

Evolving a Robust Trader in a Cyclic Double Auction Market

Peter A. Whigham
Department of Information Science
Otago University
Dunedin, New Zealand
peter.whigham@otago.ac.nz

Rasika Withanawasam
Department of Information Science
Otago University
Dunedin, New Zealand
rasika@infoscience.otago.ac.nz

ABSTRACT

A computational model of a double auction market is introduced and extended to allow a controlled cyclic behaviour in the price signal to be developed. Traders are evolved to maximise profit in this market using Grammatical Evolution, and their properties studied for a range of periods and amplitude of the trend in the price signal. The trader grammar allows decision making based on simple trading rules incorporating the concepts of moving-average oscillators and trading range break-out. The results of this investigation demonstrate that traders evolve a short waiting period between decisions, and that there underlying decision logic reflects the scale of the market price frequency. Evidence is presented that suggests evolving a robust profit-making trader, for a range of price frequency changes, requires the training data to have high frequency variation. More generally, to evolve robust solutions for any complex GP problem, a set of local models or an ensemble and state-based approach, is implied by the results.

Categories and Subject Descriptors

I.6 [Simulation and Modelling]: Miscellaneous; I.2 [Artificial Intelligence]: Distributed Artificial Intelligence—*intelligent agents*; J.4 [Social and Behaviour Sciences]: Economics

General Terms

Algorithms, Economics

Keywords

Double auction market, Grammatical Evolution, micro-economic trading, simulation

1. INTRODUCTION

The double auction market is a fundamental stock market trading concept. This market uses a limit order book for both the bid and ask prices, and a trade occurs when

the limit order books are crossed. A number of computational models have been proposed to describe the operation of these order-books and to examine the resulting time series of prices produced by these models. A review and comparison of several double auction models is described by Slanina [16]. This work concluded that, although none of the studied models reproduce all properties observed in real order markets, the Genoa artificial market [13, 14, 12] and the Maslov model [8, 11] were candidates for further investigation. These models are studied because of their simplicity and ability to produce aspects of real trading price signals without having to incorporate explicit knowledge of markets. A more complex approach has been presented by Schwartz et al. [15]. In this work several types of traders are formally represented in the model, and therefore greater control over trading behaviour is possible. Although not considered here, methods for representing knowledge, such as informed trading behaviour, is a clear line of future research in evolutionary computation and economic modelling.

Previous artificial financial market research [13, 14, 5] has considered the characteristics of the log price returns given a range of agent interactions, or the profitability of agent strategies. Evolving trader characteristics [2, 7], using either real data or through agent interactions [13], have also been investigated. A good selection of evolutionary concepts for financial trading may be found in [1].

The model presented here extends the Maslov limit-order model [8], introducing a cyclic pattern of supply and demand to the market that allows the last traded price to move through a controlled increase and decrease over time. An evolved trader interacts with this market, starting with a portfolio of shares and money, and may place limit orders to buy/sell, place a market order to buy/sell or do nothing. The cyclic pattern of prices should allow an evolved trader to identify when the market is increasing/decreasing and to subsequently make a profit. The paper will address the following questions regarding the evolved trader behaviour: what time scale of interaction is evolved for a cyclic market, and does this vary with a change in the cyclic behaviour of the market? How well does an evolved trader for one type of cyclic market perform when the number of periods varies? What is the best training strategy for evolving a trader that can generalise to a market with a range of cyclic behaviour? How does the amplitude of price change alter the profit of an evolved trader? These concepts, although based on an economic model, address the concept of how to evolve a robust GP solution given an environment that varies. The advantage of the model presented here is that

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO'11, July 12–16, 2011, Dublin, Ireland.

Copyright 2011 ACM 978-1-4503-0557-0/11/07 ...\$10.00.

a single, fundamental behaviour (i.e. the rate of change of the trend in prices) can be controlled and varied. This is analogous to changing the scale of a process, and therefore the resulting behaviour should be of general interest. Evolving robust GP solutions by evolving a hedge fund (trading) model has previously been considered by Yan and Clack [18]. However, since they used real data, divided into several different scenarios, they were limited in their possible training approaches. In addition, they had only one set of validation (test) data. Their results suggested that a robust trader has to be trained on a range of data for most generations, with a final generation trained on extreme examples. This suggests that strong selection pressure for the extremes, once the population has evolved towards the general patterns of the training data, can evolve robust individuals.

1.1 The Maslov Double-Auction Market

The Maslov market [8] is as an open market with an infinite pool of traders. Each time step a trader is drawn from this pool and interacts with an order-book to either post a bid/ask offer or to execute a trade (market order). The order book is characterised by a last traded price at time $p(t)$, and an order book for bid/ask prices. The order book for bid/ask (buyer/seller) are stored in decreasing/increasing order and, where two or more bids are equal, orders that are older are given priority. Hence the books are sorted by value and, where more than one value is the same, by time. Each time step a trader is drawn from the pool as a seller with probability p_s or buyer with probability $1 - p_s$. This trader sets a limit order with probability p_{lo} or market order with probability $1 - p_{lo}$. If the trader places a bid/ask the offer is determined as a negative/positive offset Δ from $p(t)$, where Δ is drawn from some distribution. Following Maslov [8] Δ is drawn from the uniform integer distribution $1 \dots 4$, where $\Delta = 1$ represents a single tick price change. If a limit order is placed the last traded price $p(t+1) = p(t)$. If the trader performs a market order to buy/sell the trade is matched with the top entry from the ask/bid order book, resulting in a new last traded price $p(t+1)$, drawn from the top (matched) entry of the corresponding order book. The matched entry from the ask/bid order book is then removed. If no limit order can be matched (i.e. the corresponding order book is empty) the trader places a limit order. This simple model assumes that only a single unit of stock is traded with each transaction.

Previous work on evolving a trader for this market [17] used the Maslov model to generate trading prices and evolved traders that interacted directly with this market. Although the work demonstrated that a trader could be evolved that outperformed random traders, the paper used a market where the long term price changes were zero and there was no control over the cyclic nature of the price signal. In other words, the probability of a buyer or seller being introduced from the pool of traders, p_s , was set to 0.5, resulting in a balance between supply and demand. Hence no explicit control over the long-term pattern of the trading signal could be introduced. In addition, the trader interactions with the system used a fixed number of Maslov traders (i.e. time steps) between each action of the evolved trader, rather than evolving this parameter. Given that the time between decision points in trading decisions represents a scaling factor, related to the market dynamics, it is clear that this property needs to be examined in more detail. The remainder of this paper is

structured as follows: Section 2 describes the properties of a cyclic Maslov model; Section 3 describes the Grammatical Evolution model; Section 3.1 describes the experimental setup and Section 4 presents the results. The implications of the experiments are discussed in Section 5, with concluding remarks and future work considered in Section 6.

2. A CYCLIC MASLOV MARKET

Varying the probability p_s of a seller being selected from the pool of traders will alter the dynamics of the resulting last traded price. This is obvious by considering the effect on a market if there are more people entering the market and creating market orders to buy than sell. In this case more orders are executed at the best "ask" price than the best "bid" price. Since the best "ask" price is typically higher than the last traded price, the price increases over time. Since markets do not just increase forever, some control over the pattern of rising and falling prices is required. A simple approach is to allow p_s to vary using a sine curve where the amplitude controls the maximum/minimum value of p_s , and the period to control the number of periods of buying/selling that occur during a set number of trader interactions. Figure 1 shows a typical run of traded prices for Maslov when p_s is fixed at 0.5. Figure 2 shows the traded price when p_s has two complete periods and an amplitude of 0.1. Note that for the first 1000 timesteps $p_s = 0.5$ to allow the limit-order books to be seeded. Following this initialisation p_s is increased from 0.5 to 0.6 and subsequently down to 0.4. Note that the behaviour of the pricing lags behind the change in selling probability due to the long memory properties of Maslov [8, 11]. Finally, Figure 3 shows the effect of increasing the amplitude to 0.2 over two periods, resulting in an increased range of trading prices with less variability (i.e. smoother transitions) in price. The change in price pattern is reflected in the Hurst exponent (a measure of the autocorrelation) for the price time series, where a constant $p_s = 0.5$ gives 0.30, while the Hurst exponent for Figures 2 and 3 are 0.45 and 0.59 respectively. Given real sections of double-auction price patterns commonly have Hurst exponents between 0.4 and 0.6 [4] suggests that the cyclic model produces aspects of real markets not reflected in the standard Maslov model.

3. THE GE TRADER MODEL

Traders are evolved using Grammatical Evolution (GE) [10] using a modified version of the GEVA Java framework for GE [9]. A simple approach to technical trading rules [3] was used to model the decision-making logic for the evolved trader. These rules considered two cases: the moving average-oscillator and the trading range break-out. The moving average-oscillator compares the short and long term average price. A buy/sell decision is triggered when the short term average is above/below the long-term average. The trade range break-out strategy assumes that prices have a "resistance" to move above/below a previous best/worst price. However, if the price does move above/below this resistance level prices are likely to drift, indicating a time to sell/buy. Brock et. al [3] summarises the concept by stating that specialists recommend "...buying when the price rises above its last peak and selling when the price sinks below its last trough." Both of these trading concepts rely on just a few simple measures: the current price, the average price for




Figure 1: Typical Maslov trading behaviour when sell probability is fixed at 0.5

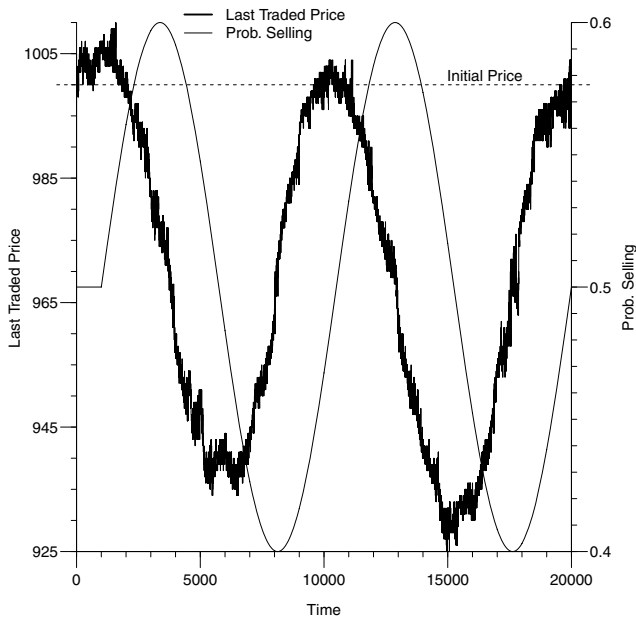


Figure 2: Maslov trading behaviour for two selling periods of amplitude 0.1




Figure 3: Maslov trading behaviour for two selling periods of amplitude 0.2

some number of time steps in the past, and the maximum price over some past time window. In addition to these basic measures, an evolving trader can select the rate at which it interacts with the market. Although a general expression grammar could be used for the trader [17] our initial experiments found that a more profitable trader was evolved when a specific grammar, based on the technical trading rules, was implemented. The following grammar was used for all experiments:

```

<DecisionPath> ::= <WaitTime>;<Option>
<Options> ::= <Buy>;<LOBuy>:<LOBid>;
             <Sell>;<LOSell>:<LOBid>
<WaitTime> ::= 10 | 15 | 20 | 25 | 30 | 35 | 40 |
               45 | 50 | 55 | 60 | 65 | 70 | 75 |
               80 | 85 | 90 | 95 | 100
<Decision> ::= lt <MAO> | lt <TRB> |
              or lt <MAO> lt <TRB> |
              gt <MAO> | gt <TRB> |
              or gt <MAO> gt <TRB>
<Buy> ::= <Decision>
<Sell> ::= <Decision>
<LOBuy> ::= <Decision>
<LOSell> ::= <Decision>
<PriceMeasure> ::= <MAO> | <TRB>
<Num> ::= 10 | 20 | 30 | 40 | 50 | 60 | 70 | 80 |
           90 | 100 | 150 | 200
<MAO> ::= avprice <Num> avprice <Num>
<TRB> ::= price maxprice <Num>
<LOBid> ::= 1 | 2 | 3 | 4

```

The <DecisionPath> represents how the evolved trader is constructed. The <WaitTime> nonterminal is used to select the number of time steps between each interaction with the

Table 1: GE Parameter Settings

Population Size	500
Generations	100
Crossover Prob.	0.9
Mutation Prob.	0.01
Max Depth	6
fixed point crossover	false
tournament size	3
elite size	2
max wraps	3
replacement type	generational
grow probability	0

Table 2: Maslov Parameter Settings

Initial Price	\$1,000
Δ	1...4
p_{lo}	0.5
p_s	varies sinusoidally
Trader Initial Wealth	\$100,000
Initial portfolio	50 shares,\$50,000
Allow Limit Order Books to empty	No
Minimum Limit Order Depth	3
Initialisation	1000 steps
Total market operations	19000

Maslov market, while <Options> lists the possible actions the trader can take. The options are tested in order, hence the decision to <Buy> is first tested, and if true this action is taken. If this is false, the option to place a limit order to buy is tested, and so on. Note that if a limit order is set, Δ is based on the fixed constant from <LOBid>. If no option is true then no action is taken. Note that the grammar allows single or multiple tests using logical "or", with the base measures of moving averages <MAO> or trade range breakouts <TRB> being less than (lt) or greater than (gt) the corresponding measures. The parameters for GE are shown in Table 1.

3.1 Experimental Design

The Maslov double-auction market requires several parameters, as shown in Table 2. Evolved traders commence with an initial portfolio of money and shares and then interact with a single limit-order book. Note that the limit order book is not allowed to empty - this is a result of the fact that when there are many more buyers than sellers (or vice-versa) the corresponding book may be drained. Although this situation can occur in real markets, it is generally when unusual market conditions (such as the 1929 New York stock market collapse [6]) exist. Since we are interested in the behaviour for a stable (but varying) market, this condition is avoided by ensuring prior to a Maslov trader operating, that at least three entries exist in both books. If this is not the case, dummy traders are introduced to enter limit orders to fill the books to this minimum depth. At the end of the market operations the wealth (fitness) of the trader is assessed by the sum of cash and value of shares. A conservative estimate of share value is taken by using the current last traded price for all shares held. Although this is not realistic (let alone that no transaction costs are considered) all traders

are treated equally and therefore the relative performance between traders is meaningful.

A number of experiments were used to address the questions presented in Section 1. Initially a GE trader will be evolved over 30 independent runs with p_s amplitude 0.1 and the number of periods ranging from 1...5. This will show, for a set variation in price and range of periods, whether there is an advantage to an evolved trader when the market increases in cyclic frequency. The best individual from the 30 runs, for each cycle, will be selected to measure generalisation. These five traders will be measured over 100 runs against the Maslov market for the range of cycles from 1...5 with amplitude 0.1. This will allow empirical results to understand whether a trader can be generalised to more than one frequency in the market. In addition, the evolved parameter <WaitTime> and the numbers used for average and maxprice measures may vary in accordance with the cycle frequency of the last traded price. Examining the values used by the best evolved traders for each period may suggest whether scaling is apparent in these solutions. The following random traders will also be measured for profit making over 100 runs: randomly buy/sell; randomly set a limit order with Δ randomly set between 1...4; and randomly buy/sell/limit order or do nothing. The random traders will use the evolved <WaitTime> value from the best individual for each cycle 1...5.

4. RESULTS

Figure 4 shows the mean \pm standard deviation performance over 30 independent runs of GE when the number of cycles is one and the amplitude is 0.1. Note that the best and mean profit (fitness) increases most rapidly over the first 20 generations, and has started to level out once 50 generations are reached. The best profit is consistently around \$25,000 with a mean profit of the population approaching \$15,000. Our first question is to examine how the traders are making money over this single period, and whether there is evidence of scaling related to the price variation. The best period one trader had the following behaviour:

```

<WaitTime> = 10
If (avprice 200 < avprice 10) Then Buy
If (price > maxprice 100) Then Limit Buy: 1
If (price > maxprice 70) Then Sell
If (price < maxprice 20) Then Limit Sell: 1
Else do nothing

```

Note that the number of Maslov iterations before the trader considers an action has evolved to be the lowest number available (i.e. 10). The trader buys if the average price over the past 200 time steps is less than the average price over the past 10 time steps - hence if the market is beginning to rise. The trader sells if the last traded price is greater than the previous prices for 70 time steps, representing a break in a price threshold. A limit order to sell occurs if the last traded price is less than the maximum price in the past 20 transactions, indicating that the market is starting to fall. Note that for both limit orders $\Delta = 1$ which ensures that the order will be match as soon as possible. The behaviour of the trader is clearly shown in Figure 5. The trader recognises that the market is falling and rapidly sells all shares, and once the market begins to rise again as many shares as possible are purchased. Towards the end of the transactions the market begins to drop again, and the trader sells




Figure 4: Best and mean fitness averaged over 30 runs for 1 period with amplitude = 0.1

some shares. The final profit is close to \$20,000. During this time period there were 437 buys, 1 limit order to buy, 2 sells, 367 limit orders to sell, and 1094 times when no action was taken. The trader behaviour is shown in detail in Figure 6. Here the buy and limit order sell actions are shown when the price is turning from a decreasing to increasing signal. Note the greatly increased number of buy actions that occur once the market starts to rise. The intermittent buy and sell actions indicate that the trader is operating on a short time scale and is reacting to minor movements in the market.

Figure 7 shows the mean performance, averaged over 30 independent GE runs, as the number of periods is increased from 1 to 5. The mean fitness clearly decreases as the number of periods increase, while the best fitness does not show any clear pattern other than the five period trader has lowest profit for all generations. This is probably a result of the noise inherent in the Maslov process, which means that for any one run it is possible that a trader does well, even though given a second evaluation it may not produce a similar profit. The surprising aspect of this result is that the best fitness occurs for 3 periods, even though the highest average profit is for one period. Given a market that moves up and down several times presents additional opportunities to make a profit, through judicious buying and selling, the three cycle trader may have discovered the appropriate frequency to operate, but the noise in the process was generally too difficult for all of the population to perform well. This lack of opportunity is also supported by the one period trader having a lower best fitness than the period two-four traders. However, once the number of periods increases to five it appears that the best and average performance is relatively poor. This will be due to the more rapid change in price signal and the interaction with the long memory of the limit-order books, resulting in a noisy signal and less clear patterning of trends. The evolved value for $\langle \text{WaitTime} \rangle$

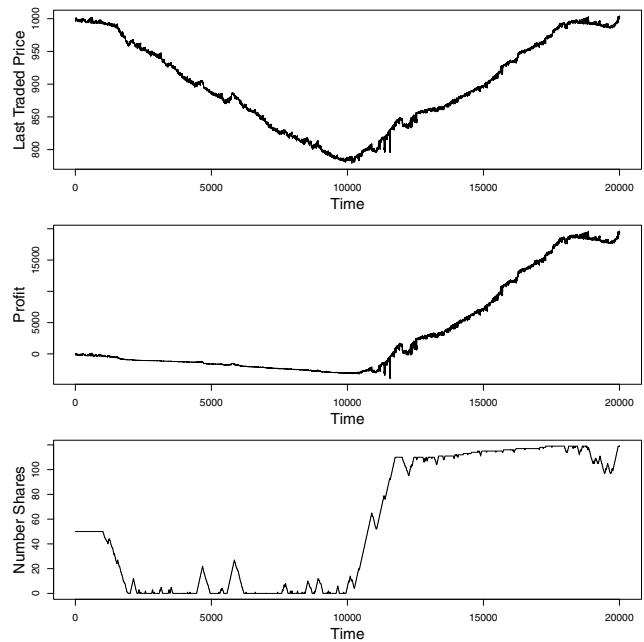


Figure 5: Example run of evolved trader for 1 period with amplitude = 0.1




Figure 6: Trading actions at the turning point of the market

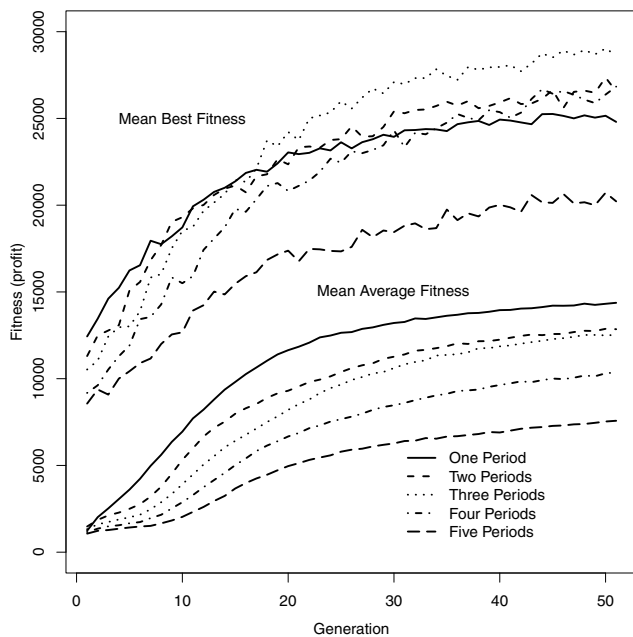


Figure 7: Mean behaviour over 30 runs evolving traders for 1-5 periods

was always 10 for all final evolved traders, for all periods. This indicates that being able to make decisions at the finest temporal scale is the most profitable approach.

The question of whether an evolved trader is biased by the number of periods in the Maslov process, the best evolved trader from each of the 30 runs for each period was selected and run 100 times with a range of periods. The mean profit for each of these runs is shown in Figure 8. Clearly those evolved traders trained with more periods in the Maslov model are better able to handle variability as the number of periods increases. Note, however, that the single (one) period trader outperforms all other traders when the period is one, suggesting that it has specialised to the scale and behaviour of this price signal. Examining the logic for the best period one trader (shown previously) the window of comparison with the price signal is quite large. For example, the buy decision is based on the average price over the past 200 trading steps. In addition, the current price is used for all decisions except for the first buy. This is likely to be a good strategy only if the price has a consistent downward or upward trend, while measuring average price behaviour over small windows is likely to be a better strategy when the price signal changes more rapidly. Examining the best traders for periods two and upwards this use of smaller windows based on average prices is confirmed. For example, the best trader for period three is:

```
<WaitTime> = 10
if (avprice 90 > avprice 200) OR
  (price > maxprice 50) Then Buy
if (price > maxprice 100) Then Limit Buy: 2
if (avprice 10 < avprice 40) Then Sell
if (avprice 70 < avprice 30) OR
  (price < maxprice 80) Then Limit Sell: 1
Else do nothing
```

This trader uses average price for a number of decisions, and for selling considers time windows of length around 30. The use of averaging to smooth the signal, combined with smaller window sizes for decision making, has meant that the evolved trader can detect changes in the price signal that represent genuine trends of rising/falling in the market. One additional comment is that the five period trader is the only evolved solution that makes a profit over all periods. Hence, although it does not make the best profit for periods less than five, it has managed to handle the noisy and changeable price signals for greater periods than those presented during training.

In addition to the evolved traders, Figure 8 shows the profit performance for random traders. Each random trader entered the market every 10 timesteps, since this was the $\langle \text{WaitTime} \rangle$ evolved by all traders. This allowed a fair comparison between the random and evolved traders. Here as the period increases the random limit order and random models approach a zero profit and outperform the evolved traders other than the period 5 trader.

The final behaviour to examine is the change in profit as the amplitude of the period is varied. Figure 9 shows the profit for the best evolved trader for 2 periods when applied to a Maslov model of period 2 with the amplitude varying from 0.0 to 0.3. Several points are worth noting. When the amplitude is zero, the Maslov model behaves without any imposed cycles, and therefore tends to rise and fall over short periods (since the probability of a buyer or seller being introduced from the pool of traders is equal). In this case the trader cannot distinguish any pattern, and the learnt behaviour regarding trends does not exist. The Hurst exponent for a Maslov price signal when the amplitude is zero is approximately 0.3. This implies that, although the price signal is not random, there is anti-persistence to the signal. Hence, if prices have moved down in the past few timesteps, they are likely to rise in the next few steps. This rapid oscillation of prices means that the evolved trader assumptions (i.e. that prices tend to follow a trend) is not observed, and therefore profits cannot be made. However, once the amplitude increases to beyond 0.06 the pattern of trends is evident in the price signal, and the trader can start to make a reliable profit. There appears to also be a linear trend in profit after the amplitude passes 0.06, with a profit increase of \$10,000 per 0.05 increase in amplitude.

5. DISCUSSION

The questions posed in the introduction can now be considered. The first issue is with scaling, where it is clear that the evolved traders do detect the frequency of the price signal and specialise to this behaviour. This is particularly apparent by considering the performance shown in Figure 8 and by examining the best evolved trader for several periods. Clearly if you want to evolve a trader that can generalise to a range of cycles the trader must be trained with a rapidly varying price signal. In addition, since all evolved traders used the lowest $\langle \text{WaitTime} \rangle$ of 10, it is clear that there is an advantage to examining the state of the market place as often as possible. In terms of performance, a trader evolved over one period will perform poorly once the number of periods increases, while a trader trained with five periods will not perform well over one period (but still make a small profit). A similar decrease in performance is noted for all traders evolved with periods greater than one. This




Figure 8: Mean behaviour for the best evolved trader for each period, and random traders, measured for periods 1-10




Figure 9: Mean profit for the best evolved 2 period trader for a 2 period Maslov model with varying amplitude

indicates that the one period model is a special case that can only be exploited by being exposed to this pattern. Future work could therefore consider training scenarios with extreme values [18] for later generations, such as single frequency and a very high frequency rates of change in price trend.

In terms of training strategy, to produce a trader that can generalise to a range of periods requires the GE trader to evolve under a high frequency market. This evolves smaller scaling windows to be used for the decision making, and therefore means that profit making can occur over a short duration of price change. However, there is a cost to this behaviour, since it cannot exploit more regular trends in the price signal and therefore does not perform well in slowly varying markets. No doubt a more profitable behaviour would be to have a trader that could identify which type of cyclic behaviour was occurring, and use this to select an appropriate trading strategy. This form of local state-based modelling would require a more complex grammar, or an ensemble-style learning model. In terms of developing a general, robust solution for varying environments in GP, the results imply that no good generalisation will outperform a specialised model. If an environment is non-stationary, a range of different models is most likely to be appropriate. Finally, as the amplitude of price change varies there is evidence from the one experiment performed here that a linear relationship exists between profit and change in amplitude. Although this intuitively makes sense this has only been performed for a single 2 period trader and therefore may not be a general statement for all period models.

6. CONCLUSIONS

This paper has presented an extension to the Maslov limit-order market model and examined the behaviour of evolved traders using Grammatical Evolution. The price signal was varied both in frequency and amplitude, and a range of evolved traders examined for these variations in market behaviour. A number of properties of the evolved traders have been shown, including the selection of fast sampling rates when making decisions about the next trade, and the generalisation ability of traders based on the price period used for training. The evidence suggests that to evolve a robust profit-making trader over a range of price behaviours it is necessary to use a high frequency price signal. In addition, it has been shown that changing the amplitude of the supply/demand balance has a linear effect on profit, however this has only been shown under one circumstance and therefore general statements about profit and amplitude cannot be made.

This is the first presented work to consider a controlled, cyclic double auction trading market and examine the properties of evolved traders over a range of conditions. As such, this is a preliminary study and suggests a number of future directions. For example, what is an appropriate training scheme for evolving a robust trader? Given the results regarding a five period trader, it may be appropriate to consider training an individual by presenting several different periods during a trading session. This would mean that a trader would have to be able to recognise a change in scale of the price signal, and to develop general measures of identifying appropriate buy and sell situations. In addition, the language used for the traders was based on very simple trading rules. Clearly further work is required to investigate

other measures and more expressive grammars to promote more sophisticated decision rules. There is also the issue of using a rule-based approach or genetic programming approach, versus some other form of model. Given the time series nature of the problem, evolving the weights for a recurrent neural network, or the use of a finite-state machine or classifier system, may also allow interesting and efficient traders to be evolved. Finally, the price signal from the cyclic Maslov model needs to be studied in more detail to understand how the scaling of the frequency and amplitude signal of p_s changes the pattern of last traded price. In particular, a comparison against real data is required to relate p_s to different market behaviours, volumes of trade, return periods and market dynamics.

7. ACKNOWLEDGMENTS

The authors would like to thank Timothy Crack and Grant Dick for useful discussions during the period of this work. The anonymous referees must also be thanked for their contribution.

8. REFERENCES

- [1] A. Brabazon, A. Brabazon, and M. O'Neill. *Natural Computing in Computational Finance: Volume 2*. Springer Publishing Company, Incorporated, 2009.
- [2] A. Brabazon and M. O'Neill. Evolving technical trading rules for spot foreign-exchange markets using grammatical evolution. *Computational Management Science*, 1:311–327, 2004.
- [3] W. Brock, J. Lakonishok, and B. LeBaron. Simple technical trading rules and the stochastic properties of stock returns. *Journal of Finance*, XLVII(5):1731–1764, 1992.
- [4] A. Carbone, G. Castelli, and H. Stanley. Time-dependent hurst exponent in financial time series. *Physica A*, 344:267–271, 2004.
- [5] S. Cincotti, S. M. Focardi, M. Marchesi, and M. Raberto. Who wins? study of long-run trader survival in an artificial stock market. *Physica A*, 324:227–233, 2003.
- [6] J. Galbraith. *The Great Crash 1929*. Pelican, 1961.
- [7] M. Lim and R. J. Coggins. Optimal trade execution: an evolutionary approach. In *IEEE Congress on Evolutionary Computation*, pages 1045–1052. IEEE, 2005.
- [8] S. Maslov. Simple model of a limit order-driven market. *Physica A*, 278:571–578, 2000.
- [9] M. O'Neill, E. Hemberg, C. Gilligan, E. Bartley, J. McDermott, and A. Brabazon. Geva: grammatical evolution in java. *SIGEVolution*, 3(2):17–22, 2008.
- [10] M. O'Neill and C. Ryan. Grammatical evolution. *IEEE Transactions on Evolutionary Computation*, 5(4):349–358, Aug. 2001.
- [11] T. Qui and G. Chen. Spread and quote-update frequency of the limit-order driven sergei maslov model. *Chin.Phys.Lett.*, 24(8):2161–2163, 2007.
- [12] M. Raberto and Cincotti. Modeling and simulation of a double auction artificial financial market. *Physica A*, 355:34–45, 2005.
- [13] M. Raberto, S. Cincotti, S. M. Focardi, and M. Marchesi. Agent-based simulation of a financial market. *Physica A*, 299:319–327, 2001.
- [14] M. Raberto, S. Cincotti, S. M. Focardi, and M. Marchesi. Traders' long-run wealth in an artificial financial market. *Computational Economics*, 22:255–272, 2003.
- [15] R. A. Schwartz, R. Francioni, and B. W. Weber. Decision making in equity trading: Using simulation to get a grip. *The Journal of Trading*, 1(1):59–74, 2006.
- [16] F. Slanina. Critical comparison of several order-book models for stock-market fluctuations. *Eur. Phys. J. B*, 61:225–240, 2008.
- [17] P. A. Whigham, R. Withanawasam, T. Crack, and I. Premachandra. Evolving trading strategies for a limit-order market. In *WCCI 2010 IEEE World Congress on Computational Intelligence*, pages 2467–2474, 2010.
- [18] W. Yan and C. D. Clack. Evolving robust GP solutions for hedge fund stock selection in emerging markets. In D. Thierens, H.-G. Beyer, J. Bongard, J. Branke, J. A. Clark, D. Cliff, C. B. Congdon, K. Deb, B. Doerr, T. Kovacs, S. Kumar, J. F. Miller, J. Moore, F. Neumann, M. Pelikan, R. Poli, K. Sastry, K. O. Stanley, T. Stutzle, R. A. Watson, and I. Wegener, editors, *GECCO '07: Proceedings of the 9th annual conference on Genetic and evolutionary computation*, volume 2, pages 2234–2241, London, 7-11 July 2007. ACM Press.