while a low confidence is conveyed through a high difference. The smallest difference is defined to be one basic point in the market. The highest difference is infinity.

The ATA uses technical analysis to obtain the interval forecast, which is constrained by the upper and lower limits: $\overline{p}_{t+1}$ and, respectively, $\underline{p}_{t+1}$.[38] Once the interval is defined one of the three possible trading actions is selected to maximize a current return. This is formalized in the ATA decision-making function:

$$D_t\left(c, p_t, p_{t-1}, ..., p_{t-n}\right) = \begin{cases} B & \text{if} \quad c < \dfrac{\underline{p}_{t+1} - p_t}{\underline{p}_{t+1} + p_t}, \\ S & \text{if} \quad c > \dfrac{p_t - \overline{p}_{t+1}}{p_t + \overline{p}_{t+1}}, \\ N & \text{if} \quad \text{otherwise.} \end{cases}$$
$$(5)$$

---

[37]The ATA does not distinguish whether a low confidence is only due to a lack of experience in a new situation or/and due to a presence of risk associated with an underlying asset, since either case leads to uncertainty.

[38]These limits are the decoded action suggested by the TiCS. Since the ATA operates with homogenized data only the limits of levels (in this case prices) are returned. In the future, the ATA can be extended to work with irregularly spaced time series. In this case the forecast of the ATA will include additionally the time interval within which the next level is predicted.

The aggregate return adjusted for transaction costs is calculated as:

$$R = \prod_{t=1}^{T} r_t, \qquad (6)$$

where the time is defined on the interval from 0 to T. The return in each time period, r_t, is calculated as:

$$r_t\left(...\right) = \begin{cases} \dfrac{p_t - (p_{t-1} + p_t)\, c}{p_{t-1}} & \text{if} \quad D_{t-1} = B, \\ \dfrac{p_{t-1} - (p_{t-1} + p_t)\, c}{p_t} & \text{if} \quad D_{t-1} = S, \\ 1 & \text{if} \quad D_{t-1} = N. \end{cases}$$
$$(7)$$

## 8. Performance of ATA Trading Strategy

Figures 3 - 5 present the results of hypothetical application of the ATA trading strategy on the TSE data. The first column of each figure shows breakeven separation surfaces. Columns 2 - 3 display the slice views corresponding to the first column. Each row represents the time span over which the returns are aggregated. The separation surface shows at which transaction costs and at which combination of trade frequency and subsample size the application of the ATA trading strategy is breakeven. Any point above the surface indicates a loss, while any point below is a gain. On the slice view fat dots show breakeven transaction costs at either fixed subsample size (column 2) or trade frequency (column 3). Color notation distinguishes the magnitude of gains (in shades of red) or losses (in shades of blue).

Table 4 presents the summary of the figures. It lists the maximum breakeven transaction costs obtained for three groups of stock liquidity. Each value of breakeven transaction costs is accompanied by the frequency of trading and the subsample size at which it was observed.

The first conclusion is that the ATA can successfully apply technical analysis to extract positive returns adjusted for transaction costs. The best performance is achieved at a trading frequency of 4 minutes, the subsample size of 50 observations, with the medium liquidity stocks. Here the revealed transaction costs are 0.32 percent of the transaction value. If under these settings the actual transaction costs would be strictly less than 0.32 percent of the transaction value, then a technical analyst would perceive this market segment at that time as weak-form inefficient.

The analysis of specific time spans reveals that the values of maximum breakeven transaction costs are increasing through time horizons.[39] This observation is valid for all liquidity pools. There are two possible explanations. The

---

[39]That is in the first time span the value of maximum breakeven transaction costs is less than in the second one, and in the second time span the value is less than in the third time span.
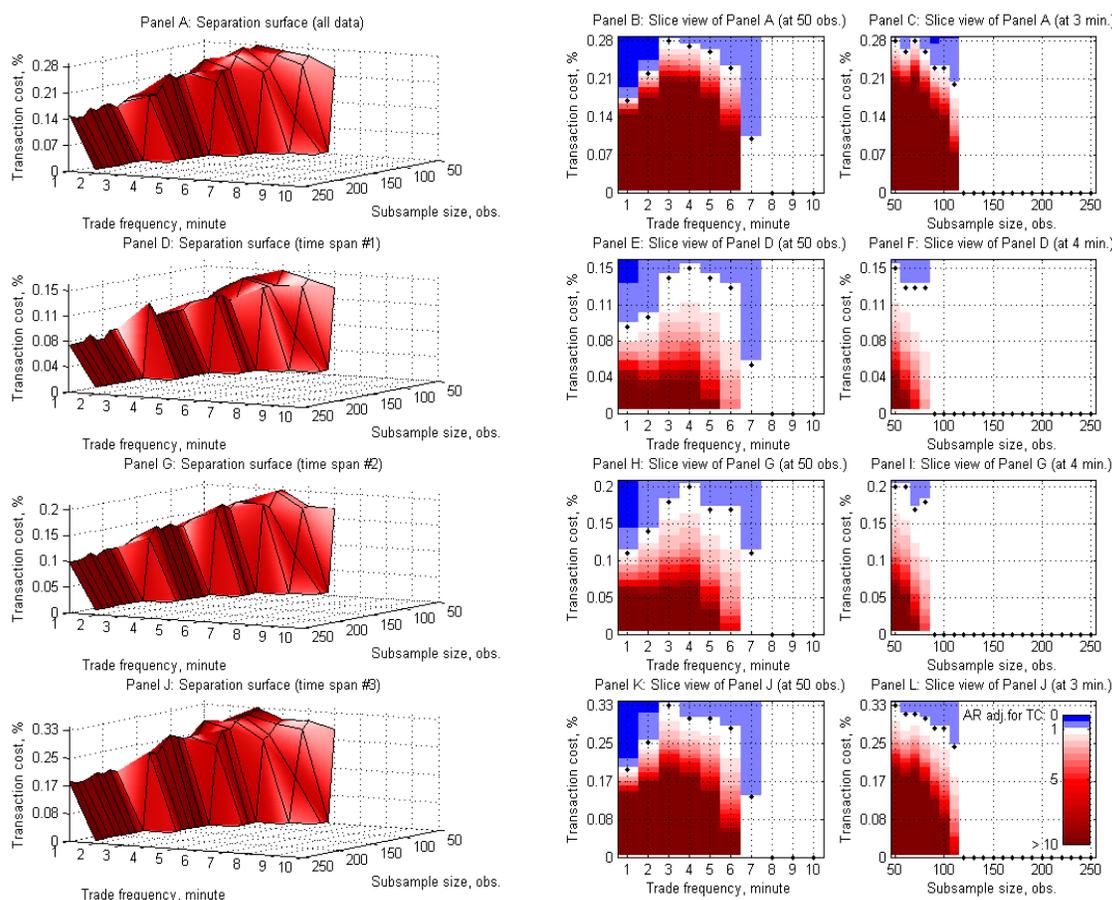
Figure 3. Application of ATA trading strategies on the pool of high liquidity stocks

first one is that the data have different structure in each time span, which, in relation to the ATA forecasting, translates into a different degree of predictability and different distribution of returns in each time span. The second explanation is a learning ability of the ATA. That is if the distribution of prices and the degree of predictability is the same for each time span, then the increase in the maximum breakeven transaction costs is due to the ATA learning.

To test the validity of the first explanation the histogram and empirical CDF of returns were produced for each time span and liquidity pool. The analysis of returns instead of raw prices allows to avoid the problem of comparing time series with different levels of prices.[40] Figure 6 presents the histograms and empirical CDF of returns. Columns from left to right correspond to high-, medium- and low-liquidity

pools, respectively. Each panel shows the overlay of histograms and empirical CDF for all three time spans. Charts show that the distribution of returns are almost identical for each time span.[41]

The BDS test is applied to compare degrees of predictability throughout all time spans.[42] Its advantage is that it tests the null hypothesis of no dependence in the data against any possible alternative. In this way the approach of the BDS test is similar to the ATA, which also looks for any possible dependence in the data. Since the ATA uses rolling subsamples of observations to perform forecasts the

---

[40]Returns are unit-free, which makes them easy to compare and aggregate.

[41]Here and in the comparison of the BDS test statistic distributions in different time spans the analysis omits the results of the statistical tests. Even though the Kolmogorov-Smirnov test rejected the hypothesis of identical distributions the high number of observations present in each sample (on average 27000 observations) the test would detect slightest difference between samples. Similar problems were observed with other nonparametric tests, which were used to compare the distribution of samples.

[42]For details on BDS test see [26, 81].
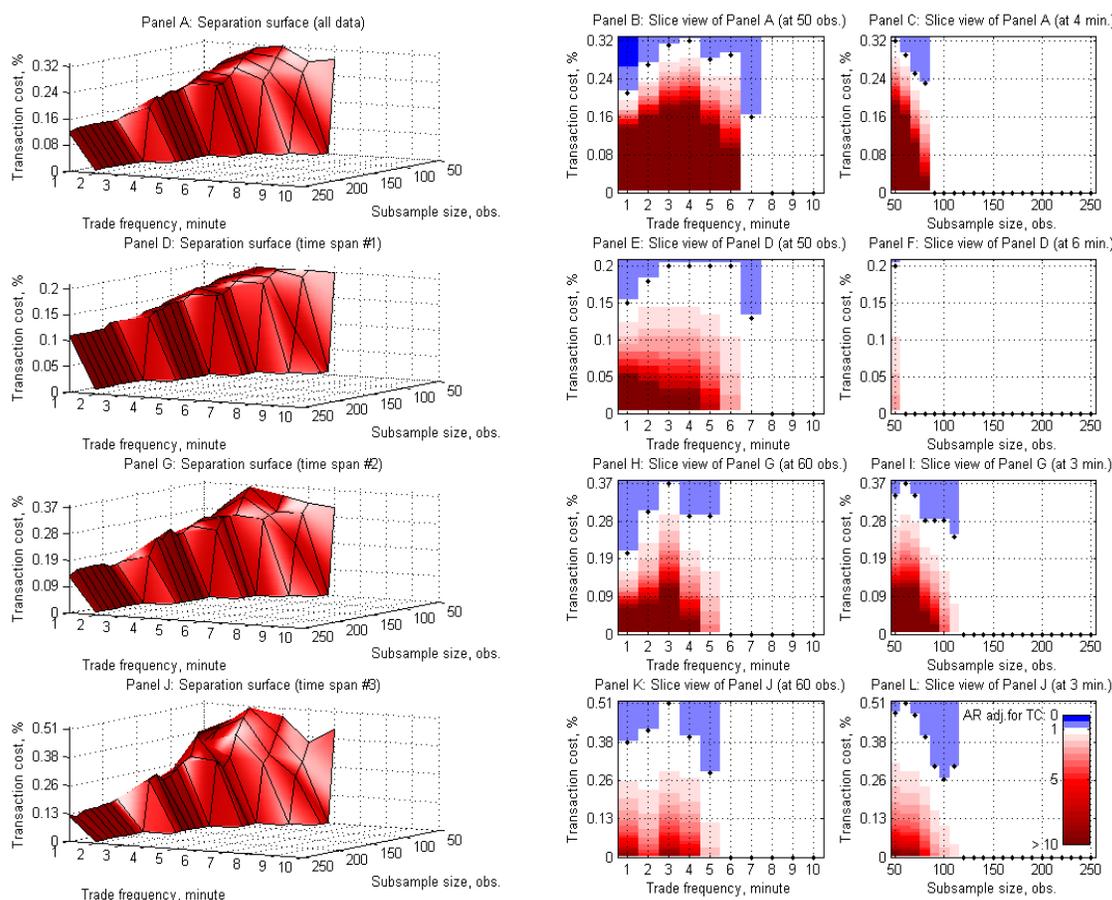
Figure 4. Application of ATA trading strategies on the pool of medium liquidity stocks

same subsamples are used in the BDS test. The histograms and empirical CDF of the BDS test statistics are presented in the last two rows of Figure 6. In the high- as well as in the medium-liquidity pools the histograms and empirical CDF are identical for all time spans.

The identity of the return distributions and degree of predictability clearly supports the second explanation. The increase in the value of maximum breakeven transaction costs is indeed due to ATA learning (see Table 4).

The conclusion is that the ATA is able to extract positive returns adjusted for and conditional on transaction costs by processing past prices. The returns, proxied by the maximum breakeven transaction costs, are increasing through time, which demonstrates the ability of the ATA to master technical analysis.

9. Conclusions

The paper presented a study of market efficiency from the viewpoint of artificial technical analyst. The artificial technical analyst had been designed to replicate the expertise of technical analysts in a robust and objective way. Its application to the Tokyo Stock Exchange data under a variety of transaction costs revealed that when transaction costs were sufficiently low the market could be perceived as inefficient and technical analysis could be rewarded with substantial profits, while under higher transaction costs the perception of market efficiency prevailed.

The efficient market hypothesis is one of the central elements of financial theory. Coined by Harry Roberts and formalized be Eugene Fama it postulates the impossibility of gaining profit on the base of information contained in past prices or returns of financial securities. In spite of its prominent role in financial theory empirical studies persistently
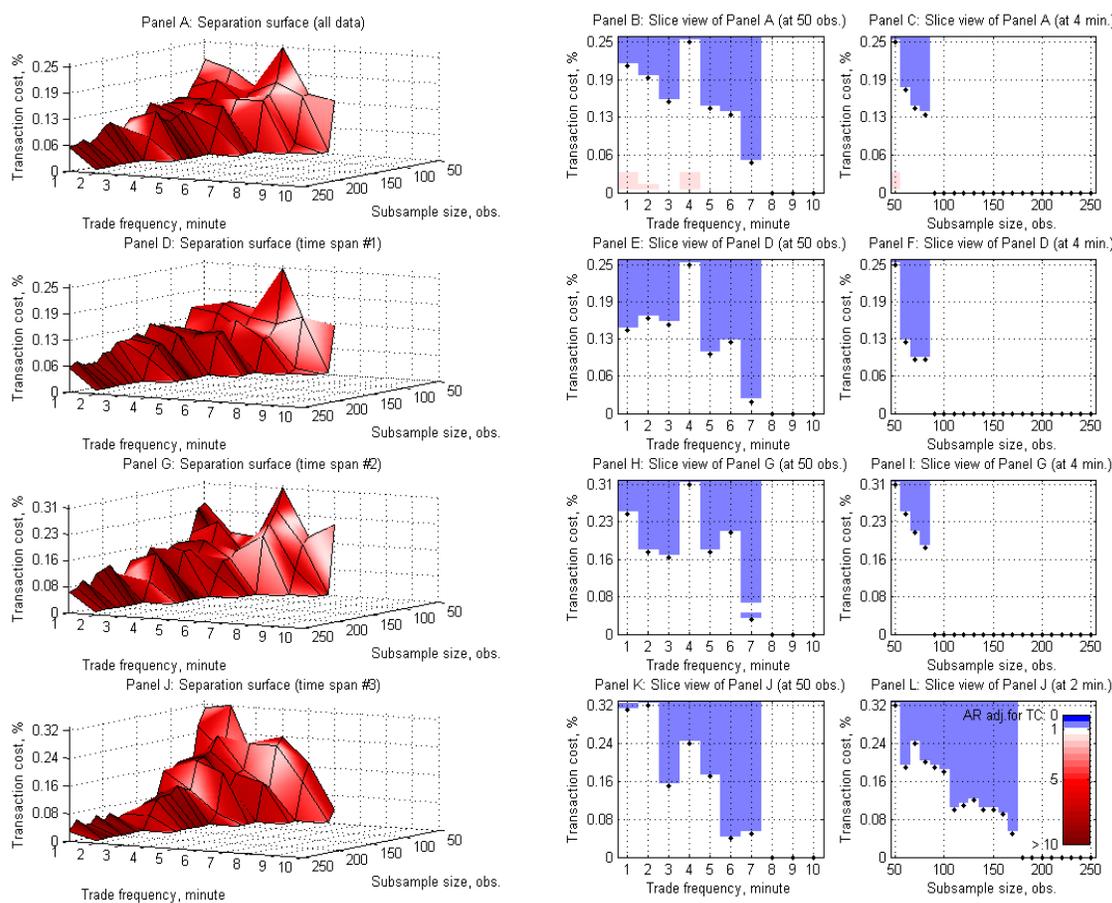
Figure 5. Application of ATA trading strategies on the pool of low liquidity stocks

show the violation of market efficiency. Moreover, the technique used by financial practitioners, technical analysis, directly contradicts the EMH.

In this study we analyzed the time series of high-frequency stock prices, collected on the Tokyo Stock Exchange over the time period from 11/03/1996 to 23/06/1998. The pool of our time series covers 2273 stocks listed on the TSE. Considering computational intensity only 30 time series were selected for the study. For this all stocks were sorted into 3 pools according to their liquidity, which can be described as low-, medium- and high-liquidity pools of stocks. In each pool 10 stocks were randomly selected. Their corresponding time series of prices were used for computations. Before reporting the computed results are averaged within a corresponding liquidity pool.

The application of the ATA allows to extract significant returns conditional on the value of transaction costs. The studied time series show that under the transaction costs of

less than 0.32 percent the application of technical analysis produces a gain. In specific time spans gains are possible even under the transaction costs of less than 0.51 percent of the transaction value.

The credibility of these results is supported by the way the ATA processes data. The ATA is an online learning technique, which does not have classic in-sample training. Each iteration is an out-of-sample analysis of a time series. Moreover, it operates only with a fixed number of recent observations. Information about the past is stored as patterns in the memory of the ATA. The ability of the ATA to derive positive returns even after adjusting for transaction costs is a clear indication of past predictability of time series of prices.

One can compare obtained values of breakeven transaction costs with values derived in other markets. Hendrik Bessembinder and Kalok Chan [19, 20] used the same rules and parameter combinations as Brock et al. [25] to esti-
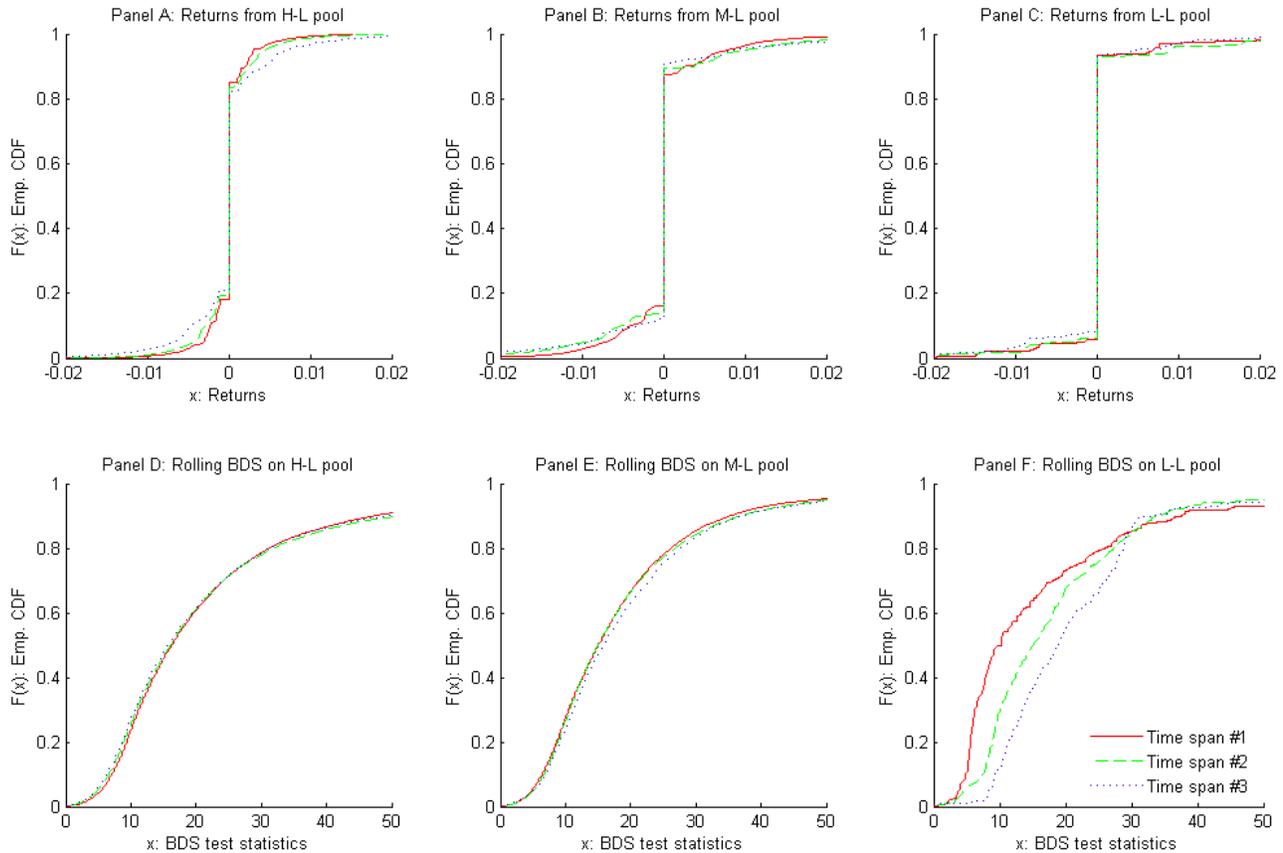
Figure 6. Time span distribution analysis of returns and rolling BDS test

		Stock pool liquidity		
		High	Medium	Low
All	TC, %	0.28	0.32	0.25
time	@ obs.	50	50	50
spans	@ min.	3	4	4
Time	TC, %	0.15	0.20	0.25
span	@ obs.	50	50	50
# 1	@ min.	4	6	4
Time	TC, %	0.20	0.37	0.31
span	@ obs.	50	60	50
# 2	@ min.	4	3	4
Time	TC, %	0.33	0.51	0.32
span	@ obs.	50	60	50
# 3	@ min.	3	3	2

Table 4. Application of the ATA trading strategy

All time spans cover data from 11/03/1996 till 22/06/1998. Time spans #1, #2 and #3 are 11/03/1996-13/12/1996, 14/12/1996-17/09/1997, and 18/09/1997-22/06/1998, respectively. Each time span on average covers 277 days.

mate breakeven transaction costs for the DJIA. For the period from 1926 to 1991 the estimated value is 0.39 percent, while for a more recent period, from 1976 to 1991, it was reduced to 0.22 percent. With optimized parameters of Brock et al. [25] methodology Stephen Taylor [152] estimated breakeven transaction costs of 1.07 percent for the DJIA, from 1968 to 1988. For twelve UK firms the value of breakeven transaction costs is 0.08 percent for the period from 1972 to 1991. More recent work of Stephen Taylor [153] reports that the average breakeven transaction costs equals to 0.31 percent for sixteen financial series (including US stock indexes, individual US stocks, and commodities), and in a range from 0.59 to 2.25 percent for four currency exchange rate series (DM/USD, GBP/USD, SF/YEN, and YEN/USD) over the same period. Thus, from the point of view of breakeven transaction costs the degree of TSE market inefficiency is comparable to US stock market.

To conclude, the application of the ATA strategy shows weak inefficiency of the TSE observed on the historical data. The results are genuine since the analysis of individual time spans did not show significant changes in the data. The degree of inefficiency is conditional on the level

of transaction costs. Under the low transaction costs technical analysts can perceive the market as inefficient, while at the higher transaction costs the perception of efficiency prevails.

References

[1] Acar, E., *Expected Returns of Directional Forecasters*, in Advanced Trading Rules, ed. E. Acar and S. E. Satchell, 2nd ed., Oxford: Butterworth Heinemann, pp. 51-80, 2002.

[2] Agrawal, A., J. F. Jaffe, G. N. Mandelker, *The Post-Merger Performance of Acquiring Firms: A Reexamination of an Anomaly*, Journal of Finance 47, pp. 1605-1621, 1992.

[3] Alexander, S. S., *Price Movements in Speculative Markets: Trends or Random Walks?*, Industrial Management Review 2, pp. 7-26, 1961.

[4] Alexander, S. S., *Price Movements in Speculative Markets: Trends or Random Walks*, in Cootner, P. (eds.) The Random Character of Stock Market Prices, Vol. 2, MIT Press, Cambridge, pp. 338-372, 1964.

[5] Allen, H. and M. P. Taylor, *Charts, Noise and Fundamentals in the London Foreign Exchange Market*, The Economic Journal 100, pp. 49-59, 1990.

[6] Allen, F. and R. Karjalainen, *Using Genetic Algorithms to Find Technical Trading Rules*, Journal of Financial Economics 51 (2), pp. 245-271, 1999.

[7] Amihud, Y. and H. Mendelson, *Volatility, Efficiency, and Trading: Evidence from the Japanese Stock Market*, Journal of Finance 46 (5), pp. 1765-1789, 1991.

[8] Armano, G., A. Murru and F. Roli, *Stock Market Prediction by a Mixture of Genetic-Neural Experts*, International Journal of Pattern Recognition and Artificial Intelligence 16 (5), pp. 501-526, 2002.

[9] Arthur, W. B., J. H. Holland, B. LeBaron, R. Palmer and P. Tayler, *Asset Pricing Under Endogenous Expectations in an Artificial Stock Market*, The Economy as an Evolving Complex System II, Eds. Arthur, Durlauf, and Lane SFI Studies in the Science of Complexity, vol. XXVII, Addison-Wesley, pp. 12-44, 1997.

[10] Asquith, P., *Merger Bids, Uncertainty and Stockholder Returns*, Journal of Financial Economics 11, pp. 51-83, 1983.

[11] Atkins, A. B. and E. A. Dyl, *Price Reversals, Bid-Ask Spreads and Market Efficiency*, Journal of Financial and Quantitative Analysis 25 (4), pp. 535-547, 1990.

[12] Bachelier, L. , *Theory of Speculation*, A thesis presented to the Faculty of Sciences of the Academy of Paris, Annales de l'Ecole Normal Superieure 27, pp. 21-86, 1900.

[13] Ball, R. and S. P. Kothari, Nonstationary Expected Returns: Implications for Tests of Market Efficiency and Market Correlation in Returns, Journal of Financial Economics 25, pp. 51-74, 1989.

[14] Batten, J. and C. Ellis, *Technical Trading System Performance in the Australian Share Market: Some Empirical Evidence*, Asia Pacific Journal of Management 13 (1), pp. 87-99, 1996.

[15] Bauwens, L., *Econometric Analysis of Intra-daily Trading Activity on Tokyo Stock Exchange*, IMES Discussion Paper, No. 2005-E-3, Institute for Monetary and Economic Studies, Bank of Japan, 2005.

[16] Beaver, W. H., *Information Efficiency*, Accounting Review 56, pp. 23-27, 1980.

[17] Beja, A. and M. B. Goldman, *On the Dynamic Behavior of Prices in Disequilibrium*, Journal of Finance 35 (2), pp. 235-248, 1980.

[18] Bekaert, G. and R. J. Hodrick, *Characterizing Predictable Components in Equity and Foreign Exchange Rates of Return*, Journal of Finance 47, pp. 467-509, 1992.

[19] Bessembinder, H. and K. Chan, *The Profitability of Technical Trading Rules in the Asian Stock Markets*, Pacific-Basin Finance Journal 3, pp. 257-284, 1995.

[20] Bessembinder, H. and K. Chan, *Market Efficiency and the Returns to Technical Analysis*, Financial Management 27 (2), pp. 5-17, 1998.

[21] Black, F. and M. Scholes, *The Pricing of Options and Corporate Liabilities*, Journal of Political Economy 81, pp. 637-659, 1973.

[22] Blume, L., D. Easley and M. O'Hara, *Market Statistics and Technical Analysis: The Role of Volume*, Journal of Finance 49, pp. 153-181, 1994.

[23] Bodurtha, J. N. and G. R. Courtadon, *Efficiency Tests of the Foreign Currency Options Market*, The Journal of Finance 41 (1), pp. 151-162, 1986.

[24] Bollerslev, T. and R. J. Hodrick, *Financial Market Efficiency Tests*, NBER Working Paper Series, No. 4108, Cambridge: National Bureau of Economic Research, 1992.

[25] Brock, W., J. Lakonishok and B. LeBaron, *Simple Technical Trading Rules and the Stochastic Properties of Stock Returns*, Journal of Finance 47 (5), pp. 1731-1764, 1992.

[26] Brock, W., D. Dechert, J. Scheinkman and B. LeBaron, *A Test for Independence Based on the Correlation Dimension*, Econometric Review 15 (3), pp. 197-235, 1996.

[27] Brown, D. P. and R. H. Jennings, *On Technical Analysis*, Review of Financial Studies 2 (4), pp. 527-551, 1989.

[28] Butz, M. V., *Rule-Based Evolutionary Online Learning Systems: Learning Bounds, Classification, and Prediction*, IlliGAL Report No. 2004034, 2004.

[29] Butz, M. V., *Documentation of XCS+TS C-Code 1.2*, IlliGAL Report No. 2003023, 2003.

[30] Campbell, J. Y., *A Variance Decomposition of Stock Returns*, Economic Journal 101, pp. 152-179, 1991.

[31] Campbell, J. Y. and R. Shiller, *The Dividend-Price Ration and Expectations of Future Dividends and Discount Factors*, Review of Financial Studies 1, pp. 195-228, 1988.

[32] Chan, K. C., *On the Contrarian Investment Strategy*, Journal of Business 61, pp. 147-163, 1988.

[33] Chung, Y. P., *A Transaction Date Test of Stock Index Futures Market Efficiency and Index Arbitrage Profitability*, Journal of Finance 46 (5), pp. 1791-1809, 1991.

[34] Clyde, W. C. and C. L. Osler, *Charting: Chaos Theory in Disguise?*, Journal of Futures Markets 17 (5), pp. 489-514, 1997.

[35] Conrad, J. and G. Kaul, *Time-Variation in Expected Returns*, Journal of Business 61, pp. 409-425, 1988.

[36] Cootner, P. A. *Refinement and Empirical Testing: Introduction*, in Cootner (1967), The Random Character of Stock Market Prices, Cambridge: MIT Press, 1967.

[37] Cowles, A., *Can Stock Market Forecasters Forecasts?*, Econometrica 1, pp. 309-324, 1933.

[38] Cumby, R. E. and J. Huizinga, *Testing the Autocorrelation Structure of Disturbance in Ordinary Least Squares and Instrumental Variables Regressions*, Econometrica 60, pp. 185-195, 1992.

[39] Dacorogna, M., R. Gencay, U. A. Mueller, R. B. Olsen and O. V. Pictet, *An Introduction to High-Frequency Finance*, Academic Press, 2001.

[40] Dawson, S. M., *Singapore Share Recommendations Using Technical Analysis*, Asia Pacific Journal of Management 2 (3), pp. 180-188, 1985.

[41] Dawson, E. R. and J. M. Steeley, *On the Existence of Visual Patterns in the UK Stock Market*, Journal of Business Finance and Accounting 30, pp. 263-293, 2003.

[42] DeBondt, W. F. M. and R. H. Thaler, *Does the Stock Market Overreact?*, Journal of Finance 40, pp. 793-805, 1985.

[43] DeBondt, W. F. M. and R. H., Thaler, *Further Evidence on Investors Overreaction and Stock Market Seasonality*, Journal of Finance 42, pp. 557-581, 1987.

[44] DeLong, J. B., A. Shleifer, L. H. Summers and R. J. Waldmann, *Positive Feedback Investment Strategies and Destabilizing Rational Speculation*, Journal of Finance 45, pp. 379-395, 1990.

[45] DeLong, J. B., A. Shleifer, L. H. Summers and R. J. Waldmann, *Noise Trader Risk in Financial Markets*, Journal of Political Economy 98 (4), pp. 703-738, 1990.

[46] Diebold, F. X., *Testing for Serial Correlation in the Presence of ARCH*, Proceedings of the American Statistical Association, Business and Economic Statistics Association, pp. 323-328, 1986.

[47] Dimson, E. and M. Mussavian, A Brief History of Market Efficiency, European Financial Management 4 (1), pp. 91-193, 1998.

[48] Edwards, R. D., Magee, J. and W. H. C. Bassetti, *Technical Analysis of Stock Trends*, 8th ed., St. Lucie Press, 2001.

[49] Fama, E. F., *The Behavior of Stock Market Prices*, Journal of Business 38, pp. 34-105, 1965.

[50] Fama, E. F., *Efficient Capital Market: A Review of Theory and Empirical Work*, Journal of Finance 25, pp. 383-417, 1970

[51] Fama, E. F., *Efficient Capital Markets: II*, The Journal of Finance 46 (5), pp. 1575-1617, 1991.

[52] Fama, E. F., *Market Efficiency, Long-Term Returns and Behavioral Finance*, Journal of Financial Economics 49, pp. 283-306, 1998.

[53] Fama, E. F. and M. E. Blume, *Filter Rules and Stock Market Trading. Security Prices: A Supplement*, Journal of Business 39, pp. 191-225, 1966.

[54] Fama, E. F. and K. R. French, *Permanent and Temporary Components of Stock Prices*, Journal of Political Economy 96, pp. 246-273, 1988.

[55] Fama, E. F. and K. R. French, *Dividend Yields and Expected Stock Returns*, Journal of Financial Economics 22, pp. 427-465, 1988.

[56] Fama, E. F. and W. Schwert, *Asset Returns and Inflation*, Journal of Financial Economics 5, pp. 115-146, 1977.

[57] Fang, Y. and D. Xu, *The Predictability of Asset Returns: An Approach Combining Technical Analysis and Time Series Forecasts*, International Journal of Forecasting 19, pp. 369-385, 2003.

[58] Fernandez-Rodriguez, F., S. Sosvilla-Rivero and M. D. Garcia-Artiles, *Using Nearest-Neighbour Predictors to Forecast the Spanish Stock Market*, Investigaciones Economicas 21 (1), pp. 75-91, 1997.

[59] Fernandez-Rodriguez, F., C. Gonzalez-Martel and S. Sosvilla-Rivero, *On the Profitability of Technical Trading Rules Based on Artificial Neural Networks: Evidence from the Madrid Stock Market*, Economic Letters 69, pp. 89-94, 2000.

[60] Fisher, L., *Some New Stock-Market Indexes*, Journal of Business 39, pp. 191-225, 1966.

[61] Flavin, M., *Excess Volatility in Financial Markets: A Reassessment of the Empirical Evidence*, Journal of Political Economy 91, pp. 929-956, 1983.

[62] Fock, J. H., C. Klein and B. Zwergel, *Performance of Candlestick Analysis on Intraday Futures Data*, Journal of Derivatives, pp. 28-40, Fall 2005.

[63] Giles, C. L., S. Lawrence and A. C. Tsoi, *Noisy Time Series Prediction Using a Recurrent Neural Network and Grammatical Inference*, Machine Learning 44 (1/2), pp. 161-183, 2001.

[64] Gilles, C. and S. F. LeRoy, *Econometric Aspects of the Variance-Bounds Tests: A Survey*, Review of Financial Studies 4, pp. 753-791, 1991.

[65] Goldberg, D. E., *Genetic Algorithms in Search, Optimization, and Machine Learning*, Reading MA: Addison-Wesley, 1989.

[66] Grossman, S. J. and R. J. Shiller, *The Determinants of the Variability of Stock Market Prices*, American Economic Review 71, pp. 222-227, 1981.

[67] Grossman, S. J. and J. Stiglitz, *On the Impossibility of Informationally Efficient Markets*, American Economic Review 70, pp. 393-408, 1980.

[68] Hansen, L. P., *Large Sample Properties of Generalized Method of Moments Estimators*, Econometrica 50, pp. 111-120, 1982.

[69] Hodrick, R. J., *Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement*, Review of Financial Studies 5 (3), pp. 357-386, 1992.

[70] Holland, J. H., *Adaptation in Natural and Artificial Systems*, Ann Arbor, MI: University of Michigan Press, 1975.

[71] Holland, J. H. and J. S. Reitman, *Cognitive Systems Based on Adaptive Algorithms*, in D. A. Waterman and F. Hayes-Roth (Eds.), Pattern-directed Inference Systems, New York: Academic Press, pp. 313-329, 1978.

[72] Holmes, J. H., Lanzi, P. L., Stolzmann, W., and Wilson, S. W., *Learning Classifier Systems: New Models, Successful Applications*, Information Processing Letters 82 (1), pp. 23-30, 2002.

[73] Houthakker, H., *Systematic and Random Elements in Short Term Price Movements*, American Economic Review 51, pp. 164-172, 1961.

[74] Hundson, R., M. Dempsey and K. Keasey, *A Note on the Weak Form Efficiency of Capital Markets: The Application of Simple Technical Trading Rules to UK Stock Prices - 1935 to 1994*, Journal of Banking and Finance 20, pp. 1121-1132, 1996.

[75] Hussman, J. P., *Time Variation in Market Efficiency: A Mixture-of-Distribution Approach*, Hussman Funds, mimeo, 1998.

[76] Jegadeesh, N., *Evidence of Predictable Behavior of Security Returns*, Journal of Finance 45, pp. 881-898, 1990.

[77] Jegadeesh, N. and S. Titman, *Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency*, Journal of Finance 48, pp. 65-91, 1993.

[78] Jensen, M., *Some Anomalous Evidence Regarding Market Efficiency*, Journal of Financial Economics 6, pp. 95-101, 1978.

[79] Kandel, S. and R. Stambaugh, *On the Predictability of Stock Returns: An Asset Allocation Perspective*, Journal of Finance 51, pp. 385-424, 1996.

[80] Kanzler, L., *A Study of the Efficiency of the Foreign Exchange, Through Analysis of Ultra-high Frequency Data*, Ph.D. Dissertation, Oxford University, 1998.

[81] Kanzler, L., *Very Fast and Correctly Sized Estimation of the BDS Statistic*, Department of Economics of Oxford University, Available at http://users.ox.ac.uk/~econlrk, 1999.

[82] Kendall, M.G., *The Analysis of Economic Time Series – Part 1: Prices*, Journal of Royal Statistical Society (Series A) 96, pp. 11-25, 1953.

[83] Keogh, E., S. Chu, D. Hart and M. Pazzani, *An online algorithm for segmenting time series*, in Proceedings of the IEEE International Conference on Data Mining, pp. 289-296, 2001.

[84] Kim, M. J., C. R. Nelson and R. Startz, *Mean Reversion in Stock Prices? A Reappraisal of the Empirical Evidence*, Review of Economic Studies 58, pp. 515-528, 1991.

[85] Kleidon, A. W., *Variance Bounds Tests and Stock Price Valuation Models*, Journal of Political Economy 94, pp. 953-1001, 1986.

[86] Kleidon, A. W., *Bias in Small Sample Tests of Stock Price Rationality*, Journal of Business 59 (2), pp. 237-262, 1986.

[87] Korczak, J. and P. Roger, *Stock Timing Using Genetic Algorithms*, Applied Stochastic Models in Business and Industry 18, pp. 121-134, 2002.

[88] Krishnakumar, K., *Micro-Genetic Algorithms for Stationary and Non-Stationary Function Optimization*, SPIE Proceedings: Intelligent Control and Adaptive Systems, 1196, pp. 289-296, 1989.

[89] Kyle, A., *Continuous Auctions and Insider Trading*, Econometrica 53, pp. 1315-1335, 1985.

[90] Ilinskaia, A. and K. Ilinski, *How to Reconcile Market Efficiency and Technical Analysis*, http://xxx.lanl.gov/abs/cond-mat/9902044, 1999.

[91] Lam, M., *Neural Network Techniques for Financial Performance Prediction: Integrating Fundamental and Technical Analysis*, Decision Support Systems 37, pp. 567-581, 2004.

[92] Lanzi, P. L., W. Stolzmann, S. W. Wilson (Eds.), *Learning Classifier Systems: From Foundations to Applications*, Vol. 1813 of LNAI, Berlin: Springer-Verlag, 2000.

[93] LeBaron, B., *Technical Trading Rule Profitability and Foreign Exchange Intervention*, Journal of International Economics 49 (1), pp. 125-143, 1999.

[94] LeBaron, B., *Technical Trading Rules and Regime Shifts in Foreign Exchange*, in Advanced Trading Rules, edited by E. Acar and S. Satchell , 2nd Edition, Butterworth-Heinemann, pp. 5-40, 2002.

[95] Lehmann, B. N. and D. M. Modest, *Trading and Liquidity on the Tokyo Stock Exchange: A Bird's Eye View*, Journal of Finance 49 (3), pp. 951-984, 1994.

[96] Lesmond, D. A., J. P. Ogden and C. A. Trzcinka, *A New Estimate of Transaction Costs*, Review of Financial Studies 12, pp. 1113–1141, 1999.

[97] Lesmond, D. A., M. J. Schill and C. Zhou, *The Illusory Nature of Momentum Profits*, Journal of Financial Economics 71, pp. 349–380, 2004.

[98] LeRoy, S. F. and R. D. Porter, The Present-Value Relation: Test Based on Implied Variance Bounds, Econometrica 49, pp. 555-574, 1981.

[99] Levich, R. M. and L. R. Thomas, *The Significance of Technical Trading-Rule Profits in the Foreign Exchange Market: A Bootstrap Approach*, Journal of International Money and Finance 12 (5), pp. 451-474, 1993.

[100] Lo, A. W. and A. C. MacKinlay, *Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test*, Review of Financial Studies 1 (1), pp. 41-66, 1988.

[101] Lo, A. W., H. Mamaysky and J. Wang, *Foundations of Technical Analysis: Computational Algorithms, Statistical Inference and Empirical Implementation*, Journal of Finance 55, pp. 1705-1765, 2000.

[102] Lukac, L. P., B. W. Brorsen and S. H. Irwin, *A Test of Futures Market Disequilibrium Using Twelve Different Trading Systems*, Applied Economics 20, pp. 623-639, 1988.

[103] Lukac, L. P., B. W. Brorsen, *A Comprehensive Test of Futures market Disequilibrium*, Financial Review 25, pp. 593-622, 1990.

[104] Maillet, B. and T. Michel, *Technical Analysis Profitability when Exchange Rates are Pegged: A Note*, European Journal of Finance 11 (6), pp. 463-470, 2005.

[105] Malkiel, B., *Efficient Market Hypothesis*, in Newman, P., M. Milgate, and J. Eatwell (eds.), New Palgrave Dictionary of Money and Finance, London: Macmillan, 1992.

[106] Mandelbrot, B., *The Variation of Certain Speculative Prices*, Journal of Business 36, pp. 394-419, 1963.

[107] Mankiw, N. G., D. Romer and M. D. Shapiro, *Stock Market Forecastability and Volatility: A Statistical Appraisal*, Review of Economic Studies 58, pp. 455-477, 1991.

[108] Matilla-Garcia, M., *Are Trading Rules Based on Genetic Algorithms Profitable?*, Applied Economic Letters 13, pp. 123-126, 2006.

[109] Michalewicz, Z., *Genetic Algorithms + Data Structures = Evolution Programs*, 3rd Ed., Berlin: Springer-Verlag, 1999.

[110] Mills, T. C., *Technical Analysis and the London Stock Exchange: Testing Trading rules Using the FT30*, International Journal of Finance & Economics 2, pp. 319-331, 1997.

[111] Mitchell, M. and E. Stafford, *Managerial Decisions and Long-term Stock Price Performance*, Journal of Business 73, pp. 287-329, 2000.

[112] Mitra, S. K., *Profiting from Technical Analysis in Indian Stock Market*, Finance India 16 (1), pp. 109-120, 2002.

[113] Morrison, A. D. and N. Vulkan, *Making Money out of Publicly Available Information*, Economics Letters 89 (1), pp. 31-38, 2005.

[114] Murphy, J. J., *Technical Analysis of the Future Markets*, New York: Prentice Hall, 1986.

[115] Murphy, J. J., *Technical Analysis of the Financial Markets: A Comprehensive Guide to Trading Methods and Applications*, New York: Prentice Hall, 1999.

[116] Neely, C. J., P. Weller and R. Dittmar, *Is Technical Analysis in the Foreign Exchange Market Profitable? A Genetic Programming Approach*, Journal of Financial and Quantitative Analysis 32 (4), pp. 405-426, 1997.

[117] Neely, C. J., *Technical Analysis in the Foreign Exchange Market: A Layman's Guide*, Federal Reserve Bank of St. Louis: Review, pp. 23-38, September 1997.

[118] Neely, C. J., *Technical Analysis and Profitability of US Foreign Exchange Intervention*, Federal Reserve Bank of St. Louis: Review, pp. 3-17, July 1998.

[119] Neftci, S. N. and A. J. Policano, *Can Chartists Outperform the Market? Market Efficiency Tests for Technical Analysis*, Journal of Futures Markets 4 (4), pp. 465-478, 1984.

[120] Nelson, C. R., *Inflation and Rates of Return on Common Stocks*, Journal of Finance 31, pp. 471-483, 1976.

[121] Nison, S., *Japanese Candlestick Charting Techniques: A Contemporary Guide to the Ancient Techniques of the Far East*, NYIF, 2nd Ed., 2001.

[122] Omrane, W. B. and H. Van Oppens, *The Performance Analysis of Chart Patterns: Monte Carlo Simulation and Evidence from the Euro/Dollar Foreign Exchange Market*, Empirical Economics 30, pp. 946-971, 2006.

[123] Osborne, M. F. M., *Some Quantitative Tests for Stock Price Generality Models and Trading Folklore*, Journal of the American Statistical Association 62 (318), pp. 321-340, 1967.

[124] Osler, C. L and P. H. K. Chang, *Head and Shoulders: Not Just a Flaky Patter*, Federal Reserve Bank of New York Staff Report, No. 4, Federal Reserve Bank of New York, 1995.

[125] Pearson, K., *The Problem of the Random Walk*, Nature 72 (1865), p. 294, 1905.

[126] Pearson, K., *The Problem of the Random Walk*, Nature 72 (1867), p. 342, 1905.

[127] Poterba, J. and L. H. Summers, *Mean Reversion in Stock Prices: Evidence and Implications*, Journal of Financial Economics 22, pp. 27-59, 1988.

[128] Potvin, J.-Y., P. Soriano and M. Vallee, *Generating Trading Rules on the Stock Markets with Genetic Programming*, Computers & Operations Research 31, pp. 1033-1047, 2004.

[129] Pring, M. J., *Technical Analysis Explained: the Successful Investor's Guide to Spotting Investment Trends and Turning Points*, 2nd Ed., New York: McGraw Hill, 1985.

[130] Richardson, M. and J. H. Stock, *Drawing Inferences from Statistics Based on Multi-Year Asset Returns*, Journal of Financial Economics 25, pp. 323-348, 1989.

[131] Rhea, R., *The Dow Theory*, Fraser Publishing, 1993.

[132] Roberts, H., *Stock-Market Patterns and Financial Analysis: Methodological Suggestions*, Journal of Finance 14, pp. 1-10, 1959.

[133] Roberts, H., *Statistical versus Clinical Prediction of the Stock Market*, unpublished manuscript, 1967.

[134] Roberts, M. C., *Technical Analysis and Genetic Programming: Constructing and Testing a Commodity Portfolio*, Journal of Futures Markets 25 (7), pp. 643-660, 2005.

[135] Roll, R., *A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market*, Journal of Finance 39 (4), pp. 1127-1139, 1984.

[136] Roll, R., *The Hubris Hypothesis of Corporate Takeovers*, Journal of Business 59, pp. 197-216, 1986.

[137] Rubinstein, M., *Security Market Efficiency in an Arrow-Debreu Economy*, American Economic Review 65, pp. 812-824, 1975.

[138] Sehgal, S. and A. Garhyan, *Abnormal Returns Using Technical Analysis: The Indian Experience*, Finance India 16 (1), pp. 181-203, 2002.

[139] Shiller, R. J., *The Volatility of Long-Term Interest Rates and Expectations Models of the Term Structure*, Journal of Political Economy 87, pp. 1190-1219, 1979.

[140] Shiller, R. J., *Do Stock Prices Move too much to be Justified by Subsequent Changes in Dividends?*, American Economic Review 71, pp. 421-436, 1981.

[141] Shiller, R. J., *Stock Prices and Social Dynamics*, Brookings Papers on Economic Activity 2, pp. 457-510, 1984.

[142] Schleifer, A., *Inefficient Markets : An Introduction to Behavioral Finance*, Oxford University Press, 2000.

[143] Skouras, S., *Financial Returns and Efficiency as Seen by an Artificial Technical Analyst*, Journal of Economic Dynamics and Control 25 (1-2), pp. 213-44, 2001.

[144] Sullivan, R., Timmerman A. and H. White, *Data-Snooping, Technical Trading Rule Performance and the Bootstrap*, CEPR Discussion Paper, No. 1976, Center for Economic Policy Research, Washington D.C., 1998.

[145] Summers, L. H., *Does the Stock Market Rationally Reflect Fundamental Values?*, Journal of Finance 41, pp. 591-601, 1986.

[146] Sweeney, R. J., *Beating the Foreign Exchange Market*, Journal of Finance 41, pp. 163-182, 1986.

[147] Sweeney, R. J., *Some New Filter Rule Tests: Methods and Results*, Journal of Financial and Quantitative Analysis 23, pp. 285-300, 1988.

[148] Taylor, S. J., *Strategies for Investors in Apparently Inefficient Futures Markets*, in Futures Markets - Modelling, Managing and Monitoring Futures Trading, ed. M. E. Streit, Oxford: Basil Blackwell, pp. 165-198, 1983.

[149] Taylor, S. J., *Modelling Financial Time Series*, Chichester: Wiley, 1986.

[150] Taylor, S. J., *Rewards Available to Currency futures Speculators: Compensation for Risk or Evidence of Inefficient Pricing?*, Economic Record 68 (supplement), pp. 105-116, 1992.

[151] Taylor, M. P., *The Use of Technical Analysis in the Foreign Exchange Market*, Journal of International Money and Finance 11, pp. 304-314, 1992.

[152] Taylor, S. J., *Stock Index and Price Dynamics in the UK and the US: New Evidence from a Trading Rule and Statistical Analysis*, European Journal of Finance 3, pp. 39-69, 2000.

[153] Taylor, S. J., *Asset Price Dynamics, Volatility, and Prediction*, Princeton University Press, 2005.

[154] Thorley, S., *The Inefficient Markets Argument for Passive Investing*, The Marriott School at BYU, mimeo, 1998.

[155] Tokyo Stock Exchange, *Tokyo Stock Exchange Fact Book 2006*, 2006.

[156] Tucker, A. L., *Empirical Tests of the Efficiency of the Currency Option Market*, The Journal of Financial Research 8 (4), pp. 275-285, 1985.

[157] Walter, C., *The Efficient Market Hypothesis, the Gaussian Assumption, and the Investment Management Industry*, Institut d'Etudes Politiques, mimeo, 2003.

[158] Weintraub, R. E., *On Speculative Prices and Random Walks: A Denial*, Journal of Finance 18 (1), pp. 59-66, 1963.

[159] Wilson, S. W., *Classifier Fitness Based on Accuracy*, Evolutionary Computation 3, 1995.

[160] Wilson, S. W., *Generalization in the XCS Classifier System*, 1998.

[161] Wong, W.-K., M. Manzur and B.-K. Chew, *How Rewarding is Technical Analysis? Evidence from Singapore Stock Market*, Applied Financial Economics 13 (7), Taylor and Francis Journals, pp. 543-551, 2003.

[162] Working, H., *A Random Difference Series for Use in the Analysis of Time Series*, Journal of American Statistical Association 29 (185), pp. 11-24, 1934.