

Comparative study of central decision makers versus groups of evolved agents trading in equity markets

Cyril Schoreels, Jonathan M. Garibaldi
Automated Scheduling, Optimisation and Planning Research Group
School of Computer Science and IT, University of Nottingham
Jubilee Campus, Nottingham, United Kingdom
czs@cs.nott.ac.uk

Abstract

This paper investigates the process of deriving a single decision solely based on the decisions made by a population of experts. Four different amalgamation processes are studied and compared among one another, collectively referred to as central decision makers. The expert, also referred to as reference, population is trained using a simple genetic algorithm using crossover, elitism and immigration using historical equity market data to make trading decisions. Performance of the trained agent population's elite, as determined by results from testing in an out-of-sample data set, is also compared to that of the centralized decision makers to determine which displays the better performance. Performance was measured as the area under their total assets graph over the out-of-sample testing period to avoid biasing results to the cut off date using the more traditional measure of profit. Results showed that none of the implemented methods of deriving a centralized decision in this investigation outperformed the evolved and optimized agent population. Further, no difference in performance was found between the four central decision makers.

1. Introduction

This paper investigates the use of centralized decision makers compared to groups of individuals in a portfolio optimisation simulation. The averaged performance of a group of trading agents has previously been shown to perform comparatively to the buy and hold strategy and professionally managed funds [10]. In this paper an attempt is made at deriving a single trading strategy in the form of a centralized decision maker based on the decisions made by a group of trained trading agents. The performance of the centralized decision maker is then compared to the averaged performance of the reference

population. Though the merit of groups of independently trading agents in a variety of market settings has been extensively studied [2, 3, 6, 12, 13], the potential benefits offered through group synergies need to be explored further.

Reh fuss, Wu and Moody compared four alternative approaches to combining forecasts into one overall prediction or decision based on a sample trading system [8]. To increase robustness and adaptivity of their system they introduced retraining of their models at each step as well as a form of exponential decay to include a time dependent factor. The first system for example was based on a combination of all individual predictors, which would generate a trading signal of the corresponding sign when the forecasted number deviated from the mean of its standard deviation. Also, a form of hybrid of their approaches was used which they argued seemed to minimize maximum losses and proved more robust than others, with their pure voting and other systems performing worse. The authors concluded that creating committees would be a useful way to decrease noise and that in decision making, it might be useful to combine decisions rather than their forecasts.

In the context of a decision support system, Vahidov and Fazlollahi developed a framework of hierarchical agent groups with a decision maker at its top, aiming for an agent-based decision support system that parallels the human problem solving process [15]. The hierarchy consists of three groups or layers of agents representing phases in their problem solving model, intelligence group, design group, choice group, where each layer must incorporate a range of aspects and different views, for instance two schools of thought of fundamental analysis and technical analysis. They found that using such a method it outperformed traditional decision support systems.

An alternative to centralizing decisions is shown by Ontañón and Plaza [7]. The authors focused on committees of agents with learning capabilities where no

agent is omniscient but has a local, limited, individual view of data. In their system, agents can solve any problem individually, however then try to create more accurate classifications through collaboration. Using a collaboration scheme based on symbolic justification of results, agent's results are aggregated using a weighted voting scheme using confidence measures as weights. Different to previous research, the system does not contain a centralized method that has control over the entire training set. It could be argued that each agent is self-contained and only through their interaction does collectivism emerge, rather than enforcing unity through a centralized control mechanism that has access to the internals of every agent, transforming them into transparent components of the overall system. They state that their method works well in different scenarios where normal voting tends to fail sometimes.

Schulenburg and Ross present an analysis of groups of agents trading in a stock market environment using historical data and a learning classifier system [9]. One of the key elements in their study was to focus on aspects of group decision-making rather than that of individual traders, based on the greater complexity of modeling a human individual with unique characteristics rather than as an abstract member of a group. Interestingly, one of their conclusions was that technical trading is a valid method to employ for trading in financial markets, but as with most things, can be less or more appropriate at times compared to non-technical trading. Technical trading being defined as using technical analysis such as indicators to determine buy/sell/hold situations, while non-technical trading encompasses a more diverse source of information not necessarily based on any calculations or numbers. Furthermore, they suggest a form of 'super trader', who could base decisions on both technical and non-technical traders to overall hopefully perform better than either. This resembles group decision making, whereby a collective decision is taken from various 'different' types of agents that is then used to make overall trading decisions as one entity. In this paper we take the first step in this direction for our future work by attempting to extract such overall trading decisions, or centralized decisions as referred to here, from a group of technical traders.

2. System description

Parts 2.1 to 2.5 of the system used in this investigation have been previously described in [11] and other previous work. However for completeness and understanding a detailed description is included.

2.1. Trading agent design

Every agent represents an individual trader with a personal portfolio and capital holdings, collectively referred to as their total assets or worth. With a fixed starting capital, at the end of every trading day each trader uses historical price data to make a decision on every security whether to buy, sell or hold. No limitations, apart from capital constraints, are set on the number of trades conducted.

Every agent uses technical indicators to generate trading signals for every security that is assessed. Indicators used are the Simple Moving Average (MA), Relative Strength Index (RSI), Price-Rate-of-Change (ROC), Stochastic Oscillator (SO), Moving Average Convergence Divergence (MACD) and Bollinger Bands (BB). Depending on the signals returned, an overall decision of buy, sell or hold is made for a security depending on the agent's decision type. The amount of capital invested is again determined by the genome, with it equally allocated between each security flagged for acquisition. Securities that are sold are converted into capital at their current closing price.

Every agent contains a string of integers, representing the agent's genome. Parameters used in the agent's analysis and decision process are determined by the agent's genome, effectively determining its trading behavior. Mapping of genome to phenotype are discussed throughout the following text. The genome for every agent consists of integers taking values 1-10, of length 28. The exceptions to this were genes 1, 2, 21 and 28, defining decision types with a cardinality of 4, the risk averseness factor with a cardinality of 2 as well as the SO D variable and BB deviation variable with a cardinality of 5. All genes are defined as in Table 1.

Table 1. Gene descriptions

Gene	Range	Description/Function
G ₁	1-4	Decision type
G ₂	1-2	Risk averseness factor
G ₃	1-10	Capital investment proportion
G ₄	1-10	Moving Average weight
G ₅	1-10	RSI weight
G ₆	1-10	Short-term ROC weight
G ₇	1-10	Long-term Price ROC weight
G ₈	1-10	SO interpretation 1 weight
G ₉	1-10	SO interpretation 2 weight
G ₁₀	1-10	MACD weight
G ₁₁	1-10	BB weight
G ₁₂	1-10	MA short-term value
G ₁₃	1-10	MA long-term value
G ₁₄	1-10	RSI time period
G ₁₅	1-10	RSI buy threshold

G ₁₆	1-10	RSI sell threshold
G ₁₇	1-10	ROC level
G ₁₈	1-10	ROC short-term value
G ₁₉	1-10	ROC long-term value
G ₂₀	1-10	SO K variable value
G ₂₁	1-5	SO D variable value
G ₂₂	1-10	SO buy threshold
G ₂₃	1-10	SO sell threshold
G ₂₄	1-10	MACD short-term value
G ₂₅	1-10	MACD long-term value
G ₂₆	1-10	MACD signal line
G ₂₇	1-10	BB time period value
G ₂₈	1-5	BB deviations number

Neither transaction costs nor interest on held capital were included, and the environment is assumed discrete and deterministic in a liquid market, meaning that an agent's actions cannot affect prices.

2.2. Technical indicators and the decision process

There are essentially three steps that every agent follows to determine whether or not to add a particular security to its acquisition or sale list. To analyse a security it first performs the indicator calculations, each using historical closing price information. Depending on the indicator, one or more values are returned which then need to be interpreted for the calculated values to gain meaning. Based on this interpretation, every agent then compiles a list of buy and sell signals for every security as the second step in this process. The final step is generating an overall buy, sell or hold decision for each security, based on the agent's decision type.

The following is a brief description of the indicators employed, and how an agent's genome is used to individualize the calculations. Technical indicators are tools used in the technical analysis of financial markets, exploiting the existence of trends to determine potential buy, sell or hold conditions. Indicators are mathematical formulae, commonly based on closing price or volume data, with price information being used exclusively in this system. Though markets are often argued to move randomly [5], regularities do appear and lead to observed phenomena such as seasonal cycles for example [14], which are exploited by the indicators.

When an agent is initialised, some of its gene values would not be appropriate for direct use in the system and need to be modified. For instance, though G₁₇ can be used directly in the technical indicator without translation, this is not possible for G₁₃. As an informal rule, the translation was based on achieving a representative average value approximate to the range used in wider literature.

A MA shows the average value of a securities price over time. The short-term and long-term moving average values are calculated over 4G₁₂ days for the short-term and (5G₁₃)+50 days for the long term. If the MA over the short-term is larger than over the long-term, it indicates an upward trend and a buy signal would be generated. If the MA over the short-term is smaller than the long-term, a downward trend is indicated and a sell signal would be generated. Additionally, the short-term MA can be compared to the current price of the security, which if greater, would indicate a downward trend and hence a sell signal should be generated. In this implementation, if the agent is risk averse, as determined by G₂, it bases its interpretation on a logical AND between those two interpretations and is therefore more reluctant to generate a buy signal. On the other hand, if the agent is risk taking, a logical OR is used and either interpretation suggesting a buy would suffice for the agent to consider this a buy signal. 1 buy or sell signal is generated by this indicator.

The RSI is a price-following oscillator that measures the magnitude of gains and losses of a single security over a specific time period to determine the current trend. It is calculated over 2.5G₁₄ days. If the calculated value lies above or below the sell (4G₁₆)+50 or buy 5G₁₅ threshold respectively, the appropriate signal will be generated. 1 buy, sell or hold signal is generated by this indicator.

The ROC indicator is based on the assumption of cyclical price movements, and considers the relative change of prices over time to indicate trends. The period considered is 2G₁₈ days for the short-term and 4G₁₉ days for the long-term. The ROC is repeated for both long- and short-term analysis, where if the calculated value lies below the negative threshold value of -G₁₇ it indicates a buy, while a calculated value above the positive threshold value of +G₁₇ indicates a sell in both instances. 1 buy or sell signal is generated for the short-term, and 1 buy or sell signal is generated for the long-term.

The SO compares a security's price relative to its price range over a given time period, using two parameters commonly defined as K and D. K is calculated over 1.5G₂₀ days, while D is a moving average of K over G₂₁ days. Multiple interpretations are possible, though the following two are used in this instance. First, it can be considered a buy signal if the K value is larger than the D value or a sell signal if D is larger than K. Second, threshold values can be added for both K and D. In that case, if K and/or D is smaller than the buy threshold of 3.5G₂₂ a buy signal is generated, or equally, if K and/or D is larger than the sell threshold of 4G₂₃ a sell signal is generated. 1 buy or sell signal is generated by the first interpretation, and 1 buy or sell signal is generated by the second.

The MACD "is a trend following momentum indicator that shows the relationship between two moving averages

of prices” [1]. It is calculated for the short-term over $2G_{24}$ days, and over $4.5G_{25}$ days for the long-term. The MACD compares its calculated value to a moving average of itself over a time period of $1.5G_{26}$ days, whereby a buy signal is generated if the MA is smaller, and a sell signal if the MA is larger. 1 buy or sell signal is generated by this indicator.

Lastly, BBs are generally used to provide a form of guideline, indicating possible trend reversals. The upper and lower bands are calculated over $3.5G_{27}$ days. The bands indicate that when the current price breaks through the lower band of $3.5G_{27}$ it is considered a buy signal, while if it breaks through the upper band of $3.5G_{27}$ it is considered a sell signal. 1 buy or sell signal is generated by this indicator.

In order to allow for different approaches that exist among real traders to selecting securities for purchase or sale, four agent decision types were implemented.

i. Decision type 1 (DT1) performs a simple comparison between the number of buy and sell signals, taking the appropriate action if one is greater than the other for any particular security. For instance, out of the 8 possible signals used here for a security, if 3 are buy and 2 are sell, an agent of this type would decide to purchase this security.

ii. Decision type 2 (DT2) follows the same principle as DT1. However it also stipulates that for a buy or sell action to occur, at least half of all signals must be in favour. In this instance for example, out of 8 possible signals, if only 3 are buy signals even though no sell signals exist, no action will be taken as it failed to reach the minimum buy threshold.

iii. Decision type 3 (DT3) sums the indicators by taking buy signals as +1 and sell signals as -1, as well as including a weighting process on each signal (G_4 to G_{11}), increasing or decreasing its impact on the final sum. Therefore a positive sum would translate into an overall buy signal, while a negative sum into an overall sell decision.

iv. Decision type 4 (DT4) follows the same principle as DT3, except for adjusting the final sum to create a threshold value which it needs to exceed prior to resulting in an overall buy or sell decision. For example, though decision type 3 would generate a buy signal for a value of +5, the threshold for decision type 4 is set at +/-10 and would therefore result in a neutral hold decision.

2.4. Genetic algorithm

The genetic algorithm forms the core functionality of the system in determining the evolution of the agents. The choice to use a genetic algorithm was based on previous research [10, 11], as well as the algorithms ability to deal with irregularities and exceptions in data. As markets are

affected by human expectations and at times irrational human behavior, a robust algorithm is essential.

A general representation of how the agent population is evolved is shown in figure 1. The elite population and killed population are each respectively the top and bottom 25% of the entire agent population. In an unpublished study, the top and bottom population sizes were varied and tested, with 25 forming a convenient and equally effective measure for both using an otherwise identical setup as described here. The alternative values for the top and bottom populations investigated were combinations of 5, 10 and 25.




Figure 1. Evolution of agent populations

For selection, elitism [6] is used, whereby a portion of the most successful agents carries forward unaltered every generation. Immigration [4] was also employed, where a portion of the worst performers are killed off and replaced by a new randomly generated group of agents immigrating into the system, constantly introducing new genetic material. This facilitates greater coverage of the search space, while also avoiding premature convergence and non-exclusion of other possible solutions not present in the original base population’s gene pool. Each agent in the mediocre population, those not killed or part of the elite, randomly selects another agent from the mediocre and elite population and uses two-point crossover to create an offspring. This offspring then replaces the parent from the mediocre population. Two randomly selected agents of the same decision type can mate, with a random part of the first agent’s genome being replaced by the equivalent section from the second agent’s genome, forming a new genome combination. In this process there exists a 25% chance of an agent mating with an agent of a different decision type, as for instance weighting genes will not have been relevant to types 1 or 2 previously.

In other research, performance tends to generally be measured as an agent’s capital and value of all holdings at the end of the trading period relative to its starting value. However, as this biases results based on the cut-off date used, a more representative picture of performance throughout the entire testing period can be obtained if the area under an agent’s total asset graph is considered as its fitness. We therefore propose that performance refers to the area under an agent’s total asset graph for the trading period being assessed.

2.5. Historical market data

The system uses historical financial data taken from the DAX-30, which is an index listing the top 30 capital weighted companies registered in the German market, with various weighting factors applied to each listed company to determine their impact on the Index. Data used in the system covers the time range from 01.01.1990 to 31.12.2002, however it does not include all securities from the DAX for its entire span. Due to changes in the constituents of the Index, daily closing price information was only available for 20 securities over the desired time span. Those included are listed in Table 2 using their Wertpapierkennnummer (German security identification number) below.

Table 2. Securities used in system

840400	648300	519000	717200	515100
760080	823212	803200	723610	575200
802200	703712	761440	766400	695200
593700	843002	750000	543900	514000

To evolve the population of 100 agents, 1465 trading days worth of data was used representing six year from 01.01.1990 to 29.12.1995. As shown previously [11], this time period was assumed to be long enough to evolve a competitive population. The out-of-sample period was taken over six year from 01.01.1997 to 31.12.2002. The fitness function used the total area under their total asset graph over the entire training duration.

2.6. Central decision maker agent

Extending previous work and prior to being able to design a centralized decision system, the issue of deriving a single sequence of decisions from a population of agents needs to be addressed. Currently, each agent in the population makes its own decisions and holds its own portfolio. However, given various possible approaches or methods of amalgamating the population's individual decisions into one, it is important to determine which of these amalgamation methods will actually result in performance similar to that demonstrated by the population overall. In other words, using one amalgamation approach may result in a sequence of decisions that do not perform very well, while another might equally well significantly outperform other strategies. The question addressed here will be whether an amalgamation of decisions from individual agents is feasible and results in comparable performance to that of the population average, as well as, investigating alternative amalgamation methods and selecting the most successful, if any, for further research. This is based on

the hypothesis that given a population of agents each holding their own portfolio and making their own decisions, it is possible to derive an amalgamation of these decisions that when applied to the market would result in performance comparable to the population's average.

Four amalgamation methods were implemented, each based on the collection of buy, sell or hold decisions made on each security by all agents for every trading day. For example, agent 1 has security 8 as a buy, agent 2 has security 8 as a sell and agent 3 has security 8 as a buy again. This would give security 8 a count of 2 buys and 1 sell for the amalgamation process. This process could loosely be compared to a simple voting system.

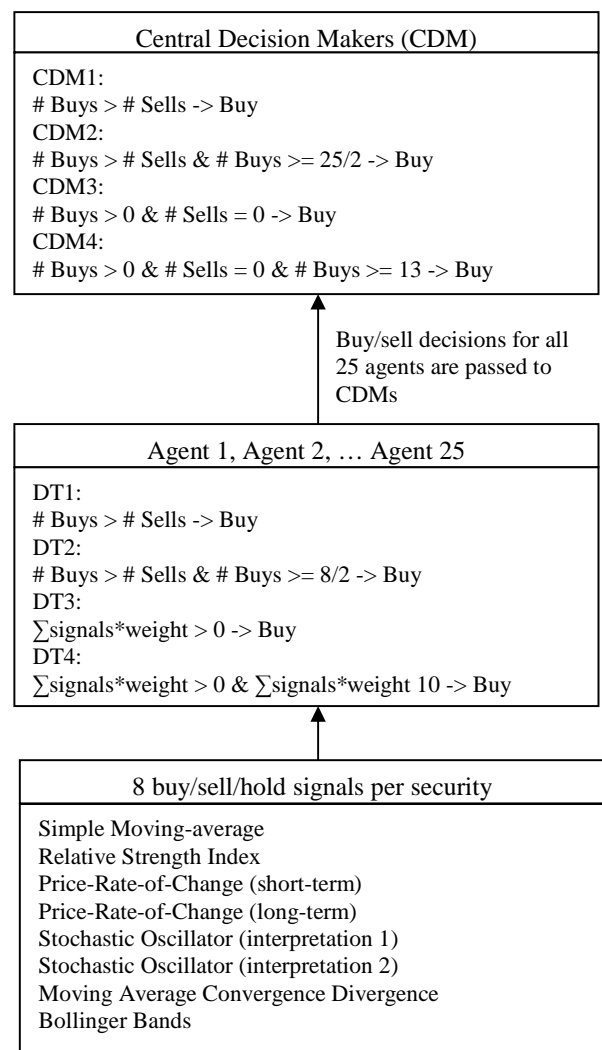


Figure 2. Decision process overview.

- i. The first method (CDM1) performs a simple comparison between the number of overall buy and sell

signals. If either is greater than the other, the corresponding action is adopted for that security.

ii. The second method (CDM2) generates a buy/sell signal if at least half the population generated a buy/sell signal respectively.

iii. The third method (CDM3) will generate a buy/sell signal if at least one agent from the population generates the respective signal and no opposing signals exist for that security.

iv. The fourth method (CDM4) follows the same principle as CDM3, except at least half the population must have generated the same signal for it to generate the same signal.

An overview of the decision process from technical analysis to the central decision making agents is shown in figure 2.

3. Experimentation and results

For comparison, ten agent populations were evolved and tested over the out-of-sample period, with their performance recorded as the experimental benchmark. Four central decision makers were implemented, each corresponding to the four decision making types described earlier. The decisions made by the elite population on

every trading day were used by the central decision makers to generate their own decisions for every experiment run. The following results were obtained, as shown in Table 3.

Table 3. Experiment results (all values x10⁶)

	CDM1	CDM2	CDM3	CDM4	Elite
1	225	254	230	230	290
2	242	225	214	214	303
3	240	271	314	314	306
4	208	215	200	201	301
5	254	209	207	207	302
6	251	277	192	192	301
7	232	226	221	222	318
8	250	251	275	275	314
9	231	208	290	290	300
10	262	230	237	235	321
Mean	240	237	238	238	306
S.D.	15.9	25.1	41.4	41.2	9.4

CDM1, CDM2, CDM3 and CDM4 all showed similar average performance while the Elite demonstrated significantly better results.

Two sets of ten runs were performed, with the first




Figure 2. Performance of Elite, CDMs mean and Buy-and-hold

providing identical starting capital to each CDM as to every trader in the elite (100,000), as shown in table 3, while the second provided each CDM with a starting capital equivalent to the sum of the elite's starting capital (25*100,000). However, the latter experiment showed very similar results that reproduced the ones observed in table 3. In conclusion, incrementing the starting capital at those levels does not affect trading performance and will therefore not be considered further here.

A graphical representation of the averaged performance of the elite and each CDM over the out-of-sample test period is shown in figure 3, also including a comparison of their performance to the buy-and-hold strategy. The buy-and-hold strategy is a common benchmark for performance comparison as it represents an investment strategy where the entire starting capital is invested in its entirety equally into all securities, without making any further trading decisions thereafter. Figure 2 shows that the elite outperformed the buy-and-hold strategy across the entire time span, while the CDMs underperformed for most of the out-of-sample data. The buy-and-hold strategy has a total area under its graph of 270x106, which is also clearly larger than that shown by the CDMs.

4. Analysis and discussion

Comparing the mean from the ten experiments for each CDM and the elite, it is quite apparent that the four amalgamation processes did not produce a phenotype that is as successful as the elite. Furthermore, when compared to the buy-and-hold strategy only the elite offers a successful alternative. Statistical analysis using Kruskal-Wallis compared CDM1 to CDM4 among one another to determine if there exists any significant difference between them. CDM1, CDM2, CDM3 and CDM4 were compared individually to the Elite using Mann-Whitney U Tests to determine if a significant difference exists between the elite and the central decision makers. Results are shown in table 5 and table 6 below.

Table 5. Kruskal-Wallis test

Experiments compared	H-value	p-value
CDM1, CDM2, CDM3, CDM4	1.181	0.7576

Table 6. Mann Whitney U tests

Experiments compared	p-value
Elite CDM1	0.0002
Elite CDM2	0.0002
Elite CDM3	0.0019
Elite CDM4	0.0019

As is clearly apparent, no significant difference in performance exists between the result sets A and B. Furthermore, both demonstrate that a significant difference in performance exists between each amalgamation method and the elite, while no difference exists between the amalgamation methods themselves. Near identical statistical results were observed for the experiments using the alternative starting capital for CDMs. As the individual agents represent already optimized trading strategies, it is not overly surprising that an amalgamation thereof should demonstrate worse performance. Using successful strategies, which can involve fundamentally different styles, for deriving an overall strategy should be anticipated to result in less coherent investment and therefore lower returns. This suggests that in order to benefit from deriving a single strategy or centralizing investment decisions from an already optimized pool of simulated traders, a less homogenous reference population with diversity in trading approaches could enable the centralized system to switch between optimized strategies. Depending on their suitability for the current period being considered, its improved performance margin may be created by trading off the strengths and weaknesses of the strategies. This however needs to be substantiated by further research.

In conclusion, this investigation has demonstrated that distilling a single or centralized decision from a group of individual decision makers using homogenous methods of analysis does not necessarily improve or maintain a comparable level of performance. The emphasis here lies on the amalgamation process itself, as it is the key factor in determining success of the approach and needs to be specifically designed and refined for the application domain.

5. References

- [1] S.B. Achelis, Technical Analysis from A to Z, McGraw-Hill Trade, 2001.
- [2] F. Allen and R. Karjalainen, "Using genetic algorithms to find technical trading rules", J. Financial Economics, 1999, Vol. 51, No. 2, pp. 245-271.
- [3] W.M. van den Bergh, K. Boer, A. de Bruin, U. Kaymak and J.Spronk, "On Intelligent-Agent Based Analysis of Financial Markets", 29th meeting of EURO Working Group on Financial Modelling, 2001.
- [4] J. Branke, "Evolutionary Approaches to Dynamic Optimization Problems - A Survey.", GECCO Workshop on Evolutionary Algorithms for Dynamic Optimization Problems, 1999, 134-137.
- [5] E.F. Fama, "Efficient capital markets: A review of theory and empirical work", J. Finance, 1970, Vol. 25, pp. 383-417.

- [6] D. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley, 1989.
- [7] S. Ontañón and E. Plaza, "Justification-based Multiagent Learning", *Proc. 20th Int. Conf. On Machine Learning (ICML-2003)*, Washington DC, USA, 2003.
- [8] S. Rehfuss, L. Wu and J. Moody, "Trading Committees: A Comparative Study", *Proc. 3rd Int. Conf. of Neural Networks in the Capital Markets*, London, UK, 1995, pp. 11-13.
- [9] S. Schulenburg and P. Ross, "Strength and Money: An LCS Approach to Increasing Returns", Lanzi, P.L., Stolzmann, W., and Wilson, S.W. (ed.), *Advances in Learning Classifier Systems*, Vol. 1996 of *Lecture Notes in Artificial Intelligence*, pp.114-137. Springer-Verlag, Berlin, 2000.
- [10] C. Schoreels and J.M. Garibaldi, "Agent based Genetic Algorithm Employing Financial Technical Analysis for Making Trading Decisions Using Historical Equity Market Data", *Proc. IEEE/WIC/ACM Int'l Conf. on Intelligent Agent Technology 2004 (IAT 2004)*, Beijing, China, 2004, pp. 421-424.
- [11] C. Schoreels and J.M. Garibaldi, "A preliminary investigation into multi-agent trading simulations using a genetic algorithm", *Proc. of the UK Computational Intelligence Workshop 2004 (UKCI 2004)*, Loughborough, UK, 2004, pp. 62-69.
- [12] C. Schoreels and J.M. Garibaldi, "A comparison of adaptive and static agents in equity market trading", *Proc. IEEE/WIC/ACM Int'l Conf. on Intelligent Agent Technology 2005 (IAT 2005)*, Compiègne, France, 2005, pp. 393-399.
- [13] L. Tesfatsion, "Introduction to the special issue on agent-based computational economics", *J. Economic Dynamics & Control*, 2001, Vol. 25, pp. 281-293.
- [14] N.S. Thomaidis and G.D. Dounias, "Behavioural Finance: A New Challenge to Computational Intelligence", *Proc. Of European Symposium on Intelligent Technologies, Hybrid Systems and their implementation on Smart Adaptive Systems (EUNITE 2004)*, Aachen, Germany, 2004, pp. 79-92.
- [15] R. Vahidov and B. Fazlollahi, "Pluralistic multi-agent decision support system: a framework and an empirical test", *Information & Management*, 2004, Vol. 41, Iss. 7, pp. 883-898.