Improved Pricing on the Stock Market with Trading Agents

Wah-Sui Almberg

Master's Thesis [1]
Department of Computer and Systems Sciences
Stockholm University / Royal Institute of Technology
March 2002

Abstract

I have investigated the feasibility of letting artificial trading agents be a part and function of the servers of a stock market. As a tentative beginning I chose to examine a particular type of trading strategy inspired by a paradox that the physicist Juan M.P. Parrondo elucidated recently. The main incentive was to find a way for clients hiring these trading agents to earn money, or at least retain invested capital, on a receding market and at the same time subdue sharply falling rates on the market as a whole. For the trading agent I designed there was a need to make a rough estimate of the probability for an eligible stock to either gain or lose money in the short term. For this purpose I designed a subagent with two adjustable parameters inspired by the mathematical notion of the second derivative. This subagent had interesting properties that were probably due to some type of 'herd' mentality prevalent on the stock market. The test results I obtained showed that the trading agent had on average 9.5 per cent more capital after a period of three months than the benchmark strategy I used had on a receding market. I found it very hard to draw any conclusions regarding the agent's core intelligence. The results I have obtained have academic value, but for practical purposes a lot remains to be done.

[1] This thesis corresponds to 20 weeks of full-time work.

1. Introduction

1.1 Background

It is not a part of the essence of a computer to understand what it is doing: every action a computer performs must be explicitly anticipated and planned for. So, we are happy to accept computers as obedient tools. However, for an increasingly large number of applications we require systems that can, to some extent, decide for themselves.

A field where such applications are becoming increasingly interesting is finance. There is a fast growing literature attempting to model financial interactions using computer agents to go beyond the restrictions of analytical methods [LeBaron, 1998]. At the same time analytical approaches are closely related to this development. A continuous interaction between computational and analytical approaches is essential to the progress within the field.

In the future financial markets might very well be important areas of application for agent-based modelling. They offer features that make them very appealing to this type of modelling. One such feature, for example, is that financial data is readily available at many different frequencies from annual to minute by minute. Naturally, there are many hurdles too. Many empirical puzzles have been difficult for standard representative agent models to explain. It is a new research area, so there are still many questions that remain unanswered [*ibid.*].

From a more practical point of view one enticing question is whether agents could beat humans as traders. There has been some research conducted on this issue, and results show that humans usually lose against computer agents [Chang, 2001]. One great advantage computers have, vis-à-vis humans, is their speed. Computer agents can respond to slight changes in prices in a fraction of a second. In a test conducted by IBM, software-based robotic agents made seven per cent more cash than people did. In these tests both agents and people had the same set-up, allowing them to trade through an unbiased software-based auctioneer [Graham-Rowe, 2001]. The auction simulated a market where buyers and sellers had a fixed amount of time to trade in a single commodity. In other tests conducted by IBM, using double auctions (the same type of auctions stock markets use), agents were on average about 5 per cent more profitable than people [Chang, 2001]. According to Dr. Steve R. White – head of this research at IBM – the agents excelled, although they were programmed with rather simple strategies. This was possible because they could quickly pounce on someone else's mistake, and because they never made mindless mistakes – selling something at a loss for example – something that humans tend to do at times. These scenarios may not be likely though in 'thin' markets, where there are only a few buyers and sellers, and where good deals are a matter of skill rather than speed.

Along the above train of thought, there is also research that shows that changes to share prices do not accord well with the degree of reliability of the information reaching the stock market [Bloomfield *et al*, 2000], a condition that agents could improve upon. Another advantage computers might have, vis-à-vis humans, is that they could be designed not to fall prey to 'herd' mentality and other psychological contingencies that to a greater or lesser degree rule humans. A closer look at market statistics has shown that the distribution of the price return – the difference between the purchase and sale price of a share – is not 'Gaussian' but 'power law' – the mathematical sign that all is not at random. This suggests that brokers follow herd instincts [Haw, 2001]. This makes sense as brokers probably act on rumours that spread across the markets. Similarly, big transactions on a stock market do create a following. Indeed, simulations with an artificial stock market, with computer agents programmed to have ‘herd instincts’, have shown statistical similarities with the human counterpart [*ibid.*].

There is also research that shows that the dynamics of share prices can be described as a random walk [Osborne, 1977]. This contradicts what I have stated above. Yet other reports show that there are departures from a simple random walk in stock prices. These anomalies are statistically significant, but practically their effects are not large. From this follows that the effort I am making tests, in a way, how good an approximation the hypothesis of stock prices as a random walk is, because it is in the departure from the random walk that profits can be found.

With the above in mind, one could envision scenarios where computer agents could stabilize the stock market. Although the crash of the stock markets around the world in 1987 may have been induced or facilitated by computer aided investment strategies [Lux, 1995], computer agents could also, for example, have a stabilizing effect if they did not have the same herd instincts as human brokers. Humans and agents would interact in a complementary way that would make the stock market less volatile. In this way the danger of having to resort to more dramatic measures such as circuit breakers, trade collars, etc. (measures to counteract extreme instability, for example, a dramatic fall of an exchange) would be reduced. Considering the increasing volatility of the stock markets around the world in recent years this would be a most welcome change for institutions running exchanges.

1.2 Problem

Assignment

For an institution running a marketplace high credibility of the marketplace is of utmost importance. To achieve this end, there are many different measures that come into consideration. Among these, measures that would decrease the transaction costs are desirable [Wennerberg, 2001]. To decrease the transaction costs, one may, for example, optimize the market structure, trading functionality, and fees. I have focused on issues concerning trading functionality.

Specifically, the assignment given to me was to investigate future possibilities for trading agents to be a part and functionality of the servers of a stock market. If they like, clients should be able to connect to a moderately customizable trading agent that takes care of trading on their behalf according to their directions. These trading agents would be running on a separate server with a direct link to the server that runs the matching service. In this way all the processes of continuous evaluation of conditions, done by the agents, would not disturb the normal functionality of the matching service. The agents would have very fast and easy access to the matching service, as they would not be hindered by slow connections on the Internet or heavy network loads. Clients could place their orders with these agents rather than placing their orders, one at a time, via a network. This would be particularly advantageous when approaching heavy load situations. With the right conditions fulfilled, extreme levels of messaging/network traffic could be reduced to levels at which measures such as throttling (braking the matching service to avoid price update dissemination congestion) could be avoided. This, in turn, might increase the number of deals completed.

Furthermore, OM's wish was that these computer-aided trading agents would also enhance, if possible, the functionality of the marketplace, for example, avoiding circuit breakers (a measure to counteract a dramatic fall of an exchange), thereby increasing the profits.

Hypothesis

Although I am closer at heart to idealism or a phenomenological approach to reality, I also see the discursive advantages of realism in present day culture. Still, I find it important not to let go of our intuitive faculties approaching a subject as complex and non-analytical as the one at hand. An understanding of the stock market is an interdisciplinary effort, where a one-sided mathematical approach is quite questionable [Mirowski, 1989]. Most importantly, the human psyche comes into play, with all its far reaching possibilities, not only logically, but even reaching into the realms of transcendental experiences, not to mention earthly crafts such as poetry and art. One might even argue that to be able to gain any kind of knowledge or understanding of the psyche of man, intuition is needed [Egeberg *et al*, 1986].

Turning back to the realist approach – as a philosophical term – we commonly call science, I have designed a model, comprising an algorithm, with a heuristic intent. My hope has been to see whether this model of understanding of certain characteristics of a receding market could be plausible to some extent. More tangibly, I wanted to see whether there are trading agents that could augment the functionality of a deterministic matching service. This was the hypothesis I wanted to test empirically.

1.3 Methodology

I chose to implement a trading agent using a trading strategy inspired by Parrondo's paradox [Parrondo *et al*, 2000]. It is a paradox that shows that several losing games, if played one after another, could result in a winning game. There are many ways to construct such games. The paradox stems from the fact that although an individual game is losing in the long run, it may still gain ground during shorter intervals. By switching back and forth in between such games, hopefully switching to a game that will gain ground in the immediate future, one could actually create an overall winning game.

I wanted to see whether the theoretical results that Parrondo *et al.* had obtained could be adapted to a real stock market or not. Due to the nature of the problem, I decided to use an experimental approach. To get a quantitative measure of the results I obtained I used a benchmark strategy called the 'Buy and Hold' strategy. It is the same measure Boman *et al.* used in a similar study a year ago [Boman *et al*, 2001a].

Before I continue my exposition, I feel the need to point out that the algorithm I have implemented taken together with the stock market differs too much from the game of Parrondo *et al.* for my work to be considered a test of the practical use of their game. Rather, it should be looked upon as a separate game, although theirs has inspired it.

As the model I have designed has its roots in probabilistic reasoning, it has been natural to test the model experimentally using a statistical approach. The complexity of the stock market also makes it difficult to use any other type of approach, i.e., if one is looking for quantitative facts. This shows the non-analytical strength of artificial agents, finding new domains of knowledge not yet explored by the analytical approach.

If the model corresponds well with its intent, the trading agent I implemented should be able to earn money during an overall falling trend, or at least do better than the 'Buy and Hold' strategy.

In the concluding chapter I have also tried to throw some light on the issue of whether the strategy of the trading agent I implemented would be suitable for agents running on the servers of a stock market or not, and what effects it might have on the stock market as a whole.

As time and other constraints have set limits for my project, it can only be regarded as a tentative beginning of a deeper understanding of market-based agents.

1.4 Readers' Guide

In Chapter 2, the theoretical and practical considerations regarding the trading strategy I implemented will be expounded. Following that, there is a chapter on the experimental set-up. In this chapter an overview of the algorithm I constructed will be given and issues regarding the test data will be discussed. The benchmark strategy will also be described in the same chapter. Besides a detailed description of the algorithm and its sub-algorithm, the fourth chapter contains the results of the tests. Finally, there is a concluding chapter discussing different aspects of the results, and issues related to them. It also comprises conclusions and personal reflections. An acknowledgement and an appendix end the whole thesis.

2. The Trading Strategy

2.1 Theoretical Considerations

There are good reasons for investigating agent strategies, run by artificial trading agents, that could earn money on a receding market. Agent strategies of this type could give people opportunities to earn money, or at least lose less money, during a fall of an exchange. This would give them an extra incentive to stay on the stock market, investing in shares, and thus potentially mitigating a fall, as more capital would be invested in shares. Many equity funds also have adamant regulations that force them to stay invested in shares during a fall. Trading agents could make it possible for them to cut losses. As these agents might engage frequently in transactions, they could also increase the liquidity of the market. All this harmonizes with my aim to design an agent that contributes to the stability of the stock market.

Looking closer at the behaviour of a stock market, a well-known stylized fact is the 'leverage' effect, which corresponds to a negative correlation between past returns and future volatility [Black, 1976]. Black observed that the volatility of stocks tends to increase when prices drop. An effect particularly important for option markets. For individual stocks this correlation is moderate and decays exponentially over 50 days, while for stock indices it is much stronger, but decays faster (around 10 days) [Bouchaud *et al*, 2000]. Intuitively it seems plausible that the greater the volatility is, the greater the chances are that a strategy that benefits from the volatility of individual stocks would make money. In the contrary case, were there no volatility at all, such a strategy would not work.

There are also reasons to believe that a trading agent based on such a strategy might work in a real setting. The physicist Sergei Maslov has shown that principles similar to Parrondo's paradox could be applied to stock markets [Marsili *et al*, 1988; Maslov *et al*, 1998]. Specifically, he showed that the value of a portfolio could increase in a receding market. According to his scheme the entire portfolio has to be sold and repurchased periodically in the same proportions the different shares had originally in the preceding portfolio. The effect of this is that the gains from stocks that are doing well, at the moment, are distributed towards underachieving stocks, which hopefully will do well in the coming period. Continuing this research, Boman *et al.* have studied Parrondo strategies (strategies inspired by Parrondo's paradox) for artificial traders [Boman *et al*, 2001a]. Testing artificial traders on data from the Stockholm stock market, they found that in a receding market a Parrondo strategy could outperform 'Buy-and-Hold', value, and trend investor strategies.

Although a theoretical examination shows that the paradoxical games of Parrondo *et al.* are bound to lose if the conditions are not ideal [Parrondo *et al*, 2000], the more erratic movements of a real stock market, combined with artificial traders, may give more tempered results. On a more abstract level, there are research results that show that rationality is not a sufficient condition for survival in a financial market [Blume and Easley, 1990]. To strengthen this point, other results show that the survival of irrational traders in a market is quite possible [Schleifer *et al*, 1991]. The whole issue of rationality versus irrationality is quite controversial though, but it is an issue worth bringing up while doing research on artificial traders and their efficiency. Parrondo's paradox is easily seen as rational, but with a trading agent working in an environment with the dynamics of a real world stock market it is a different story. The fixed probabilities of the paradoxical games of Parrondo *et al.* for either gaining or losing capital would be substituted for probabilities induced by the movements of share prices. The movements of share prices, as anyone might realise, are quite difficult to predict. But introducing irrationality in this way might be the very thing that may make a trading strategy work in the long run.

2.2 Practical Considerations

There are also practical issues of some bearing. As it is, OM financial markets offer clients several order types, but all of these are non-temporal. For OM's part it would be interesting to look into the possibilities for temporal alternatives. This is one of the reasons I have chosen to take a closer look at a paradoxical game where the rules are linked to the history of the game rather than the capital [Parrondo *et al*, 2000].

One difficulty that a trading agent might face, depending on the functional details of the algorithm the agent implements, is the cost of transactions. If a strategy needs to trade frequently to be successful, high transaction costs would render such a strategy useless. Therefore I have striven to keep the number of transactions at a minimum. One way to circumvent this problem could be to introduce transaction fees based on a different set of criteria for agent-initiated transactions. Starting or ending a full run of an agent could, for example, be equivalent to an ordinary order to buy or sell a stock, with transaction fees of similar magnitude, while transactions executed during a run of an agent, transactions that the agent initiates itself, would entail much lower fees. How these criteria should be formulated is an issue that lies beyond the scope of this thesis, but I have suggested their use by keeping the cost of transaction very low in the tests I have conducted.

3 Overview of the Experimental Set-Up

3.1 The Trading Agent

For the trading agent I have implemented there is a need to make a rough estimate of the probability for an eligible stock to either rise or fall in value in the short term. For this purpose I designed a subagent with two adjustable parameters, based on a function with stock prices as variables, conceptually somewhat similar to the mathematical notion of the second derivative. If stock prices of the most recent history of the stock of interest translated to a function value above a certain threshold value set in advance, it would be taken as a sign of a coming rise of the value of the stock. Conversely, if it translated to a function value below a set threshold, it would be taken as a sign of a coming fall of the stock value. As it turned out, this type of subagent has interesting properties, making it possible to make estimates of use to the trading agent.

Four subagents of the type I have described above together with a random function constitute the trading agent I have designed. By implementing a random function, two trading agents with the same settings could get different results on the same sample of the stock market. This design has two advantages. Firstly, if two customers would choose the same set of parameters for their agents, the agents still would not follow suit, thus avoiding unnecessary volatility on the market. Second, it makes it possible to run several agents at the same time, with the same settings, resulting in a mean performance less volatile than the performance of a single agent.

3.2 Empirical Tests

I have tested the performance of the trading agent I have implemented on historical financial data from the Stockholm stock market. The samples I have used have been selected manually. One reason for this is that it is impossible to define and set limits to a relevant theoretical population that could also be used for practical purposes. One of many reasons for this, in turn, is that the future development and behaviour of the stock market is not available for our scrutiny and sampling. If I had decided on a historical population, it would still be too time-consuming to list all the samples of the past. Apart from the problems of listing, samples dating back a few years or more might be of questionable value. The present behaviour of the stock market might be quite different from that of the past. For practical reasons, I had to get the financial data for my tests from the Trust-database (historical quotes from the Nordic stock markets adjusted for splits, new issue of shares, etc.). For the given reasons, the samples I have used only encompass the Stockholm stock market and span the period between 26th February 1999 and 30th October 2001. I only used the closing rates. The stocks I used in my tests were all picked from the A-list and the Attract40-list of the Stockholm stock market, as these stocks are the most frequently traded, and therefore most representative of the stock market as a whole. Even so, there were stocks that did not register a single trade for whole days during the period in question. These stocks were left out.

With all the limitations of the experimental design indicated above, the results of the tests can only serve as a feeble indicator of the plausibility of the hypothesis; it could be different from an intuitive and practical point of view.

3.3 Financial Data

The data used for the tests are listed in the appendix. Ideally, I would have liked to have samples consisting of 10 different stocks spanning a year for all the tests, but for reasons already stated in section 3.2, I had to divide the material into periods of 63 days of trading with each sample consisting of 8 different stocks. For the tests on samples taken from the newly introduced evening exchange (closing 20:00 instead of 17:30), spanning the period between the 26th of February and the 15th of October

2001, I had to over-sample with each period overlapping the previous and following samples by 21 days of trading.

3.4 The Benchmark Strategy

To get a comparative measure of the performance of a trading agent or a subagent on a specific sample, I have compared it to the performance of a benchmark strategy on the same sample. This benchmark strategy, which is called the 'Buy and Hold' strategy, divides the same amount of money as the trading agent is allotted at the beginning of a run into as many equally sized shares as there are stocks in the sample. Using these equal amounts of money, stocks of the sample are each bought in the same proportions. Then these stocks are held on to until the end of the period of the sample in question. The total value of these stocks is then taken as a benchmark of performance for that particular sample.

Furthermore, by dividing the final sum, achieved by a trading agent on a particular sample, with that of the 'Buy and Hold' strategy on the same sample, I could compare the agent's performance on different samples. It would have been difficult otherwise, as every sample contains different stocks, and has therefore also differing overall rises or falls of stock value.

4 Algorithms, Results, and Associated Conclusions

4.1 The Subagent Algorithm

The subagent I have sketched in section 3.1 uses the three most recent stock quotes of each stock of the chosen sample of the stock market for its calculations. If t is a time-variable and if f(t) denotes the most recent quote of a specific equity line, f(t-1) the one before that, and f(t-2) the first quote in this series of consecutive quotes, the function I have used is the following:

$$F(t) = (f(t) + f(t-2) - 2*f(t-1)) / (f(t) + f(t-2)) \qquad (1)$$

The time elapsed between two quotes is 1, of whatever unit it may be, in my test samples 1 day. The numerator of the function can be seen as an incremental form of the second derivative of the graph of a stock's price over time. I chose this construct as I thought it plausible that stockbrokers and traders would react to changes of the speed with which share prices went up or down. I divided this rough form of a second derivative by (f(t) + f(t-2)) to compensate for the varying absolute values of the different stocks. In this way a stock with an absolute value of 10 crowns (SEK), for example, and with the same value for its numerator as a stock with an absolute value of 100 crowns (SEK) would result in a greater absolute value using function (1) above. It seemed wise, as this would make it possible to direct the subagent's efforts towards the cheaper shares, while each share would give possibly the same absolute return rate as the more expensive shares. In essence the denominator is the mean value of f(t) and f(t-2). There are also other alternatives to get the mean price of a stock, but as the distribution of function values over a sample spanning one year was Gaussian, as will be shown in Chapter 4, I settled with the form the function had.

At the beginning of each run the subagent was given a set amount of money. Then, at the first checkpoint in time (the third day in my tests as I used time increments of one day), it checked all the stocks in a given sample (in my samples 8 or 10 stocks) for the stock with the highest value of function (1). If that value exceeded a threshold value, set before the start of the run, it would buy that particular share using all the money allotted. Otherwise the algorithm continued to the next checkpoint in time (the following day), with its new set of stock quotes pertaining to that particular time. When a share had been thus purchased, the agent held on to that share until the share's function value fell below another threshold value, also set in advance. After having sold the share, the procedure was repeated until the end of the period of the sample in question.

4.2 Results and Conclusions Pertaining to the Subagent

To begin with, I tested the subagent for a large set of parameter values of the threshold of purchase and the threshold of sale of function (1). To get an idea of the range of values of this function, I collected all the function values of Sample 2 (see appendix) during one year. In Chart 1 the number of occurrences of values within consecutive intervals is plotted against the mid-value of the interval in question.

**Chart 1. Frequency of values of function (1) in equally spaced intervals.
Sample 2, 28 Feb 00 - 23 Feb 01.**

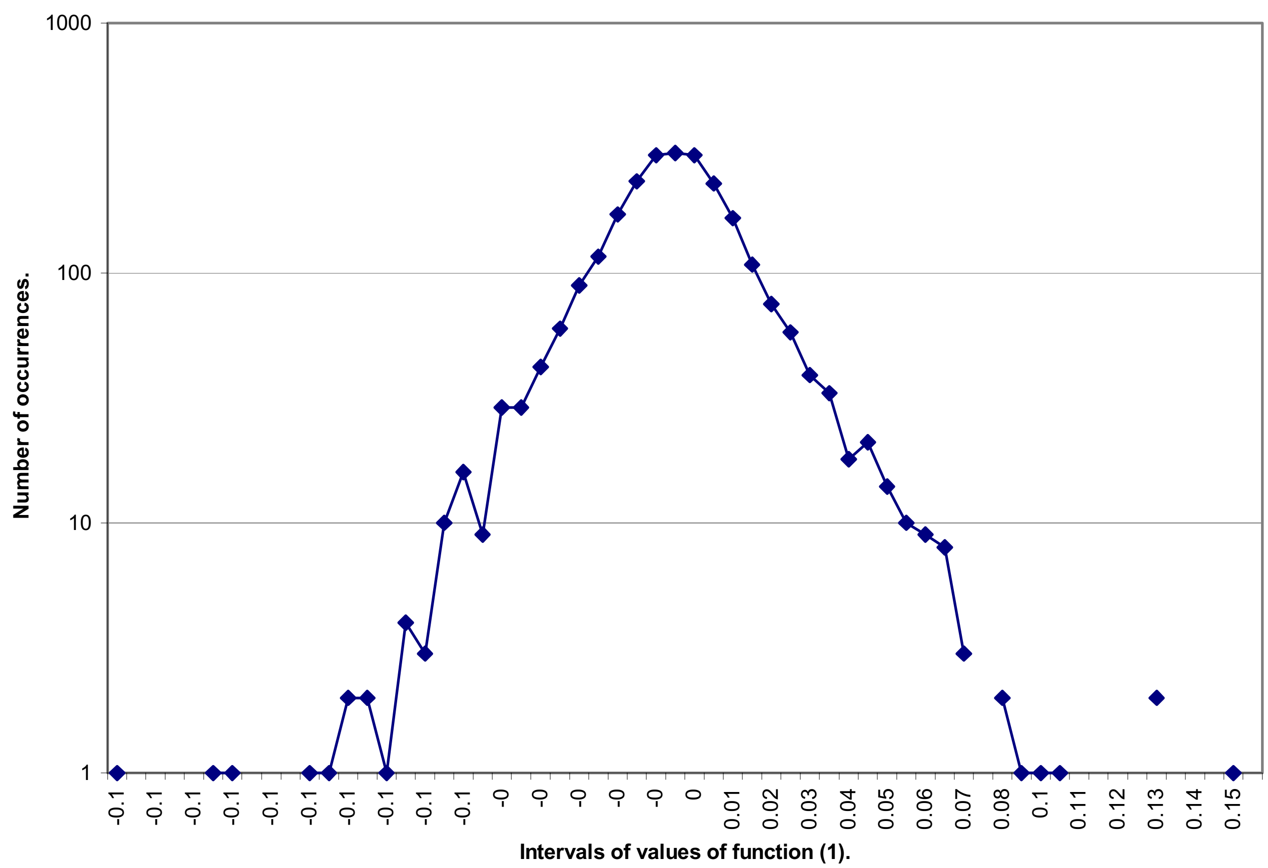

I have used a logarithmic scale; and, as can be seen, the distribution is close to a Gaussian distribution tending towards a fat-tailed distribution (a distribution with large or infinite valued standard deviation). If this would be generally the case, it would mean that we would be looking at a stable distribution of values, thus making it easier to harness the subagent. Except for a few stray values, most values lie between –0.10 and 0.10, omitting the unit, as it has no bearing.

Having obtained this result, I proceeded to test for all pairs of values of thresholds of purchase and thresholds of sale, spanning the interval between –0.15 and 0.15 (I chose to have an extra margin of 0.05) with an interval of 0.01 in between values. The outcome of these tests can be seen in Chart 2 and Chart 3 (next page). In each case the subagent was given an initial amount of 10,000 Swedish crowns (SEK). I made the rather simplifying assumption that the total cost of transaction would be 0.05 per cent of the absolute worth of the stock traded, for both purchase and sale. I opted for this low cost for a few reasons. Firstly, the Internet based broker Avanza (www.avanza.se) offers a brokerage of 0.12 per cent. That should set an upper boundary for the transaction fee. Second, as a trading agent initiates all transactions on its own, starting or stopping a full run of these agents could be looked upon as a single separate type of order, just like an ordinary limit order, for example. There would not be any need for any communication with outer sources except at the very start and at the very end of a run.

Apart from the transaction fee, there is also the implicit transaction cost of the spread (the difference between the highest bid and the lowest ask price in a double auction). It would have been impossible for me to simulate a real stock market, and thereby the cost of the spread, in a realistic way. Instead I made the idealistic assumption that I could sell or buy stocks at the rate I had used for my predictive calculations, i.e., the rates would not be affected by the agents' orders and the stock market would be extremely liquid.

In the charts (2 and 3) every crossing of two lines in the landscapes represents a pair of parameters (threshold values). The performance of the corresponding 'Buy and Hold' strategy is also given in the headings of the individual charts.

Searching for patterns in these landscapes, I immediately recognised the plateaux and valleys in the lower and right quarters of the charts (looking at the charts as they are displayed). In Chart 3, corresponding to a fall of the exchange, there is a plateau in the right quarter and a valley in the lower quarter. This corresponds to high values for both the threshold of sale and the threshold of purchase, and to high values for the threshold of sale and low values for the threshold of purchase, respectively. Here I must point out that at the very edge of the right corner, just a few or no transactions have taken place. The threshold of purchase is set at a level too high for any transaction to be initiated, except for possibly a few stray transactions. If bought, a stock would immediately be sold owing to the very high-set threshold of sale. In the chart pertaining to a rise of the exchange (Chart 2), the reverse scenario arises. In this chart there is a plateau in the lower quarter and a lowland plain in the right quarter. The plateau corresponds to a high value for the threshold of sale and a low value for the threshold of purchase, and the valley corresponds to high values for both the threshold of sale and the threshold of purchase.

The results are not too astonishing. They just state that it is always good to sell early so that you can switch to the best performing stocks, that in bad times you should only invest in stocks with really good prospects, and that in good times one should aim at always remaining invested in shares.

I decided to take a closer look at this pattern, focusing on the area close to the centre of the charts with their high peaks and deep valleys, i.e., the area where most values of function (1) were centred, to see whether it had any kind of generality or not. For this purpose I used samples 3 to 58 (appendix), in all 56 samples. 28 samples

Chart 2, 26 Feb 99 - 25 Feb 00, sample 1. Interval in between threshold values 0.01.
'Buy and Hold': 16070 SEK.

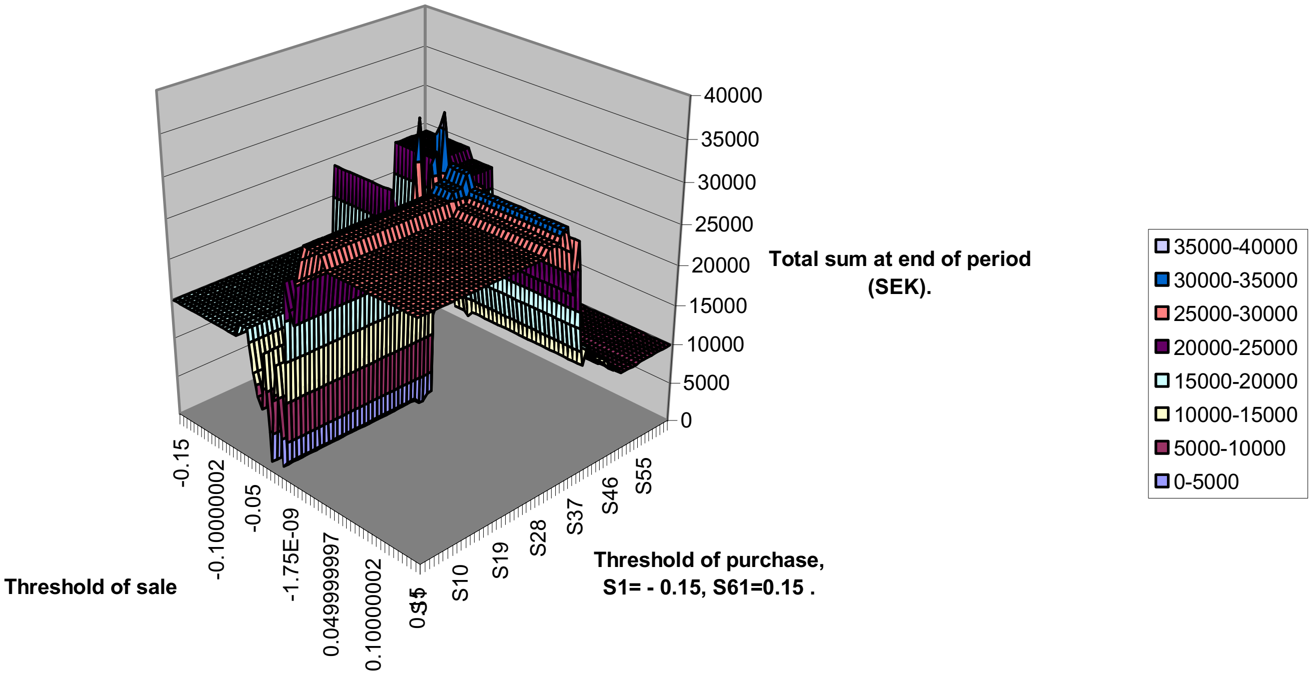

Chart 3, 28 Feb 00 - 23 Feb 01, sample 2. Interval in between threshold values 0.01.
'Buy and Hold': 6960 SEK.

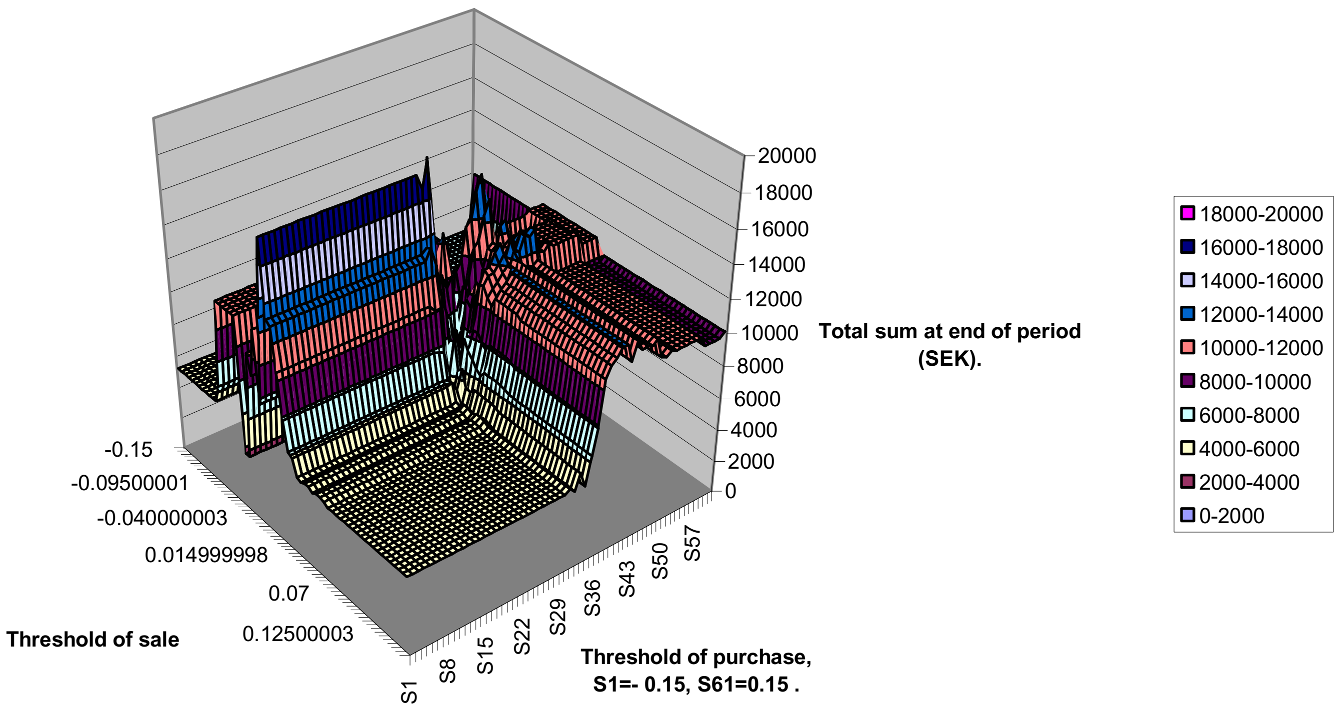

covered the period between 26th February 1999 and 24th February 2000 (an overall rise of the exchange), and the remaining 28 samples the period between 25th February 2000 and 23rd February 2001 (an overall fall of the exchange). A few of the samples of the period starting 26th February 1999 had overall falling rates, and similarly, a few samples of the period starting 25th February 2000 had rising rates. Charts 4 and 5 (next page) show the average outcome of the respective 28 samples. As mentioned earlier, each sample contains 8 stocks and covers about 3 months (63 days of trading). As can be seen, the pattern seems to persist. Chart 4 depicts a landscape of an overall rise of the exchange and Chart 5 an overall fall of the exchange. To get a more precise understanding of this pattern I have also tested for the standard deviation and other measures of accuracy of 8 pairs of parameters for both sets of samples. The resulting values can be seen in Table 1 and Table 2.

Table 1. Margins of error for the mean values of the final sums divided by corresponding final sum of the 'Buy and Hold' strategy. 26 Feb 99 – 24 Feb 00. Average of 28 samples. Transaction fee set at 0.05 per cent. Results in bold type.

	Valley				Plateau			
Sub-agent Sale purchase	1 - 0.01 0.04	2 - 0.01 0.06	3 0.01 0.04	4 0.01 0.06	5 0.01 0.01	6 0.01 0.03	7 0.03 0.01	8 0.03 0.03
Mean value	**0.99**	**0.95**	**1.00**	**0.94**	**1.02**	**1.05**	**1.01**	**1.06**
Standard deviation of mean	**0.07**	**0.03**	**0.06**	**0.03**	**0.05**	**0.08**	**0.05**	**0.08**
Corrected standard deviation	**0.33**	**0.12**	**0.30**	**0.11**	**0.24**	**0.38**	**0.24**	**0.39**

Table 2. Margins of error for the mean values of the final sums divided by corresponding final sum of the 'Buy and Hold' strategy. 25 Feb 00 – 23 Feb 01. Average of 28 samples. Transaction fee set at 0.05 per cent. Results in bold type.

	Plateau				Valley			
Subagent Sale purchase	1 - 0.01 0.04	2 - 0.01 0.06	3 0.01 0.04	4 0.01 0.06	5 0.01 0.01	6 0.01 0.03	7 0.03 0.01	8 0.03 0.03
Mean value	**1.06**	**1.07**	**1.06**	**1.08**	**1.02**	**1.00**	**0.99**	**0.98**
Standard deviation of mean	**0.06**	**0.06**	**0.06**	**0.06**	**0.07**	**0.06**	**0.07**	**0.05**
Corrected standard deviation	**0.28**	**0.29**	**0.28**	**0.29**	**0.36**	**0.27**	**0.33**	**0.26**

Chart 4, Average of 28 samples, 26 Feb 99 - 24 Feb 00. Intervals in between threshold values 0.01.

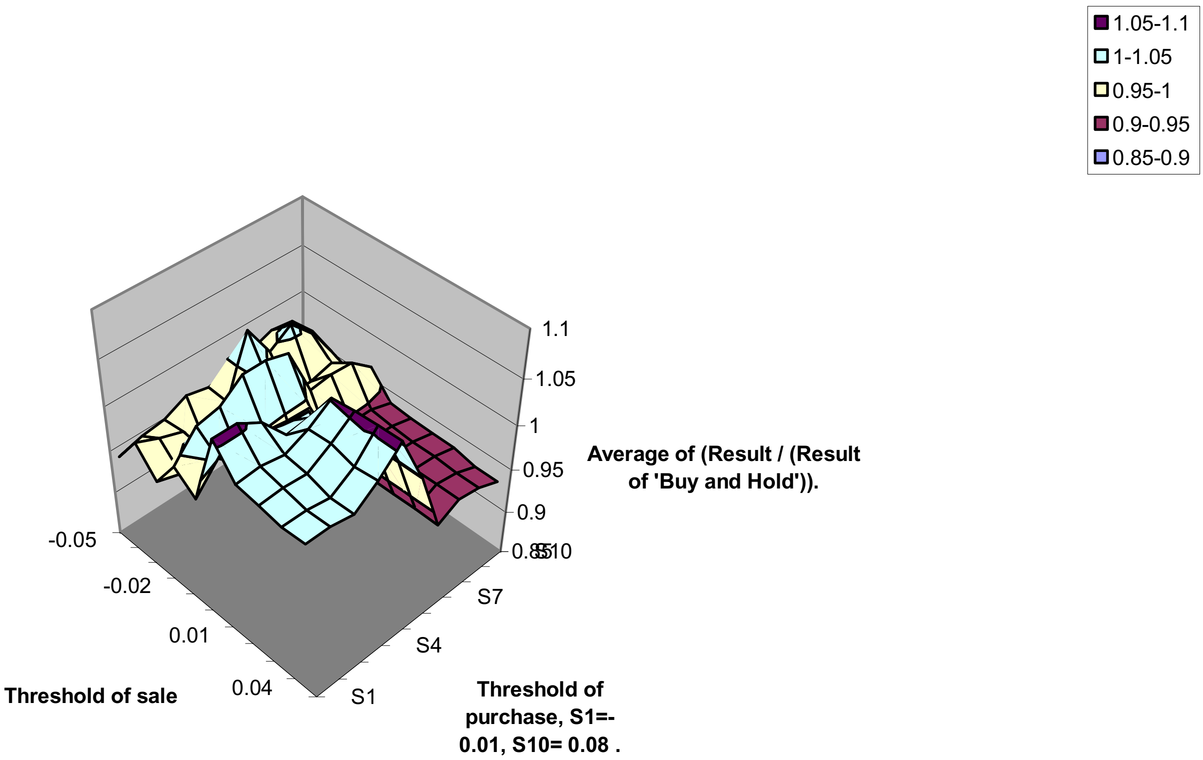

Chart 5, Average of 28 samples, 25 Feb 00 - 23 Feb 01. Intervals in between threshold values 0.01.

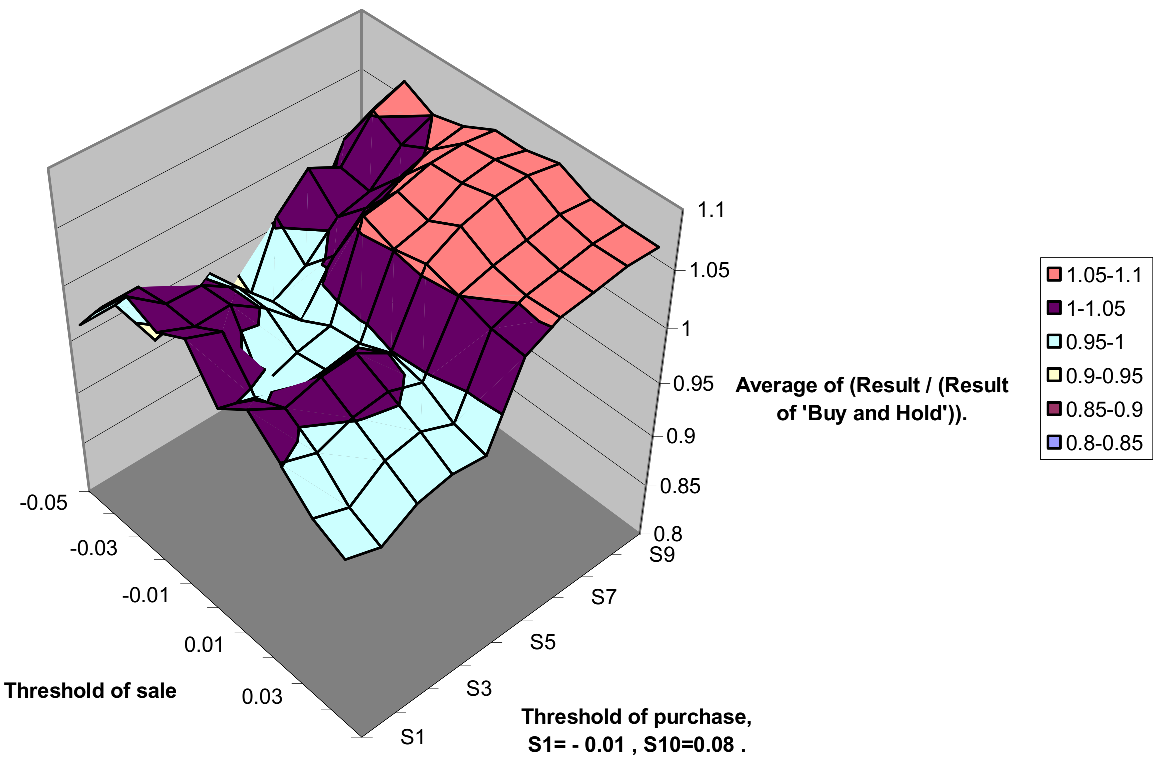

As the tables (1 and 2) show, the standard deviations of the mean values are large enough to cancel out the pattern of plateaux and valleys I was looking for, i.e., the tests cannot be considered conclusive. Extensive testing, with larger sets of samples, is needed for more definite answers. Still, with the samples I had, the mean values support the hypothesis of plateaux and valleys.

To get an idea of whether the standard deviation of the subagents was converging towards a finite value or not, I looked at its value for different numbers of samples. Table 3 shows the outcome.

Table 3. Standard deviation of subagents for different number of samples. The samples are the same as those in table 1 and 2.

	24 samples	26 samples	28 samples
Rise of the exchange Threshold of sale 0.01 Threshold of purchase 0.01	**0.252**	**0.245**	**0.238**
Fall of the exchange Threshold of sale 0.01 Threshold of purchase 0.06	**0.310**	**0.298**	**0.288**

In both cases the standard deviation can be seen to vary a little. It can also be seen to diminish with the number of samples. The evidence is too weak to be conclusive; so I have to keep the question – whether the distribution at hand is fat-tailed or not – open.

4.3 The Trading Agent Algorithm

This algorithm consists of four subagents, of the type described above, and a computer generated random function. Only one of the four subagents or the random function is run at a time. The four subagents and the random function come into play according to the following scheme. At every juncture, after having run the random function, or one of the subagents has sold its stock, or at the very start of a run of the trading agent, there is a fifty per cent chance of turning to the computer generated random function, and a fifty per cent chance of running one of the four subagents. An overview of the algorithm can be seen in Flowchart 1 (next page). The random function has equal probability of resulting in either a win or a loss. This function incurs no real loss or gain of capital to the trading agent, as would be the case of the subagents. If the overall game is set to run one of the four subagents, the subagent is chosen according to the scheme of Table 4 (next page). This means, for example, that S1 is chosen if both the previous run and the one before that have been losses, incurred by either one of the four subagents or the random function of the computer.

In the hope of being able to test for a paradoxical behaviour similar to the one of the paradoxical games of Parrondo *et al.*, I gave the subagents S1 and S4 threshold values of winning subagents, and S2 and S3 threshold values of losing subagents. Then I compared the performance of this combined trading agent with the average performance of the four subagents. I will discuss this further below.

Table 4. The choice of subagent.

The run before last	Last run	Subagent chosen for the coming run, each with its own pair of limit values
Loss	Loss	S1
Loss	Win	S2
Win	Loss	S3
Win	Win	S4

Flowchart 1. Overall view of the algorithm at work.

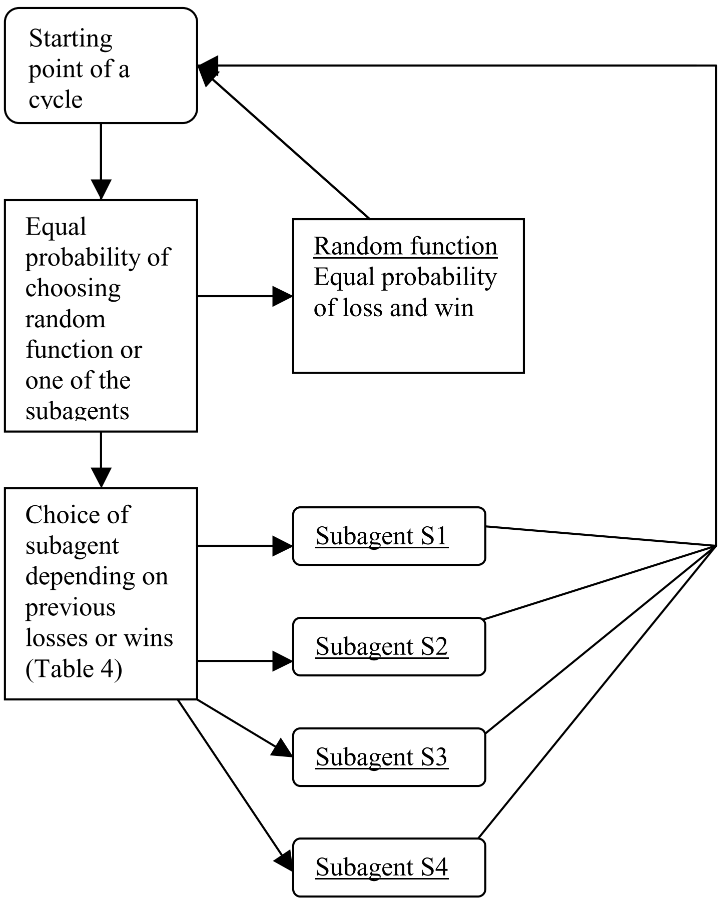

4.4 Results and Conclusions Pertaining to the Trading Agent

I tested the trading agent using the landscapes I had obtained (Charts 4 and 5). I placed subagents S1 and S4 on the plateau of the falling market and the other two in the valley. I also tested for three other combinations to see if I could get any new insights this way. Here I must emphasise that these tests did not test whether the trading agent would work in a real setting, or not. I used material that I knew would give S1 and S4 good results. The only thing I tested was whether the combined agent would do better than the average of its constituent subagents would, given that these subagents performed according to a certain scheme. If the plateaux and the valleys could be anticipated then, and not until then, the results of the tests of the trading agent could be of practical use. In Table 5 the resulting values are shown.

Table 5. Performance of four variations of the trading agent compared to the average of their four constituent subagents. The results are given in bold type. All values are relative to the 'Buy and Hold' strategy. The trading agents, corresponding to the four columns, were run 10,000 times on every sample to obtain a mean. Transaction fee was set at 0.05 per cent. The resulting improvements of the trading agents compared to their constituent subagents are highlighted with a shaded background. The values in column 2 pertain to the trading agent that adheres to the scheme of Table 4.

	Column 1	Column 2	Column 3	Column 4
Rates	Rising	Falling	Rising	Falling
Period, Samples 3 to 58	26 Feb 99 – 24 Feb 00	25 Feb 00 – 23 Feb 01	26 Feb 99 – 24 Feb 00	25 Feb 00 – 23 Feb 01
Subagent S1	In the valley	On the plateau	On the plateau	In the valley
Threshold of sale	0.01	0.01	0.02	0.02
Threshold of purchase	0.06	0.06	0.03	0.03
Subagent S2	On the plateau	In the valley	In the valley	On the plateau
Threshold of sale	0.0	0.0	0.0	0.0
Threshold of purchase	0.0	0.0	0.07	0.07
Subagent S3	On the plateau	In the valley	In the valley	On the plateau
Threshold of sale	0.02	0.02	0.01	0.01
Threshold of purchase	0.03	0.03	0.06	0.06
Subagent S4	In the valley	On the plateau	On the plateau	In the valley
Threshold of sale	0.0	0.0	0.0	0.0
Threshold of purchase	0.07	0.07	0.0	0.0
Average of agent	***0.955***	***1.082***	***0.970***	***1.073***
Average of constituent subagents	***0.998***	***1.037***	***0.998***	***1.037***
Mean value of improvement	**- 0.043**	**0.045**	**- 0.028**	**0.036**
Standard deviation of mean value of improvement	**0.021**	**0.019**	**0.017**	**0.018**

As can be seen in this table, the trading agent does on average 4.5 per cent better on a 3 month period than the average of its constituent subagents do, on a market of falling rates (column 2). The result seems quite remarkable as the performance of the combined agent reaches the same level as those of the two winning subagents do separately (compare with results of Chart 5 and Table 2). Letting the subagents S2 and S3 be the winning type of subagent, with rates falling, and with S1 and S4 going below the 'Buy and Hold' strategy, the good results persist (column 4). Looking at the period of rising rates (column 1), it seems the good performance during the period of falling rates has been traded for bad performance during the period of rising rates. The pattern is the same as the analytical results of [Parrondo *et al*, 2000].

To give an idea of how volatile the performance of a single trading agent could be, I tested one on two separate periods, one period of rising rates and one period of falling rates. The results are shown in Table 6.

Table 6. Standard deviation of a 1000 runs over 63 days of trading in 2 samples with the trading agent. Initial amount 10,000 (SEK). Transaction fees set at 0.05 per cent. Results in bold type.

Rates	Rising	Falling
	30 Aug 99 – 24 Nov 99	29 May 00 – 28 Aug 00
Subagent S1 Threshold of sale Threshold of purchase	In the valley 0.01 0.06	On the plateau 0.01 0.06
Subagent S2 Threshold of sale Threshold of purchase	On the plateau 0.0 0.0	In the valley 0.0 0.0
Subagent S3 Threshold of sale Threshold of purchase	On the plateau 0.02 0.03	In the valley 0.02 0.03
Subagent S4 Threshold of sale Threshold of purchase	In the valley 0.0 0.07	On the plateau 0.0 0.07
Mean value (SEK)	**9895**	**9685**
Mean value of 'Buy and Hold' strategy (SEK)	**11135**	**9414**
Corrected Standard Deviation (SEK)	**156**	**275**
Standard deviation of mean value (SEK)	**4.9**	**8.7**

For the sample of the falling exchange the corrected standard deviation is 275 SEK. If the distribution of the results of the 1000 runs would be Gaussian, it would mean that the confidence of a value to lie within the range 9685 - 275 SEK to 9685 + 275 SEK would be 68 per cent. This means that an agent, on a single run, might very well get results below the level of the 'Buy and Hold' strategy. It should also be noted that the standard deviation of the mean value is as low as 8.70 SEK. This margin of error is independent of the type of distribution at hand, i.e., it does not have to be Gaussian. Thus, the mean value should be quite accurate. This should not be too surprising. As the volatility of the results is due to the random function of the agent, the combined results of several agents should be less volatile. This is similar to the concept behind the use of the Sharpe ratio (a ratio that is used to find the best possible proportions of different securities in a portfolio, that also can contain cash, for optimal reward/risk ratio) [Sharpe, 1994]. Table 6 also shows that the margins of error of the period of rising rates reach the same magnitudes as those of the period of the falling rates.

Finally, as a last run of tests, I tested the trading agent on previously unseen financial data, except for the observation that it was data taken from a period of falling rates. The result is shown in Table 7. As mentioned earlier I had to resort to over-sampling in these tests. As the closing rates of the period in question were obtained at a later hour, the movements of the rates might have been a little different. Institutional investors might not have been active to the same degree, for example; liquidity might at times have dropped drastically resulting in erratic rates, etc. From the viewpoint of testing this was good, as testing on material different in character tests in a better way the generality of the model or theory tested, or in our case, the trading agent.

Table 7. Performance of the trading agent 26 Feb 01 – 26 Sept 01, a period of fall of the exchange. Calculations based on closing rates of evening exchange (20:00). Initial amount 10,000 (SEK). Results in bold type. Values for standard deviation pertaining to performance divided by 'Buy and Hold' (shaded background).

Number of runs on each sample	100				1000	
Brokerage (per cent)	0.05	0.12	0.12	0.12	0.05	0.12
Number of samples	21	21	*10*	*5*	21	21
Mean value of performance (SEK)	**9975**	**9887**	**10391**	**10299**	**9965**	**9888**
Mean value of 'Buy and Hold' strategy (SEK)	**9289**	**9276**	**9376**	**9510**	**9289**	**9276**
Mean value of performance divided by 'Buy and Hold'	**1.095**	**1.086**	**1.127**	**1.083**	**1.094**	**1.087**
Standard deviation of mean	**0.042**	**0.041**	**0.060**	**0.065**	**0.042**	**0.042**
Corrected standard deviation	**0.195**	**0.190**	**0.189**	**0.146**	**0.191**	**0.190**

In Table 7, it can be seen that even with just a hundred runs on 21 samples the standard deviation of the mean is about the same as that of a thousand runs. I also conducted tests with the cost of transaction set at 0.12 per cent, the same percentage as the brokerage of the Internet broker Avanza. With a cost of transaction set at 0.05 per cent, and still running the trading agent a 100 times, the agent achieved an average final sum 9.5 per cent higher than that of the corresponding 'Buy and Hold' strategy. Translated to a yearly basis that would be 44 per cent. Similarly, with a transaction fee of 0.12 per cent, and still a hundred runs, it would amount to 39 per cent. I have not accounted for the margins of error here. I leave that to the reader. Also, one must keep in mind that the agent does not have the money invested in shares all the time. Probably its good performance is, to some extent, due to this fact. It would be interesting to see how much time the agent keeps its capital invested relative to the length of the period it is active.

I also tested the trading agent on 5 and 10 samples to see how the number of samples would affect its results. These tests included only stocks from the OMX-list (most traded stocks on the Stockholm stock market), which makes it a bit devious to compare their results with the results of the other tests. As can be seen, they indicated that the fewer the samples were, the greater the volatility (standard deviation of mean) would be. It can also be noted that the trading agent lost some of its capital in the tests with 21 samples, but not in those of 5 and 10 samples.

5 Discussion

There are infinitely many issues to bring up in a discussion as the world is mirrored in the stock market. Of course, this is an insurmountable task. Nevertheless, there are obvious characteristics and minor details that might be worthwhile considering or that should be brought up.

5.1 Scientific Validity

There are a few points that I need to add regarding the tests and the parameters chosen for the artificial agents. When I tested the subagent for different pairs of parameters, resulting in charts 2 and 3, I could have chosen to use narrower interval spacing. Instead of 0.01 units between each value I could have used 0.001 units, as even a change of the threshold values of this magnitude could make as large a difference to the final outcome as leaps of 0.01 units. For the tests I conducted this had no significance though, as I was only looking for the general contours of the landscapes, but if a subagent would be put in use it would make a difference, as there would be more opportunities for both profits and losses.

Taking a closer look at the plateaux of both the period of rising rates and the period of falling rates one can see that they do not overlap. Actually, the edges of the plateaux are adjacent. The borderline is approximately located along a line corresponding to the value 0.035 for the threshold of purchase. This means that there were no subagents that could gain money both in the receding as well as in the rising market. This suggests that it might not be possible to make a profit on this construct alone. Comparing the two plateaux, one can also see that in the chart of the falling rates (Chart 5), the threshold of sale is slightly lower than that of the chart of the rising rates. For the chart of rising rates the edge of the plateau can be drawn along the line corresponding to a threshold value of 0 for a sale. For the fall of the exchange the corresponding value is –0.04. I cannot explain this phenomenon. Perhaps, to be able to earn money during a slump, there is a need to hold on to stocks, even though they are not doing exceedingly well.

Another point that merits some mentioning, for the sake of future research, is time and the possibility for a third parameter. Instead of being fixed, the increments of time could be adjustable, or even dynamic.

As a benchmark for the tests I ran, I could have chosen to use a ‘Buy and Hold’ strategy that sold or bought its stocks when the agent did. This would have had the advantage of giving a measure of how well the stocks the agent owned did while in its possession. As I have already mentioned above, with the benchmark strategy I used, the agents automatically lost or gained ground relative to the ‘Buy and Hold’ strategy whenever they did not own stocks. Still, I needed a Benchmark strategy of this type to be able to compare agents with different parameter settings, and results pertaining to different samples.

I must admit I do not know why the trading agent did better than its constituent subagents did in a falling trend. All it does is to look back at previous performance in a rather undisciplined manner, incorporating a virtual random game. Also, it did not perform well in a rising trend, just like the paradoxical game of Parrondo *et al.* [Parrondo *et al*, 2000]. Whether this could reduce volatility on the stock market as a whole or not, I cannot tell. At least it should give an incentive to people to stay on the stock market during a slump, investing in shares. This could subdue a fall of the exchange. One has to keep in mind though that the trading agent would only own shares part of the time it is running. Furthermore, a lot of money leaves the market due to the cost associated with the numerous transactions the agents initiate. This is another good reason for keeping the transaction fees of the agents at a lower level than ordinary transactions do.

As it turned out, the cost of transaction was of paramount importance. The agent I have implemented engage in transactions frequently, typically around a hundred a year. It would have been interesting to see how a comparatively large change of the cost of transaction would have affected the performance of the agents. Having looked at just a few samples, my estimate is that the total cost of transaction cannot exceed a figure of somewhere in between 0.15 and 0.5 per cent for the agent to remain profitable.

Another option for people who want to stay on the market, in spite of falling rates, or who want to beat the index in a rise, would be to just combine freely subagents with differently set parameters. In this way people could decide for themselves how much risk they would like to be exposed to. I would not be surprised if a set of subagents put simply together could achieve results similar to those of the trading agent I put together by design.

What I have done is a continuation of what Boman *et al.* did in their study of the Parrondo strategy for artificial traders [Boman *et al*, 2001a]. One could say that their agents depended on how rates would change (which could be looked upon as a rough form of the first derivative) and was inspired by a capital dependent version of the paradoxical games of Parrondo [Harmer and Abbott, 1999]. I have gone one step further. I have based my agents on a function conceptually similar to the second derivative, and it was inspired by a temporal version of the paradoxical games of Parrondo *et al.* [Parrondo *et al*, 2000]. Furthermore, I have incorporated the cost of transactions in my tests. I have also increased the number of samples. The results I have obtained are valid within the framework of the set-up of my experiments, as I have described them. As I have also shown, they have some academic significance. The question whether the work I have done has any practical significance or not, I have to leave to future research.

5.2 Commercial Applicability

In auctions, time is very important, especially facing a reality where split seconds might make the difference between loss and profit. If a customer is confined to sending his orders via the Internet he is at a great disadvantage compared to competitors that have direct links to the servers of a stock market. There is research that shows the importance of monitoring network resources on a competitive market [Boman, 2001]. Artificial agents could remedy injustices that unequal network resources create.

As I have already indicated in a previous section, it would be desirable with trading agents that could subdue overreactions. Investors would like stock prices to be as close to the equilibrium price as possible. A volatile market carries with it a greater risk for not trading at this price, which is important to investors. The subagent I have constructed may reduce the volatility of stocks as it probably buys and sells stocks a little earlier than the ‘herd’, thereby facilitating a turn of a trend earlier than it would otherwise.

Looking at the whole idea of having trading agents implemented on the servers of a stock market, I see a danger in having agent strategies open to public scrutiny, which they would have to be to some extent. The reason for this is that if a strategy is known it is always easy to find a counter strategy. I do not know whether this could become a real threat in our case or not. Maybe the complexity of the system, which the agents would form together with a stock market as a whole, is too great for anyone to hit upon a working counter-strategy exclusively benefiting from the strategy of our agents.

For future research, other issues that need to be addressed are how taxes would affect profits, and the judicial positions of OM and clients using these agents.

5.3 Conclusions

The test results of the trading agent turned out to be quite interesting. Very much like the theoretical game of Parrondo *et al.* [Parrondo *et al*, 2000], it did better than its constituent subagents on a market of falling rates and worse than the constituent subagents on a market of rising rates. It was very hard for me to draw any conclusions regarding the core intelligence of the trading agent. All I could do was to observe. Although uncertain, the final run of tests showed that the trading agent achieved a result 44 per cent greater than that of the ‘Buy and Hold’ strategy, translated to a yearly basis.

Keeping the assignment and the hypothesis of the first chapter in mind, the results I have obtained show that trading agents could very well be a part and functionality of the servers of a stock market and at the same time augment the functionality of a deterministic matching service.

The research I have been doing is still very much at its beginning stages. The results I have obtained have academic value, but for practical purposes a lot remains to be done.

5.4 Personal Reflections

Finally, I would like to add a few words on efficiency of a marketplace. As I have already mentioned earlier, it is probably important to use one’s intuitive faculties in trying to understand a field as complex as finance. Intuitively, I feel that the only possible way to really create a market that does not overreact to events, where prices are based on the real worth of companies, and where companies do not have to worry about the going rate of their stock, is to focus on ethics. For this we have to go back to our schools and universities. This could only be a long-term solution though. For our own generation it is probably too late. In the mean time I think we should face reality and take a look at the character traits of people buying and selling stocks. My intuitive insight is that these people (basically all of us including myself) need to experience the stock market as boring, at all times. They should feel that there are no shortcuts to wealth. In this way the volatility of stock prices would decrease. People would act in a way that would be more beneficial for our society as a whole. At the same time I do realise that it would not be too realistic to assume that people would adopt such a view. Until then, let us have some fun running artificial agents! People love sports and excitement! Besides, considering the fees stockbrokers and equity funds charge they need to be challenged.

Acknowledgement

First of all, I would like to extend my gratitude to David Lybäck, who, always sharing his insights with me, has, I feel, gone much further than duty required. Through the many discussions we have had, I have become aware of many important aspects of the subject I have been grappling with.

I am also grateful to Dr. Magnus Boman who was instrumental to this project at the very outset. He has given me important advice steering my work towards safer waters. He has this tacit, maybe subconscious, sense for where good research can be found, research that is both safe and at the same time fruitful. He gave me the courage to go into a field I was not too well acquainted with at the outset. Actually, I realised through my own work that knowledge can even be a hindrance as it confines one's actions to a world limited by that which can be preconceived.

I would also like to thank Dr. Mats Danielson for the time he has given me discussing different aspects of research and finance, and life at large. A man of solid reading, he has shared with me his clear vision of research methodology.

I must also express special thanks to my mother, Dr. Shiu-Pang Almberg, ever full of encouragement, for making time to proofread this material.

OM AB has provided me with the financial means and facilities needed for the project.

Finally, and most of all, I stand in great gratitude to life itself. I cannot pretend that I have been happy and relaxed all way through, but there has been peace and a sense of meaning accompanying my work. My wish is that the work I do – I have the action itself in mind and not its results – one day will become an utter celebration of life. Kahlil Gibran has expressed this state of being beautifully in the novel 'The Prophet'. A ploughman asks the prophet to speak to the crowd of work, and he says:

> "You work that you may keep pace with the earth and the soul of the earth.
> For to be idle is to become a stranger unto the seasons, and to step out of life's procession that marches in majesty and proud submission towards the infinite.
> When you work you are a flute through whose heart the whispering of the hours turns to music."

These words have stuck to me. 'When you work you are a flute through whose heart the whispering of the hours turns to music.' [Gibran, 1923, p.13].

References

[Black, 1976] F. Black. *Studies of stock price volatility changes.* Proceedings of the 1976 American Statistical Association, Business and Economical Statistics Section, 177.

[Bloomfield *et al*, 2000] Robert Bloomfield, Robert Libby, Mark W. Nelson. *Underreactions, overreactions and moderated confidence*. Journal of FINANCIAL MARKETS, **3**:113-137, 2000.

[Blume and Easley, 1990] L. Blume, D. Easley. *Evolution and market behavior*. Journal of Economic Theory, 1990.

[Boman *et al*, 2001a] Magnus Boman, David Lybäck, Stefan Johansson. *Parrondo strategies for artificial traders*. World Scientific, March 2001.

[Boman, 2001] Magnus Boman. *Trading Agents*. AgentLink News. **6**:15-17, 2001, http://www.agentlink.org , 2001-10-17.

[Bouchaud *et al*, 2000] Jean-Philippe Bouchaud, Andrew Matacz, Marc Potters. *The leverage effect in financial markets: retarded volatility and market panic*. http://www.arxiv.org/list/cond-mat/0101?100 , 2000-09-25.

[Chang, 2001] Kenneth Chang. *New Age Bidding: Against Computers, Humans Usually Lose*. The New York Times, January 2001.

[Egeberg *et al*, 1986] Sonja Egeberg, John Halse, Espen Jerlang, Ann J. Jonassen, Suzanne Ringsted, Birthe Wedel-Brandt. *Utvecklingspsykologiska Teorier, chapter 6: Abraham Maslow och den humanistiska psykologin.* Munksgaard, 1986, In Swedish.

[Graham-Rowe, 2001] Duncan Graham-Rowe. *Robots beat human commodity traders*. http://www.newscientist.com/news/news.jsp?id=ns99991131 , 2001-08-30.

[Harmer and Abbott, 1999] Gregory P. Harmer, Derek Abbott. *Losing strategies can win by Parrondo's paradox*. NATURE, Volume 402, December 1999.

[Haw, 2001] Mark Haw. *phenomena: Stockbrokers may act like sheep*. http://www.nature.com/nsu/010104/010104-1.html , 2001-03-14.

[LeBaron, 1998] Blake LeBaron. *Agent Based Computational Finance: Suggested Readings and Early Research*. Journal of Economic Dynamics and Control. Volume 24, 2000.

[Lux, 1995] Thomas Lux. Herd behaviour, Bubbles and Crashes. Economic Journal. **105**:881 – 896, 1995.

[Marsili *et al*, 1988] Matteo Marsili, Sergei Maslov, Yi-Cheng Zhang. *Dynamical optimization theory of a diversified portfolio.* Physica A **253**:403-418. 1988.

[Maslov *et al*, 1998] Sergei Maslov, Yi-Cheng Zhang. *Optimal investment strategy for risky assets.* International Journal of Theoretical and Applied Finance. **1**(3):377-387, 1998.

[Mirowski, 1989] Philip Mirowski. *More Heat than Light: Economics as social physics. Physics as nature's economics.* Cambridge University Press, New York,1989.

[Osborne, 1977] M.F.M. Osborne. *The Stock market and Finance from a Physicist's viewpoint.* Crossgar Press, Minneapolis, West Lake of the Isles Parkway, 1977.

[Parrondo *et al*, 2000] Juan M.R. Parrondo, Gregory P. Harmer, Derek Abbott. *New Paradoxical Games Based on Brownian Ratchets*. PHYSICAL REVIEW LETTERS, December 2000.

[Schleifer *et al*, 1991] J.B. DeLong, A. Schleifer, L.H. Summers. *The survival of noise traders in financial markets*. Journal of Business, 1991.

[Sharpe, 1994] William F. Sharpe. *The Sharpe Ratio*. Reprinted from: The Journal of Portfolio Management, Fall 1994. http://www.stanford.edu/~wfsharpe/art/sr/sr.htm , 2002-02-25.

[Wennerberg, 2001] Christer Wennerberg, Vice President Business Development, OM fixed income exchange. Interview at OM AB's head office. Stockholm. September 2001.

All links verified as of the 1st of March 2002.

Appendix

Table 8. Samples 1 and 2. The names of the stocks abbreviated according to standard.

	99-02-26 until 00-02-25	00-02-28 until 01-02-23
ELUX-B, ALIV-SDB, EURO, HOLM-B, NOKI, SAND, WM-B, SKF-B, TEL2-B, SDIA	Sample 1	Sample 2

Table 9. Samples 3 to 30. The names of the stocks abbreviated according to standard.

	99-02-26 until 99-05-31	99-06-01 until 99-08-27	99-08-30 until 99-11-24	99-11-25 until 00-02-24
ALIV-SDB, HM-B, ASSA-B, ATCO-A, ELUX-B, ERIC-B, EURO, FSPA-A	Sample 3	Sample 4	Sample 5	Sample 6
NDA, NOKI-SDB, SAND, SCA-B, SDIA, SEB-A, SECU-	Sample 7	Sample 8	Sample 9	Sample 10

B, SHB-A				
TEL2-B, VOLV-B,WM-B, STE-R, ASDO, GAMB-B, INDU-A, KINV-B,	Sample 11	Sample 12	Sample 13	Sample 14
SWMA, AXFO, BURE, GIND-B, GUNN, HLDX, HOGA-B, HUFV-A	Sample 15	Sample 16	Sample 17	Sample 18
SAPA, SAS, SLT-B, TICK, TREL-B, WIHL-B, ACTI-B, BCOR	Sample 19	Sample 20	Sample 21	Sample 22
ENEA, IBS-B, IFS-B, INT-B, MAXM, MTG-B, RATO-B, RROS	Sample 23	Sample 24	Sample 25	Sample 26
HOLM-B, INVE-B, SKA-B, SKF-B, SCV-B, SSAB-A, LDEX, NCC-B	Sample 27	Sample 28	Sample 29	Sample 30

Table 10. Samples 31 to 58. The names of the stocks abbreviated according to standard.

	00-02-25 until 00-05-26	00-05-29 until 00-08-28	00-08-29 until 00-11-23	00-11-24 until 01-02-23
ALIV-SDB, HM-B, ASSA-B, ATCO-A ELUX-B, ERIC-B, EURO, FSPA-A	Sample 31	Sample 32	Sample 33	Sample 34
NDA, NOKI-SDB, SAND, SCA-B, SDIA, SEB-A, SECU-B, SHB-A	Sample 35	Sample 36	Sample 37	Sample 38
TEL2-B, VOLV-B,WM-B, STE-R, ASDO, GAMB-B, INDU-A,	Sample 39	Sample 40	Sample 41	Sample 42

KINV-B				
SWMA, AXFO, BURE, GIND-B, GUNN, HLDX, HOGA-B, HUFV-A	Sample 43	Sample 44	Sample 45	Sample 46
SAPA, SAS, SLT-B, TICK, TREL-B, WIHL-B, ACTI-B, BCOR	Sample 47	Sample 48	Sample 49	Sample 50
ENEA, IBS-B, IFS-B, INT-B, MAXM, MTG-B, RATO-B, RROS	Sample 51	Sample 52	Sample 53	Sample 54
HOLM-B, INVE-B, SKA-B, SKF-B, SCV-B, SSAB-A, LDEX, NCC-B	Sample 55	Sample 56	Sample 57	Sample 58