Chapter 13

CO-EVOLVING TRADING STRATEGIES TO ANALYZE BOUNDED RATIONALITY IN DOUBLE AUCTION MARKETS

Shu-Heng Chen[1], Ren-Jie Zeng[1] and Tina Yu[2]

[1]*National Chengchi University, Taipei, Taiwan 116.*

[2]*Memorial University of Newfoundland, St. John's, NL A1B 3X5 Canada.*

Abstract We investigate double-auction (DA) market behavior under traders with different degrees of *rationality* (*intelligence* or *cognitive ability*). The rationality of decision making is implemented using genetic programming (GP), where each trader evolves a population of strategies to conduct an auction. By assigning the GP traders different population sizes to differentiate their cognitive ability, through a series of simulations, we find that increasing the traders' intelligence improves the market's efficiency. However, increasing the number of intelligent traders in the market leads to a decline in the market's efficiency. By analyzing the individual GP traders' strategies and their co-evolution dynamics, we provide explanations to these emerging market phenomena. While auction markets are gaining popularity on the Internet, the insights can help market designers devise robust and efficient auction e-markets.

Keywords: bounded rationality, zero-intelligence, agent-based modeling, human subject experiments, auction markets design, double-auction markets, macroeconomics, trading strategies, software agents, market simulation, market efficiency.

1. Introduction

While genetic programming (GP) has been applied to solve various practical problems in the pursuit of financial rewards, there are other types of applications where GP has made important non-financial contributions. This chapter presents one such kind of application where GP has been used to co-evolve trading strategies in an artificial double-auction (DA) market. The objective of the study is to understand the DA market behavior under traders with various degrees of *rationality* (*intelligence* or *cognitive ability*) as well as the learning

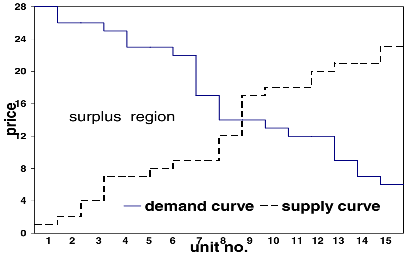

Figure 13-1. Demand and supply curves.

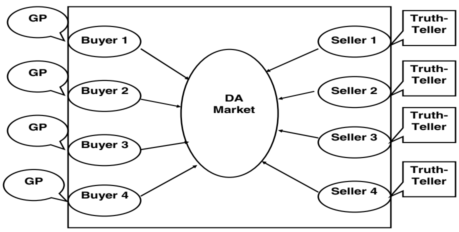

Figure 13-2. The 8 traders' DA market.

behavior of individual traders. Such insights are important to the design of robust and efficient real-world DA markets, where the maximum possible profit is realized. As the number of Internet auction markets (e.g., eBay and Amazon) increases and these markets become more diverse, this GP application can enhance the design of future auction e-markets.

The remainder of this chapter is organized as follows. Section 2 explains DA markets and Section 3 discusses the motivations for this work. In Section 4, our approach to using GP population size to implement traders' rationality is presented. Section 5 details the setups needed to carry out the experiments. The DA aggregate market behavior results are then discussed in Section 6. We analyze the trading strategies that GP evolved in Section 7. Finally, Section 8 concludes the chapter.

2. Market Efficiency and Double Auction Markets

In a standard market environment, the demand for and supply of a commodity can be satisfied under different market prices. When the price is high, the demand decreases and when the price is low, the supply decreases. Adam Smith's *invisible hand* will determine the market price such that both demand and supply are satisfied.

Figure 13-1 gives the demand and supply curves of a standard market. The demand curve gives the maximum price that consumers can accept for buying a given commodity (e.g. cotton, electricity), and the supply curve gives the minimum price at which the producers are willing to sell that commodity. For example, in Figure 13-1, the maximum price that buyers can pay for the second unit is 26, and the minimum price that sellers would accept is 2.

Transactions only take place when the market price is below the demand curve and above the supply curve. The area between the demand and supply curves is the surplus region generated by the transactions. For example, for the first unit, the highest acceptable price for the buyers is 28, and the lowest acceptable price for the sellers is 1. So, the difference between the two (surplus) is $28 - 1 = 27$. Similarly, for the second unit, the highest acceptable price for the buyers is 26, and the lowest acceptable price for the sellers is 2; hence the surplus is $26-2 = 24$. The distribution of the surplus depends on the transaction price, which must lie between the two acceptable limits. For example, if the transaction price of the first unit is 20, the surplus distributed to the consumer is $28 - 20 = 8$, and the rest, $20 - 1 = 19$, is distributed to the producer.

In an efficient market, a commodity's price is between the demand and supply curves so that all potential surpluses can be fully realized. When the market price is outside the curves, no transaction can occur. Consequently, the commodity stays with the sellers leaving consumers unsatisfied. This market is not efficient. DA is one type of market structure that results in high market efficiency, and hence is very popular in the world. For example, the New York Stock Exchange (NYSE) and the Chicago Mercantile Exchange are organized as DA markets.

In a DA market, both buyers and sellers can submit bids and asks. This is opposed to only buyers shouting bids (as in an *English Auction*) or only sellers shouting asks (as in a *Dutch Auction*). There are several variations of DA markets. One example is the *clearinghouse* DA of the Santa Fe Token Exchange (SFTE) (Rust et al., 1993) on which this work is based.

On the SFTE platform, time is discretized into alternating *bid/ask* (BA) and *buy/sell* (BS) steps. Initially, the DA market opens with a BA step in which all traders are allowed to simultaneously post bids and asks. After the clearinghouse informs the traders of each others' bids and asks, the holders of the *highest bid* and *lowest ask* are matched and enter into a BS step. During the BS step, the two matched traders carry out the transaction using the *mid-point* between the *highest bid* and the *lowest ask* as the transaction price. Once the transaction is cleared, the market enters into a BA stage for the next auction round. The DA market operations are a series of alternating BA and BS steps.

3. Motivation and Related Work

Economic markets are complex systems where the market behavior is the result of the aggregate of many traders' behavior. Due to the interdependence of the traders' actions, it is not easy to isolate or analyze any individual trader's auction strategies based on the market's overall performance. However, understand these trading strategies can help identify and isolate the causes of market inefficiency. Recently, simulation approaches have been developed to study individual traders behavior in economic markets.

When conducting DA simulations for market analysis, human traders may reflect the real market more closely than software agents. However, when the study involves traders' rationality, there is no consensus on how to represent and measure human cognitive ability. One simple approach is using traders' background information, such as education and experiences, to determine their intelligence. Unfortunately, studies have shown that such information does not reflect human rationality in decision making (Duffy, 2006). By contrast, software agents, such as GP systems, can represent traders with different learning characteristics through different parameter settings. Recently, more and more studies have been using software agents to model and simulate DA markets, through an approach called agent-based modeling (ABM). Both ABM and human subject experiments implement controlled "laboratory" conditions to isolate the sources of aggregated market phenomena. The two approaches complement each other and provide information to advance our knowledge of DA market properties (Duffy, 2006; Chen and Tai, 2003).

The interest in investigating DA market behavior under traders with different degrees of rationality was strongly influenced by the concept of *bounded rationality* introduced by Herbert Simon. Prior to that work, much theoretical work on economics was based on the assumption that individuals are *perfectly rational* in their decision making. This assumption is invalid as humans engage in "rational choice that takes into account the cognitive limitations of the decision-maker, limitations of both knowledge and computational capacity" (Simon, 1997). Due to the bounded rationality, humans make suboptimal yet acceptable decisions by creating partial strategies (simple heuristics) based on imperfect information.

Gode and Sunder (Gode and Sunder, 1993) were among the first to study the impact of bounded rationality on DA market performance. In that work, traders had zero-intelligence in that all bids and offers were randomly generated within their budget constraints (i.e., traders were not permitted to sell below their costs or buy above their values). The DA market with only zero-intelligence traders was able to achieve almost 100% market efficiency. Based on these experimental results, Gode and Sunder argued that the rationality of individual traders accounts for a relatively small fraction of the overall market efficiency.

To investigate the generality of Gode and Sunder's result, ABM researchers have conducted similar experiments using zero-intelligence traders in various types of DA markets. For example, Cliff and Burten (Cliff and Bruten, 1997) studied a DA market with asymmetric supply and demand curves. They found that zero-intelligence traders gave rise to poor market efficiency. They then assigned the traders with the ability to use the closing price in the previous auction round to determine the current bid. Such traders, whom they referred to as zero-intelligence-plus, performed better and improved market efficiency. Thus, individual traders' cognitive ability does impact overall market efficiency.

In this work, we are interested in understanding market behavior under traders with different degrees of intelligence. Would increasing traders' intelligence improve the market efficiency? Does a trader learn better in the presence of another intelligent trader? Is a trader using strategies that exploit the market in an unfair manner? The answers to these questions can help DA market designers implement rules, such as transaction pricing and the order of bid and offer, so that the market is fair and efficient under traders with different qualities.

4. Methodology

Although the representation of human intelligence is still a debatable issue, we find GP population size to be a reasonable parameter indicating a trader's cognitive ability. In recalling the *learning from experience* analogy of (Arthur, 1993), a trader's strategies are influenced by two factors: the trader's original ideas of how to bid/ask, and the experiences he/she learned during the auction process. In GP learning, the population is the brain that contains the possible strategies to be used for the next bid/ask. It is reasonable to argue that a bigger population size gives the learner a larger reservoir to store and process new strategies. We therefore assign traders with different population sizes to differentiate their cognitive ability when evolving trading strategies.

The trading strategies are represented as rules. We provide three types of information for GP to construct these rules (see Table 13-1):

- Past experiences: terminals 1 – 9 and 16 – 17
- Time information: terminals 10 – 11
- Private information: terminals 12 – 14

In this implementation, only transaction information on the previous day/auction round is provided for GP to compose bidding strategies. In our future work, we will investigate if more profitable strategies will evolve when GP is provided with longer memory.

The three types of information are combined using logical and mathematical operators to decide the bidding prices. Table 13-2 lists these operators.

Each GP trader has a population of strategies to use during an auction. These strategies evolve during GP experimental runs. Each GP generation consists of multiple days and on each day one strategy is randomly selected to carry out a series of auction rounds. The fitness of a strategy is the surplus it accumulated from the transactions at the end of the day when all auction rounds were completed. Since the number of days in each generation is twice the population size, each strategy is likely to be selected at least once and has a fitness value.

The quality of the strategy depends on how well it plays with other traders' strategies during the auction. When multiple self-interest GP traders are in the market, they co-evolve their strategies with the goal being to realize the most surplus. Co-evolution dynamics is not easy to control and analyze. A healthy co-evolution dynamics would develop an "arms race" spiral in which each population spurs the other(s) to advance and the result is continuous learning for all populations. However, this does not always happen. Sometimes, the improvement in one population leads to the deterioration of other populations.

Table 13-1. Terminal Set.

Index	*Terminal*	*Interpretation*
1	PMax	the highest transaction price on the previous day
2	PMin	the lowest transaction price on the previous day
3	PAvg	the average transaction price on the previous day
4	PMaxBid	the highest bidding price on the previous day
5	PMinBid	the lowest bidding price on the previous day
6	PAvgBid	the average bidding price on the previous day
7	PMaxAsk	the highest asking price on the previous day
8	PMinAsk	the lowest asking price on the previous day
9	PAvgAsk	the average asking price on the previous day
10	Time1	the number of auction rounds left for today
11	Time2	the number of auction rounds that have no transaction
12	HTV	the highest token value
13	NTV	the second highest token value
14	LTV	the lowest token value
15	Pass	pass the current auction round
16	CASK	the lowest asking price in the previous auction round
17	CBID	the highest bidding price in the previous auction round
18	Constant	randomly generated constant number

Table 13-2. Function Set.

Function	*Function*	*Function*	*Function*	*Function*
+	-	*	%	min
>	exp	abs	log	max
sin	cos	if-then-else	if-bigger-then-else	

It is also possible that all populations will settle into a mediocre stable state, reaching a local optimum and being unable to move beyond it (Koza, 1992).

In this research, we will investigate the global market outcome produced as a result of the co-evolution of multiple GP traders (see Section 6). In addition, the changes in behavior of an individual GP trader due to his/her co-evolution with another GP trader will be analyzed (see Section 7).

The DA Market Environment

The simulation was conducted using the DA system implemented in (Chen and Tai, 2003). We have customized the system in the following areas to meet our research goals:

Table 13-3. Token Value Table.

buyer 1	buyer 2	buyer 3	buyer 4	seller 1	seller 2	seller 3	seller 4
23	28	23	26	2	8	1	7
22	26	14	25	7	12	4	9
12	17	9	14	17	21	9	20
7	13	6	12	18	23	18	21

- There are 4 buyers and 4 sellers, each of whom is assigned 4 private token values for trading. For buyers, these are the 4 highest prices that they are willing to pay to purchase 4 tokens and, for sellers, they are the 4 lowest prices that they accept to sell these tokens. Table 13-3 lists these values.
- We regard all sellers as being truth-teller, who always gave the assigned true token value during an auction. For buyers, however, we have varied the number of GP buyers from 1 to 4. In this way, we can study the macro market behavior and micro individual behavior changes under the co-evolution of a different number of GP buyers. Figure 13-2 gives the setup of this market design.

Our DA market design is based on the *clearinghouse* DA of the Santa Fe Token Exchange (SFTE). Each DA market simulation is carried out with a fixed number of GP generations (g), where each generation lasts n ($n = 2 \times pop\_size$) days. On each day, 4 new tokens are assigned to each of the buyers and sellers. The 8 traders then start the auction rounds to trade the 16 tokens. A buyer would start from the one with the highest price and then moving to the lower priced ones while a seller would start from the one with the lowest price and then move to the higher priced ones. The day ends when either all 16 tokens have been successfully traded or the maximum number of 25 auction rounds is reached. Any un-traded tokens (due to no matching price) will be cleared at the end of each day. The following day will start with a new set of 16 tokens.

During each auction round, if the trader is a truth-teller, the bid/ask price is always the current token value (HTV for buyers and LTV for sellers). If the trader uses GP to evolve strategies, one randomly selected strategy from the population will be used to decide the bidding price. The GP trader may decide to $pass$ the round without giving a bid. The same 8 strategies will play 25 auction rounds, during which GP traders may give different bidding price if the selected strategy uses information from the previous round/day.

Once all 8 traders have presented their prices, the highest bid is compared with the lowest ask. If the bid is higher than the ask, there is a match and the transaction takes place using the average of the bid and ask as the final price. The profits for these 2 strategies (the difference between the transaction and the given private token prices) are recorded. The fitness of a strategy is the accumulated profit from the transactions of its 4 tokens during the 25 auction rounds.

$$F = \sum_{i=1}^{m} |TokenValue_i - TransactionPrice_i|,$$

where m is the number of tokens traded using the strategy. Since one strategy is randomly selected each day to carry out the auction, after n (n= 2 $\times$ pop_size) days, each strategy in the population will most likely be selected at least once and has a fitness value. This fitness value decides how each strategy will be selected and mutated to generate the next generation of new strategies.

When evaluating a GP trading strategy, it is possible that the output price is outside the private price range. That is, the traders may sell below the cost or buy above the value of a token. We might interpret this as a strategic way to win the auction. Since the final price is the average of the winning bid and ask, the trader might still make a profit from the transaction. However, such a risk-taking approach has been shown to make the market unstable and to reduce the market efficiency (Chen and Tai, 2003). We therefore enforce the following rule on the price generated by a GP trading strategy.

if Bid $>$ 2$\times$HTV then Bid=HTV

if Ask $<$ (1/2)$\times$HTV then Ask=HTV

These rules protect the market from becoming too volatile and also allow GP to evolve rules that take on a small amount of risk to generate a profit.

Market Evaluation Criteria

In this work, we analyze the following 8 market characteristics based on the daily transactions results. Let $V_{b,i}$ be buyer b's private value of token i, $V_{s,i}$ be seller s's private value of token i and P_i be token i's transaction price. Also, the potential daily surplus (PDS) is the daily market profit when all traders are truth-tellers who use the assigned token values to conduct the auction. This is the maximum possible daily market profit from all transactions.

1. ***Individual Profit***: The accumulated daily profit of each individual.

   - The profit of buyer b is:

     $\pi_b = \sum(V_{b,i} - P_i)$, **i is the token number b bought.**

   - The profit of seller s is:

     $\pi_s = \sum(P_i - V_{s,i})$, **i is the token number s sold.**

2. ***Consumer Surplus(CS):*** The total daily profit of all 4 buyers.

   $CS = \sum \pi_b$, b=1,2,3,4

Table 13-4. GP Parameters.

Parameter	*Value*	*Parameter*	*Value*
tournament selection size	5	elitism size	1
Initialization method	grow	maximum tree depth	5
population size	10,20,30,40,50	number of days	2×pop_size
crossover rate	1.0	subtree mutation rate	0.005
no. of generation	300	point mutation rate	0.045
no. of runs for each setup	90	no. of buyers that are GP trader	1,2,3,4

3. ***Producer Surplus(PS)***: The total daily profit of all 4 sellers.

   $PS = \sum \pi_s$, s=1,2,3,4

4. ***Total Surplus(TS)***: The daily profit of all transactions.

   TS=CS+PS

5. ***Individual Realized Surplus Ratio (IR)***: The proportion of the potential daily surplus (PDS) that is realized by each individual.

   Buyer b's IR=$\frac{\pi_b}{PDS}$
   Seller s's IR=$\frac{\pi_s}{PDS}$

6. ***Consumer Realized Surplus Ratio (CSR)***: The proportion of the potential daily surplus (PDS) that is realized by all 4 buyers.

   CSR=$\frac{CS}{PDS}$

7. ***Producer Realized Surplus Ratio (PSR)***: The proportion of the potential daily surplus (PDS) that is realized by all 4 sellers.

   PSR=$\frac{PS}{PDS}$

8. ***Total Realized Surplus Ratio(TR)***: The proportion of the potential daily surplus (PDS) that is realized by all transactions.

   TR=$\frac{TS}{PDS}$

5. Experimental Setup

Table 13-4 gives the GP parameter values used to perform simulation runs. With 5 different population sizes and 4 different ways to assign GP buyers, the total number of setups is 20. For each setup, we made 90 runs. The total number of simulation runs made was 1,800.

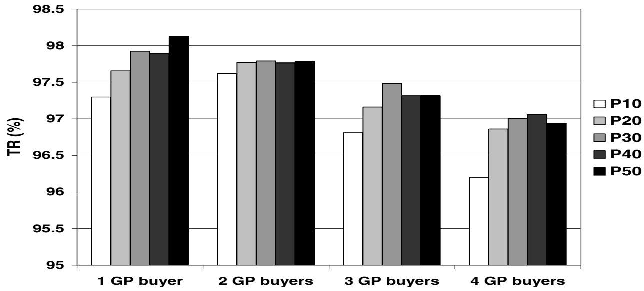

Figure 13-3. Total realized market surplus (TR) under different experimental setups.

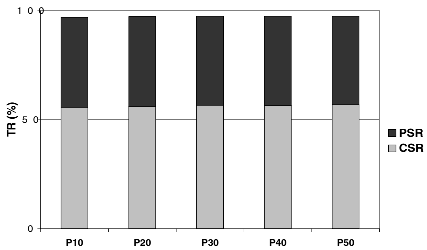

Figure 13-4. CSR vs. PSR under different GP population sizes.

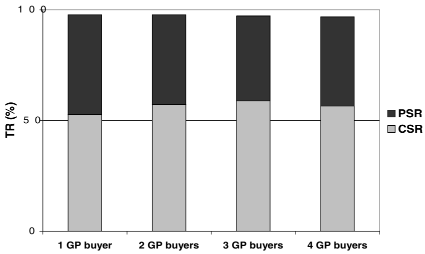

Figure 13-5. CSR vs. PSR under different numbers of GP buyers.

6. Macro Market Behavior Results

The market's total realized surplus ratios (TR) (averaged over 90 runs) under different experimental setups are presented in Figure 13-3. As shown, TR is high (more than 95%) regardless of the size of the GP population and the number of GP buyers in the market. One possible reason is that the market has naive truth-telling sellers, which make the market less competitive. Once the sellers are replaced with more sophisticated traders, the market's efficiency might decline. This will be investigated in our future work.

The results also show that, in general, TR *increases* as the GP population size increases, but *decreases* as the number of GP buyers in the market increases. In other words, all things equal, replacing a GP buyer with another more "intelligent" GP buyer increases the aggregated market efficiency. However, replacing a naive truth-telling buyer with a GP buyer decreases the market efficiency. To understand why this is happening, we calculated the market profit distribution between buyers (CSR) and sellers (PSR) under different GP population sizes (Figure 13-4) and different numbers of GP buyers (Figure 13-5).

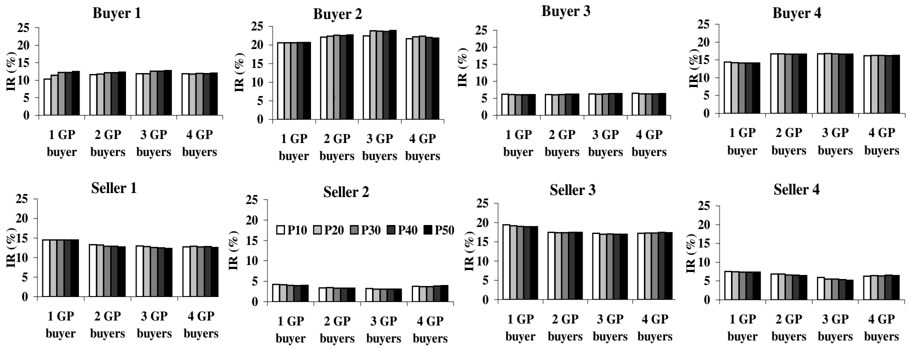

Figure 13-6. Individual realized market surplus (IR) under different experimental setups.

In both figures, buyers allocated more of the total market profit than the truth-telling sellers (CSR>PSR). This indicates that GP produces strategies that are more profitable than the naive truth-telling strategy. Meanwhile, when the number of GP buyers increases, CSR increases and PSR decreases. This indicates that the co-evolution of the GP buyer populations has generated more profitable strategies and strengthened the GP buyers' ability to reduce the sellers' profits. The gain in terms of CSR is, however, less than the loss of PSR. Consequently, the market's overall profits have declined.

By contrast, the naive truth-telling traders have no ability to learn. A "smarter" GP buyer (who has a larger population size) can therefore exploit the simple bidding pattern of these naive traders and increase its own profit. This improved CSR is not completely at the cost of PSR. Consequently, the total market efficiency is increased.

When considering the aggregated market behavior, both the increase in the GP population size and the number of GP buyers empower the buyers to exploit the naive truth-telling sellers. Increasing the GP population size impacts the market at the individual level, as it assigns a higher-degree of intelligence to a single buyer. By contrast, increasing the number of GP buyers impacts the entire buyer group, as all GP buyers co-evolve their trading strategies together. The latter approach therefore has a stronger impact on the market's profit allocation and produces a larger increase in CSP than the first approach.

Figure 13-6 shows the profits of each individual (IR) under different population sizes and numbers of GP buyers. On the buyers' side, either increasing the population size or the numbers of GP buyers has a positive impact on the profit. When a buyer is switched from truth-telling to GP, the buyer's profit increases. Meanwhile, other GP buyers also obtain increased profits. This again

Table 13-5. Distribution of daily profit generated by 18,000 strategies (pop_size 10).

Profit	-0.5	0	8	10.5	12.5	14	14.5	17	18	18.5	21	Total
Count	18	429	346	5,666	72	8	3,829	16	4	335	7,277	18,000

indicates that the co-evolution dynamics benefits all GP buyers. On the sellers' side, however, the effect is exactly the opposite. All truth-telling sellers face declining profits when "smarter" GP buyers or new GP buyers join the market; their profits are taken away by these newcomers.

Buyer 3 and seller 2 have a noticeably lower IR than the other traders. This is partly due to their tight budget constraints. As shown in Table 13-3, buyer 3 can not pay much for his 4 tokens and seller 2 faces high costs of his 4 tokens. This makes it hard for them to compete with other traders during an auction. Even when they win the auction, the profits from the transactions are low.

7. Individual GP Trading Strategies Analysis

We also analyze individual GP buyer strategies using the experimental results from two different settings. The first is when there is only one GP buyer in the market and the second is when there are two GP buyers in the market. In the first case, the focus is on how GP population size influences the learning of strategies. In the second case, besides population size, we also investigate the co-evolution dynamics of two GP buyers and its impact on the learning of strategies. We address these two cases in two sub-sections.

The data used to conduct our analysis are obtained from the runs with population sizes 10 and 50, which give rise to the largest difference in trader's "intelligence". We collected data on all evolved strategies and their daily profit (π) generated during the last 10 generations. We consider these strategies to be more "mature", and hence to better represent the GP buyers' trading patterns. With 90 runs for each setup, the data set is large enough for us to conduct the statistical analysis.

One GP Buyer Case

When the population size is 10, each generation is $2 \times 10 = 20$ days long. On each day, one strategy is picked randomly from the population to conduct the auction. The total number of auction days for the last 10 generations is therefore $20 \times 10 = 200$. Since we made 90 runs for this setup, the number of strategies picked to conduct the auction is $200 \times 90 = 18,000$. Table 13-5 gives the distribution of daily profit (π) generated by these strategies.

Table 13-6. The 3 most used strategies and the associated information (pop_size 10).

Strategy	Profit	Count	Ratio(Count/18,000)
`PMinBid`	21	7,277	0.4043
`PMinBid`	10.5	5,666	0.3148
`HTV`	14.5	3,829	0.2127
Total		16,722	0.9318

As shown in the table, the strategies that generated profits 21, 10.5 and 14.5 are used for a total of 93% of the auction. It is therefore reasonable to assume that they represent the GP buyer's trading strategies. Table 13-6 presents these 3 strategies and their associated information.

The most highly used strategy is `PMinBid`: the lowest bidding price on the previous day. Since the strategy that provided the lowest bidding on the previous day can be different, the generated bidding price can also be different. Consequently, this strategy can produce different profit. In our case, `PMinBid` generates a profit of 21 on 7,277 days and a profit of 10.5 on 5,666 days. After careful examination, we find that if the strategy used on the previous day to give the lowest bidding price is `HTV` (the highest token value) or `NTV` (the second highest token value), `PMinBid` will give a profit of 21 for the current day. If the strategy was `PMaxAsk` (the highest asking price among all transactions on the previous day), `PMinBid` gives a profit of 10.5 for the current day.

The other highly used strategy is `HTV`: the highest token value, which is also equal to truth-telling, as we always trade the token with the highest value first. This strategy gives a profit of 14.5, which is between the two possible profits produced by `PMinBid`.

Although using historical information (`PMinBid`) can be unstable (sometimes giving a profit of 21 and sometimes of 10.5), the buyer was found to be willing to take the risk and used it on 70% of the auction days, instead of being a safe truth-teller (`HTV`), who gained a stable and good profit (14.5). This risk taking did pay off: the number of days that `PMinBid` gave a profit of 21 was 20% higher than the number of days that gave a profit of 10.5. Overall, the strategy gave a higher profit than telling the truth.

The more intelligent buyer (GP with a population size of 50) has developed slightly different trading strategies under an identical setting. In addition to `PMinBid` and `HTV`, another class of strategies has also emerged. These strategies are more complex (longer in length) and give a slightly higher profit of 22 (we will call this group of strategies p-22). Table 13-7 lists the 4 most frequently used strategies by the more intelligent GP buyer together with the associated information. Note that the number of days for each generation is $2 \times 50 =$

Table 13-7. The 4 most used strategies and the associated information (pop_size 50).

Strategy	Profit	Count	Ratio(Count/90,000)
p-22	22	22,381	0.2487
PMinBid	10.5	10,976	0.1220
PMinBid	21	29,779	0.3309
HTV	14.5	7,827	0.0870
Total		70,963	0.7885

100. The total number of auction days (also the number of strategies picked to conduct auction) during the last 10 generations for all 90 runs is therefore $100 \times 10 \times 90 = 90,000$.

We examine this new class of strategies and find that they can be reduced to one of the following two forms:

- `Min PMinBid HTV`
- `If_Bigger_Then_Else PMinBid HTV HTV PMinBid`

These two strategies are identical semantically: combine `HTV` and `PMinBid` and choose the one which gives the lower price. The emergence of this class of strategies indicates that the "smarter" GP buyer is able to fine-tune existing strategies to produce more profitable strategies. This is probably due to its larger population size which allows more diversity in the strategies pool. Operators, such as `Min` and `If_Bigger_Then_Else`, have more opportunities to survive even if the strategies that contain them did not give rise to a very high profit. At a later time, GP is able to use them to fine-tune (through crossover and mutation) existing good strategies to become better ones.

The "smarter" GP buyer also used `PMinBid` more wisely: the percentage of its usage that generated a profit of 21 increased (40% for profit 21 and 31% for profit 10.5 when population size is 10 vs. 33% for profit 21 and 12% for profit 10.5 when population size is 50). This is probably due to this strategy utilizing historical auction information. With the emergence of the more profitable p-22 strategies, `PMinBid` is more likely to make better decisions.

The dependence between the `PMinBid` and p-22 strategies is not one-way. The existence of the p-22 strategies also relies on the presence of `PMinBid` and `HTV` (otherwise, the two less profitable strategies would have been replaced by p-22 strategies). This is not difficult to comprehend as p-22 strategies are composed of `PMinBid` and `HTV`. The co-existence of these 3 groups of strategies (p-22, `PMinBid` and `HTV`) is therefore mutually beneficial. The "intelligent" GP buyer has learned that information and has maintained the diversity (since it has a large enough population space) to make profitable auction decisions.

Table 13-8. Distribution of daily profit of the 2 most used strategies for two GP buyers.

	Pop_size 10				Pop_size 50			
Buyer	Profit	Count	Ratio	Total	Profit	Count	Ratio	Total
1	17	13,042	0.7246		17	53,704	0.5967	
	24	392	0.0218	0.7463	24	11,054	0.1228	0.7195
2	35	9,702	0.5390		35	38,358	0.4262	
	37	1,004	0.0558	0.5948	37	17,731	0.1970	0.6232

Two GP Buyers Case

With 2 GP buyers in the market, the learning dynamics becomes more complicated. Table 13-8 gives the profit distributions of the 2 strategies most used by the two GP buyers under population sizes of 10 and 50.

For GP buyer 1, the strategies most used are the ones that result in a profit of 17 (p-17) and for GP buyer 2, the strategies most used are those that result in a profit of 35 (p-35). This pattern is the same under population sizes 10 and 50. Meanwhile, when the population size was increased from 10 to 50, GP buyer 1 increased the usage of the strategies that realized a profit of 24 (p-24) and GP buyer 2 increased the usage of the strategies that produced a profit of 37 (p-37). These two patterns suggest the following correlation:

- GP buyer 1 used p-17 to auction against the p-35 of GP buyer 2
- GP buyer 1 used p-24 to auction against the p-37 of GP buyer 2

While examining the data, we found that many strategies produced the same profit. Meanwhile, GP buyer 1 used p-24 against many different strategies of GP buyer 2, not just p-37. Nevertheless, the p-17 vs. p-35 and p-24 vs. p-37 pairs appeared frequent enough in the auction rounds that we decided to conduct more of an investigation.

For the p-17 vs. p-35 pair, the most frequently used strategies were `HTV` of GP buyer 1 and `PMax` (the highest transaction price on the previous day) of GP buyer 2. Recall that GP buyer 1 was a risk-taker using `PMinBid` the most when no other player in the market had learning ability. However, in the presence of another buyer with GP learning ability, GP buyer 1 became a truth-teller. Why did GP buyer 1 make this switch? One possible reason is that `HTV` is a more profitable strategy than `PMinBid` against `PMax`.

We then examined the data and found that `PMinBid` vs. `PMax` would give GP buyer 1 a profit of 20.5 and GP buyer 2 a profit of 29.5. `PMinBid` is actually a better strategy than `HTV` for GP buyer 1 (but not for GP buyer 2). Why didn't GP buyer 1 learn the better strategy? We have two hypotheses:

- 300 generations did not provide enough time for GP to learn the `PMinBid` strategy under the presence of another GP buyer whose strategies were also evolving. In the other experiment where there was only one GP buyer in the market, the earliest time `PMinBid` appeared in the population was at generation 60. In some of the runs, the strategy was not seen until generation 230. Using a trial-and-error learning style, GP needed time to obtain information and improve. If the number of generations had been greater, GP buyer 1 might have learned the strategy. We will test the hypothesis in our future work.

- The market environment for GP buyer 1 to learn the `PMinBid` strategy was not provided. The simulation data showed that GP buyer 1 used `PMinBid` only when GP buyer 2 used `HTV` (i.e., truth-telling). When there is only one GP buyer in the market, all other buyers used `HTV` (i.e., truth-telling). Under such market environment, the GP buyer learned and used `PMinBid` the most to conduct auction. However, once the market has changed such that buyer 2 had the ability to learn and found `PMax` was a better strategy than truth-telling, `HTV` was seldom used. Without the required market environment, GP buyer 1 did not learn that `PMinBid` was a better strategy to use.

For the p-24 vs p-37 pair, 4 combinations appeared frequently:

`Min PMax HTV` vs. `PMax`

`Min PMinBid HTV` vs. `PMax`

`Min PMax HTV` vs. `PMinBid`

`Min PMinBid HTV` vs. `PMinBid`

Recall that when there was only one GP buyer in the market, `Min PMinBid HTV` was not a strategy used by GP buyer 1 when its population size was 10. It was only when the population size was increased to 50 that this more sophisticated strategy emerged. However, under the presence of another GP buyer, buyer 1 was able to learn this more profitable strategy under the population size of 10. This indicates that the competitive co-evolution dynamics between GP buyer 1 and GP buyer 2 has promoted the emergence of this strategy without an increase in population size.

When the population size was increased from 10 to 50, the use of p-24 and p-37 increased, while the use of p-17 and p-35 decreased. This indicates that the two more "intelligent" GP buyers learned to use a more profitable strategy to conduct the auction. This co-evolution dynamics produced a win-win result as both gained more by using the p-24 and p-37 strategies more often in the auctions. However, this was not the case for the p-17 vs. p-35 pair as GP buyer 2's gain was at the expense of GP buyer 1. As mentioned previously, the co-evolution dynamics led GP buyer 2 to use `PMax` more often as it resulted in a higher profit than did truth-telling. Without the presence of GP buyer 2's `HTV`, GP buyer 1 did not learn the more profitable `PMinBid` and used `HTV` which resulted in a lower profit.

To sum up, with the presence of another GP buyer in the market, GP buyer 1 changed its strategy in conducting auctions in the following ways:

1. By switching from `PMinBid` to the less profitable truth-telling strategy.
2. Capable of learning the more sophisticated and profitable `Min PMinBid HTV` strategies.

The aggregated result of the 2 changes was positive and the profit of GP buyer 1 increased. Meanwhile, by switching from truth-telling to being a GP learner, buyer 2 also gained more profit from the co-evolution of the two GP buyers. When the intelligence of both GP buyers was increased, they learned to use the more profitable strategies more often and both gained more profits.

8. Concluding Remarks

The behavior of a DA market is the result of the aggregated behavior of multiple traders. Due to the interdependence of the traders' actions, it is not easy to isolate or analyze any individual trader's auction strategies based on the market's overall performance. While the human subject experiment technology has been developed to analyze emerged market phenomena, the representation of human cognition is an unsettled issue. ABM provides an alternative to conduct market simulation in a controlled environment. The results of the ABM simulation provide a basis that can be validated by human subject experiments.

This research conducts DA market simulations on an ABM platform. By using GP to implement the bounded rationality of traders' decision making, we were able to study the macro market performance changes under traders with different degrees of rationality. Meanwhile, the individual GP buyer's strategies and the co-evolution of multiple GP buyers' strategies were analyzed.

In terms of the macro market performance, we observed that the DA market efficiency increases when the individual trader's intelligence increases. However, increasing the number of intelligent traders in the market led to a decline in market efficiency. This is because the more intelligent GP trader developed strategies to collect extra profits that did not conflict with others' profits. On the contrary, under multiple GP traders with learning ability, the market became more competitive. As a result, the GP traders developed aggressive strategies that damaged other traders' profits. The overall market efficiency was reduced.

The learning behavior of the individual GP trader differs when the market environment is different. When the market is stable (i.e., all other traders are truth-tellers), the GP trader learns more simple rules that require less risk to gain more profit than that gained by being a truth-teller. When another GP trader joined the market, however, it was found that the market became dynamic in that both GP traders revised their strategies constantly to outdo each other. The co-evolution of the two GP traders' strategies generated both positive and negative impacts on the two traders. In one instance, both GP traders co-operated and applied strategies that benefited each other. In another instance, one GP trader applied strategies that prevented the other GP trader from learning more profitable strategies to protect his own profit. The third instance was where one GP trader learned more sophisticated strategies in this dynamic environment, which he did not learn under the stable environment.

All of the observed market performance and individual traders learning behavior make intuitive sense. Under the ABM platform, GP agents demonstrated human-like rationality in decision making. We plan to conduct human subject experiments to validate these results in the near future.

Acknowledgment

A preliminary draft of this paper was presented at the GPTP-2008 – Genetic Programming Theory and Practice (GPTP) Workshop, Ann Arbor, Michigan May 15-17, 2008. The authors are grateful for the comments received from the participants at the conference, especially, John Holland and Robert Reynolds. The paper is revised in light of two very helpful referee reports provided by B. Worzel, M. Kim, Y. Becker, P. Fei and U-M O'Reilly to whom the authors are grateful. The research support in the form of NSC grant No. NSC. 95-2415-H-004-002-MY3 is also gratefully acknowledged.

References

Arthur, W. Brian (1993). On designing economic agents that behave like human agents. *Journal of Evolutionary Economics*, 3:1–22.

Chen, Shu-Heng and Tai, Chung-Ching (2003). Trading restrictions, price dynamics and allocative efficiency in double auction markets: an analysis based on agent-based modeling and simulations. *Advances in Complex Systems*, 6(3):283 – 302.

Cliff, Dave and Bruten, Janet (1997). Zero is not enough: On the lower limit of agent intelligence for continuous double auction markets. Technical Report HP-97-141, HP Technical Report.

Duffy, John (2006). Agent-based models and human subject experiments. In Tesfatsion, Leigh and Judd, Kenneth L., editors, *Handbook of Computational Economics*, chapter 19, pages 949 – 1011. Elsevier.

Gode, Dhananjay K. and Sunder, Shyam (1993). Allocative efficiency of markets with zero-intelligence traders: markets as a partial substitute for individual rationality. *Journal of Political Economy*, 101:119 – 137.

Koza, John R. (1992). Genetic evolution and co-evolution of game strategies. In *International Conference on Game Theory and Its Applications*.

Rust, John, Miller, John, and Palmer, Richard (1993). Behavior of trading automata in a computerized double auction market. In Friedmand, D. and Rust, John, editors, *The Double Auction Market: Institutions, Theories and Evidence*, pages 155 – 198. Addison-Wesley.

Simon, Herbert A. (1997). Behavioral economics and bounded rationality. In Simon, Herbert A., editor, *Models of Bounded Rationality*, pages 267 – 298. MIT Press.