Co-evolution of Auction Mechanisms and Trading Strategies: Towards a Novel Approach to Microeconomic Design

Steve Phelps[1] **Simon Parsons**[1,2] **Peter McBurney**[1] **Elizabeth Sklar**[3]

[1]Dept. of Computer Science, University of Liverpool, Liverpool, UK
sphelps,p.j.mcburney@csc.liv.ac.uk

[2]Center for Coordination Science, Sloan School, M.I.T., Cambridge, MA, USA
sparsons@mit.edu

[3]Dept. of Computer Science, Columbia University, New York, NY, USA
sklar@cs.columbia.edu

Abstract

Mechanism design is the economic theory of the design of effective resource allocation mechanisms, such as auctions. Traditionally, economists have approached design problems by studying the analytic properties of different mechanisms. An alternative is to view a mechanism as the outcome of some evolutionary process involving *buyers*, *sellers* and the *auctioneer*. As a first step in this alternative direction, we have applied genetic programming to the development of an auction pricing rule for double auctions in a wholesale electricity marketplace. For this purpose we adopted the multi-agent simulation model of Nicolaisen, Petrov and Tesfatsion.

1 Introduction

Much recent work in the field of Multi-Agent Systems (MAS) has focused on resource allocation problems, for example (Fatima & Wooldridge 2001; Jennings *et al.* 2001). These problems are especially difficult to solve efficiently in an open system if the values which agents place on resources, or the values of their human principals, are private and unobservable. In such a situation, the difficulty facing somebody wishing to give the resources to those who value them most highly is that participating agents cannot necessarily be relied upon to report their private values truthfully; there is nothing to prevent "greedy" agents from exaggerating their resource requirements.

Auction mechanisms attempt to overcome this difficulty by having agents support their value-claims with hard cash. Such mechanisms can be designed so as to induce agents to reveal their true valuations, thereby enabling the allocation of resources to those agents who genuinely value them most highly. Designing mechanisms to achieve specific economic requirements, such as achieving market efficiency or maximising social welfare, against self-interested intelligent traders, is no trivial matter as can be seen from accounts of the auction design process for the recent radio spectrum auctions in the UK (Klemperer 2002). The economic theory of mechanism design approaches the task of designing efficient resource allocation mechanisms by studying the formal, analytical properties of alternative mechanisms (Jackson 2000; Sandholm 1999). Because of the complexities involved in market design problems, economists are increasingly turning to computational methods in an attempt to take an engineering approach to "microeconomic design" (Roth 2001).

Our approach applies the notion of *co-evolutionary machine learning* (Hillis 1992; Angeline & Pollack 1993; Miller & Cliff 1994) to the microeconomic design problem. In successful applications of co-evolution, populations of agents interact with each other, an "arms race" spiral develops wherein each population spurs the other(s) to advance and the result is continuous learning for all populations. However, this has been notoriously difficult to achieve. Often populations settle into a *mediocre stable state*, reaching a local optima and being unable to move beyond it.

Consequently, there is a growing body of work examining the dynamics of co-evolutionary learning environments in an attempt to identify phenomena that contribute to success (Cliff & Miller 1995; Pollack & Blair 1998; Ficici & Pollack 1998; Blair, Sklar, & Funes 1999). The following aspects are of particular importance:

1. choice of representation for individuals within each population;
2. definition of a fitness function for determining which individuals in a population will reproduce;
3. methodology and proportion for reproduction;

4. selection of learning experiences for individuals (i.e., who interacts with whom, how many times and how frequently);

5. avoidance of collusion[1] wherein members of different populations can work together to make non-optimal moves that may produce better short-term results for each but cause the populations as a whole to get stuck in local optima; and

6. a clearly defined vision of the landscape and how to measure progress so that one can even recognize if a local (or indeed global) optimum has been reached.

We see efficient mechanisms evolving through repeated interactions between participants who may also be evolving individually — thus we believe that the co-evolutionary methodology is highly appropriate for our problem. In our work, we are using genetic programming (GP) (Koza 1992) to represent individuals with different roles in an auction: the auctioneer, and the two types of traders (buyers and sellers). Through the interactions of the traders, individual and group trading strategies evolve, as well as auction mechanisms themselves. We view the mechanisms as "hosts" and the trading strategies as "parasites"; as greedy, non-truthful strategies emerge, it would be hoped that the mechanism population will adapt defenses, and that strategy-proof, incentive-compatible mechanisms would evolve.

Such an approach is the long-term aim of our research, and to our knowledge we are the first to apply genetic programming and co-evolution to mechanism design. Here, we report our initial work towards this aim. To provide a multi-agent test-bed for such an approach we have adopted the wholesale electricity market auction simulation model of (Nicolaisen, Petrov, & Tesfatsion 2001), hereafter referred to as NPT. In Section 2, we describe the NPT model, and our work to replicate their results. Section 3 then describes our use of genetic programming to co-evolve trading strategies for buyers and sellers in these auctions. Section 4 presents some of our preliminary results in using genetic programming to evolve auction pricing rules. The final section concludes with a brief description of our future research.

2 The NPT model

In the NPT experiments (Nicolaisen, Petrov, & Tesfatsion 2001), a number of traders buy and sell electricity in a discriminatory-price continuous double auction. Every trader has a private value for the electricity that they trade; for buyers this is the price that they can obtain in a secondary retail market and for sellers this reflects the costs associated with generating the electricity. Trade in electricity is affected by capacity constraints; every trader has a finite maximum capacity of electricity that they can generate or purchase for resale. The market proceeds in rounds of trading. In each round, all the traders are given the opportunity to bid in a random order. Each trader selects a price and submits a bid or an ask at that price and with a quantity equal to their generating capacity. Trade proceeds until the maximum number of auction rounds is reached.

Agents use a myopic reinforcement learning algorithm which is a modification of the Roth-Erev algorithm (Roth & Erev 1995); the learner chooses possible actions from K possible mark-ups, and receives a reward directly proportional to the profits that result from this offer. The learner chooses actions by generating random numbers according to a probability distribution built up linearly from the cumulative rewards for each possible action. The modified Roth-Erev algorithm (MRE) has three main parameters: r the recency parameter; e the experimentation parameter and $s(1)$ the scaling parameter.

NPT is interested in the market power that can be exercised by buyers or sellers under different market conditions. Market power is defined as the difference between actual profits earned versus the theoretical profits available in competitive-equilibrium, expressed as a ratio of the equilibrium profits. The different market conditions are represented by two parameters: *relative concentration* (RCON) and *relative capacity* (RCAP). RCON is the ratio of the number of buyers (NB) to the number of sellers (NS) and RCAP is the relative generating capacity of each group.

2.1 NPT results

The main results from NPT are summarised in Table 1. Each cell of the table corresponds to particular values for RCON and RCAP. The outcome under these conditions is summarised by the variables:

- *Buyer MP* – market power exercised by buyers
- *Seller MP* – market power exercised by sellers
- *Efficiency* – ratio of total profits earned to total profits theoretically available in competitive equilibrium, expressed as a percentage

Because traders use stochastic strategies, the sensitivity of these outcomes to particular values of the pseudo-random number generator (PRNG) seed is

[1] Note that this is not necessarily the same as the notion of collusion in auction theory.

tested by running the experiment 100 times with different PRNG seeds on each iteration. For each variable we present the average result, followed by the standard deviation in parentheses.

These results are significant because they indicate that there are market biases inherent in the discriminatory-price auction rules. For example, one would expect that Seller MP should increase as RCAP increases, but this is not what is found by experimentation. NPT concludes that the inherent market-structure is responsible for failure of this hypothesis.

2.2 Replication of results

This scenario was selected for our research because of the focus on market power. As agents evolve successful extra-marginal strategies, their market power indices will increase. For example, if sellers were able to employ collusive price-fixing strategies, we should expect to see their market power indices grow. Different auction rules may have differing abilities to counter this kind of tactic; hence, market power outcomes are an important quantative metric to consider in assessing auction designs.

We began our implementation of the NPT model by attempting to replicate the results presented in that paper. The software used to run the auction experiments was written in Java. The software is available under an open-source license at `http://jasa.sourceforge.net/`. The 4-heap algorithm (Wurman, Walsh, & Wellman 1998) was used to maintain auction state; all price information was encoded using double-precision floating point variables and all quantity information was encoded using integers. The Mersenne Twister PRNG was used to generate the random numbers required for MRE.

The replicated results are presented in Table 2. Although similar market power and mean efficiency outcomes are obtained, the standard deviations we obtained for the efficiency outcomes are considerably larger than those reported in NPT. These results give us some confidence that our experimental setup is accurate, although we are continuing to try and track down the source of these increased standard deviations.

3 Co-evolution of Trading Strategies using Genetic-Programming

We next compare the reinforcement learning algorithm used by NPT with co-evolution of trading strategies using genetic programming. In this work, we evolve a separate population of strategies for each trader in the electricity market scenario. These strategies evolve in competition with the simultaneously evolving strategies of other traders. For these experiments we made use of a Java-based evolutionary computation system called ECJ.[2] The scenario is similar to the NPT experiments, but instead of using the modified Roth-Erev algorithm to select prices, agents select prices by evaluating a function that was evolved using genetic programming (GP).

ECJ implements a strongly-typed GP (Montana 1993) version of Koza's (Koza 1992) original system. For all of the GP experiments in this paper, the standard Koza parameters were used in combination with the standard Koza GP operators. The functions given in Tables 3 and 4 were used as the GP function-set, and the initial populations are generated randomly using these functions. As is usually the case with GP, individuals are tree structures composed of these functions. We used six populations of GP-evolved strategies, that is one population for each buyer and seller in the market. Each population contained 100 tree-individuals. When breeding trees for the next generation, the crossover operator is applied with a probability of 0.9, and the reproduction operator is applied with a probability of 0.1, as per standard Koza GP (Koza 1992). Individuals are selected for breeding using tournament selection, with a tournament size of 7.

To evaluate the fitness of individuals in each generation, one member of each population was randomly selected. The strategies that corresponded to these trees were then played against each other in a 10-round version of the electricity market, and each individual's fitness was set in proportion to the profits obtained for the corresponding strategy. This continued until all individuals in all populations had been evaluated, giving a fitness measure for each individual. Note that wherever evaluation of the tree resulted in a negative price, or in a division by zero exception, the price was set to 0 and this was used as the requisite bid or ask. These fitness values, established by competition between populations are then used, as described above, to select which individuals from a single population will get to reproduce (both in terms of being copied to the next generation and undergoing crossover).

Initially, we are interested in whether high-efficiency outcomes are sustained in this experiment. As with the NPT experiments, high levels of market efficiency indicate that overall, traders are successfully "discovering" profits that are available in the market. We would not necessarily expect to see stability, or gradual improvement, of each strategy's individual profits in this co-evolutionary scenario. But if overall mar-

[2] `http://www.cs.umd.edu/projects/plus/ec/ecj/`

Relative Concentration	Relative Capacity 1/2		stdev	1.00		stdev	2.00		stdev
2	Buyer MP	-0.13	(0.09)	Buyer MP	-0.15	(0.09)	Buyer MP	0.10	(0.30)
	Seller MP	0.55	(0.38)	Seller MP	0.38	(0.33)	Seller MP	-0.10	(0.25)
	Efficiency	99.81	(0.02)	Efficiency	96.30	(0.05)	Efficiency	99.88	(0.06)
1	Buyer MP	-0.22	(0.12)	Buyer MP	-0.13	(0.10)	Buyer MP	0.13	(0.33)
	Seller MP	0.80	(0.53)	Seller MP	0.28	(0.35)	Seller MP	-0.10	(0.26)
	Efficiency	92.13	(0.09)	Efficiency	94.59	(0.07)	Efficiency	100.00	(0.00)
1/2	Buyer MP	-0.21	(0.12)	Buyer MP	-0.14	(0.08)	Buyer MP	0.09	(0.24)
	Seller MP	0.67	(0.46)	Seller MP	0.30	(0.31)	Seller MP	-0.07	(0.19)
	Efficiency	91.84	(0.09)	Efficiency	94.24	(0.07)	Efficiency	100.00	(0.00)

Table 1: Original NPT market power and efficiency outcomes for the best-fit MRE algorithm with 1000 auction rounds and parameter values s(1) = 9.00, r= 0.10, and e = 0.20. Refer to the original NPT paper for a detailed description of the MRE parameters: r the recency parameter; e the experimentation parameter and $s(1)$ the scaling parameter.

Relative Concentration	Relative Capacity 1/2		stdev	1.00		stdev	2.00		stdev
2	Buyer MP	-0.33	(0.07)	Buyer MP	-0.27	(0.08)	Buyer MP	0.10	(0.11)
	Seller MP	1.12	(0.31)	Seller MP	0.72	(0.32)	Seller MP	-0.15	(0.10)
	Efficiency	94.46	(3.87)	Efficiency	95.04	(3.43)	Efficiency	96.71	(0.51)
1	Buyer MP	-0.39	(0.07)	Buyer MP	-0.28	(0.08)	Buyer MP	0.10	(0.08)
	Seller MP	1.19	(0.40)	Seller MP	0.76	(0.30)	Seller MP	-0.15	(0.07)
	Efficiency	91.01	(7.61)	Efficiency	95.34	(3.26)	Efficiency	96.63	(0.47)
1/2	Buyer MP	-0.38	(0.09)	Buyer MP	-0.27	(0.08)	Buyer MP	0.04	(0.07)
	Seller MP	0.84	(0.45)	Seller MP	0.72	(0.29)	Seller MP	-0.10	(0.06)
	Efficiency	84.86	(9.93)	Efficiency	94.62	(3.92)	Efficiency	96.79	(0.42)

Table 2: Replicated market power and efficiency outcomes for the best-fit MRE algorithm with 1000 auction rounds and parameter values s(1) = 9.00, r = 0.10 and e = 0.20

Function	*Arguments*	*Return-type*	*Description*
+	(+ `number number`)	`number`	Addition
−	(− `number number`)	`number`	Subtraction
/	(/ `number number`)	`number`	Division
∗	(∗ `number number`)	`number`	Multiplication
1	none	`number`	1
DoubleERC	none	`number`	A double precision floating point ephemeral random constant in the range (0..1).
QuoteBidPrice	none	`number`	The current bid quote
QuoteAskPrice	none	`number`	The current ask quote

Table 3: GP functions common to all function-sets

Function	*Arguments*	*Return-type*	*Description*
<	(< `number` `number`)	`boolean`	Less-than function
=	(= `number` `number`)	`boolean`	Equals function
>	(> `number` `number`)	`boolean`	Greater-than function
True	none	`boolean`	True
PrivateValue	none	`number`	The agent's private valuation for electricity
Nand	(Nand `boolean` `boolean`)	`boolean`	Not-and boolean operator
IfElse	(IfElse `boolean` `number` `number`)	`number`	Return 2nd argument if condition is true, otherwise return 3rd argument.

Table 4: Additional GP functions used in evolving trading strategies

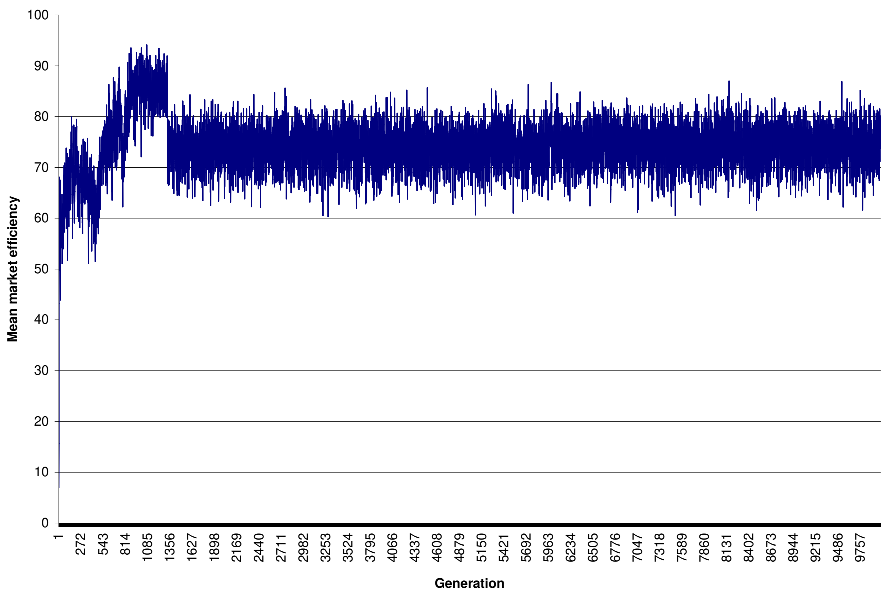

Figure 1: Evolution of mean efficiency for RCON=1 and RCAP=1 over 10,000 generations using a fixed discriminatory-pricing auctioneer, and 6 sub-populations of co-evolving strategies each of size 100.

ket efficiency does decline temporarily, we would expect the co-evolving strategy set as a whole to adapt and reaquire the "lost" profits; thus if strategy sub-populations are successfully adapating to new market conditions, we would expect to see market efficiency remain stable at near to 100%.

Figure 1 shows the evolution of the mean market efficiency for each generation of the experiment in the case RCAP=1 and RCON=1 over 10,000 generations. Note that by generation 2000, the market efficiency has become stable, and the mean efficiency is 74.3.

The use of co-evolution to evolve trading strategies is not new in experimental economics; for example, see (Price 1997). Our interest in co-evolving strategies was to verify that such an approach worked for this scenario, and also as a step towards the use of co-evolutionary techniques to evolve trading strategies and auction rules—in other words to evolve mechanisms along with the best way to trade within them. To our knowledge no one has yet done this, and our preliminary work towards doing this will be the focus of the next section.

4 Co-evolution of Auction Pricing Rules and Strategies

An additional population of *auctioneers* was introduced into our experiment. These agents were derived from the auctioneer classes that we implemented for our previous experiments, but instead of using the standard code to set the clearing price for a given transaction, they used a function that was evolved using GP. The set of functions used for the auction

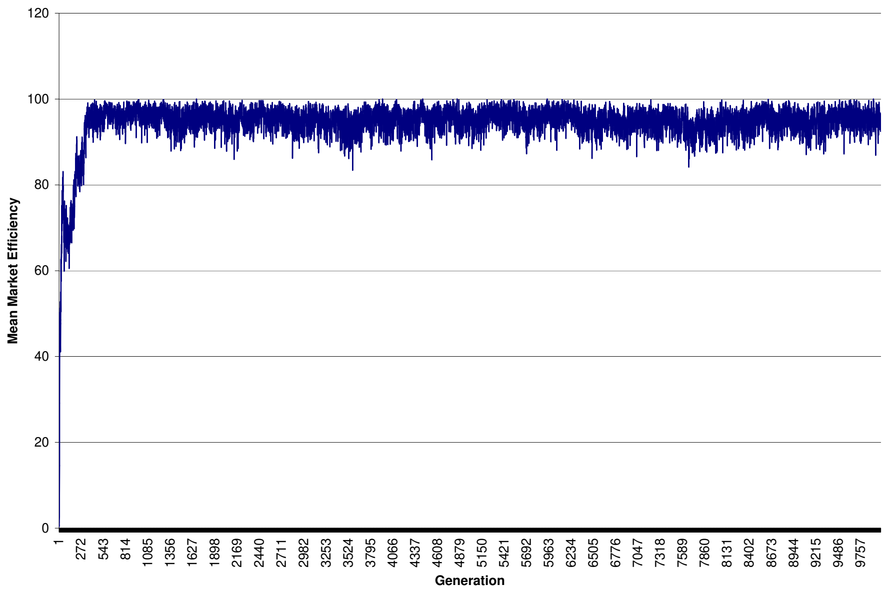

Figure 2: Evolution of mean efficiency for RCON=1 and RCAP=1 over 10,000 generations using an auctioneer with a GP-evolved pricing rule, and 6 additional populations of co-evolving strategies.

pricing rule consists of those functions in Tables 3 and 5. The fitness for the auctioneer population was set proportional to the total profits earned in the market.

Intuitively, the auctioneer population can be thought to be "learning" auction-pricing rules that maintain market efficiency in the face of co-evolving strategies. Our hypothesis is that in this version of the experiment, in which there are a small number of traders with fixed private values, the most robust auction pricing rule is the one that sets the price for electricity at the equilibrium price, regardless of what traders actually bid. We believe that the auctioneer population should discover this rule; i.e it should discover the equilibrium price for the market. This is because private values are fixed, and the auctioneer population has indirect access to meta-information about the market — market efficiency — that is based on the (in-practice unobservable) private values. Of course, this pricing rule would not work in practice, because in practice private values are not from a fixed, predefined set. However, this hypothesis gives us a basis to test some of our assumptions about this experiment. Future work will consider scenarios in which agents with randomized private-values enter and leave the market.

The experimental set-up was a slight variation of the previous experiment. We added a seventh population, auctioneers, and evaluated their fitness by running auctions with randomly selected buyer and seller individuals (again picking one random individual from each of the six populations) and looking at the overall profits obtained. The same auctions were used to evaluate the buyers and sellers, though their fitness was still based on local profit.

Figure 2 shows the evolution of the mean market efficiency for each generation of this version of the experiment in the case RCAP=1 and RCON=1 over 10,000 generations. As can be seen from the graph, the adaptive auctioneers are able to improve mean market efficiency when compared to the fixed discriminatory-price auctioneer used in the previous section — the mean efficiency for the adaptive auctioneer is 94.5. In addition, the market reaches stability more quickly, after only 500 generations.

Figure 3 shows the function tree evolved for the auctioneers' pricing rule in the final generation, and Table 6 shows the trading strategy-set for that auction. We are currently investigating whether our hypothesis regarding the discovery of the equilibirum price is borne out by this experiment.

5 Conclusions and Further Work

In this paper we have reported on the preliminary stages of work aiming to explore the evolution of economic auction mechanisms. In our initial work, we have adopted a multi-agent systems test-bed involving auctions in an electricity marketplace, first developed in NPT. In that work, the trading agents in the auctions were equipped with a modified Roth-Erev learning model, enabling them to change their bids on the

(+ (+ -0.8958983 QuoteAskPrice) (/ (/ QuoteAskPrice (+ (+ (+ (+ (+ -0.8958983 QuoteAskPrice)
(+ QuoteAskPrice (+ -0.8958983 QuoteAskPrice))) (+ (+ (+ -0.8958983 QuoteAskPrice) (+ QuoteAskPrice
(+ (+ -0.8958983 (+ QuoteAskPrice -0.8958983))
-0.8958983))) (+ (+ (+ QuoteAskPrice (+ -0.8958983 QuoteAskPrice)) (+ (+ QuoteAskPrice (+ QuoteAskPrice
QuoteAskPrice)) (+ -0.8958983 (+ QuoteAskPrice (+ -0.8958983 (+ -0.8958983 QuoteAskPrice))))))
(+ (+ (+ -0.8958983 (+ -0.8958983 QuoteAskPrice))
-0.8958983) (/ (+ QuoteAskPrice -0.8958983) (+ (+ (+ (+ QuoteAskPrice -0.8958983) -0.8958983)
-0.8958983) (+ -0.8958983 QuoteAskPrice)))))))
(+ (+ (+ (+ QuoteAskPrice -0.8958983) (+ (+ QuoteAskPrice -0.8958983) (+ QuoteAskPrice
-0.8958983))) (+ QuoteAskPrice (+ QuoteAskPrice -0.8958983))) (+ (+ (+ -0.8958983 QuoteAskPrice)
(+ -0.8958983 (+ -0.8958983 QuoteAskPrice))) (+ (+ (+ (+ (+ -0.8958983 -0.8958983) (+
(+ QuoteAskPrice -0.8958983) (+ -0.8958983 -0.8958983))) (+ -0.8958983 -0.8958983))
QuoteAskPrice) (+ (+ QuoteAskPrice -0.8958983) (+ QuoteAskPrice QuoteAskPrice)))))) (+ -0.8958983
(+ (+ (+ (+ (+ QuoteAskPrice -0.8958983) QuoteAskPrice) (+ QuoteAskPrice QuoteAskPrice))
-0.8958983) (+ (+ QuoteAskPrice QuoteAskPrice) (+ (+ QuoteAskPrice (+ -0.8958983 QuoteAskPrice))
-0.8958983)))))) (+ (+ (+ QuoteAskPrice -0.8958983)
(+ (+ (+ -0.8958983 QuoteAskPrice) (+ (+ QuoteAskPrice (+ (+ (+ -0.8958983 -0.8958983)
-0.8958983) -0.8958983)) (+ -0.8958983 (+ -0.8958983 QuoteAskPrice)))) (+ (+ (/ -0.8958983
(+ QuoteAskPrice -0.8958983)) (+ QuoteAskPrice (+ QuoteAskPrice -0.8958983))) (+ QuoteAskPrice
(+ QuoteAskPrice -0.8958983))))) (+ -0.8958983 QuoteAskPrice))))

Figure 3: Best auction pricing rule in final generation of a 10,000 generation experiment for population size 100, $RCON = 1$, $NS = 3$ and $NB = 3$

Function	*Arguments*	*Return-type*	*Description*
AskPrice	none	`number`	The price of the ask (offer to sell) currently being matched in the auction
BidPrice	none	`number`	The price of the bid currently being matched in the auction

Table 5: Additional GP functions used in evolving auctioneer pricing rules

basis of their experiences in successive auction rounds.

We first sought to replicate the results of NPT, and were able to obtain similar average results; our standard deviations, however, were much larger than those reported in NPT. We then compared the use of the MRE learning algorithm with an experiment in which strategies are co-evolved using genetic programming. The genetic programming approach was able to produce reasonably high efficiency outcomes. Finally we presented some of our preliminary work on evolving *auction designs* using genetic-programming which again was able to produce relatively high efficiency outcomes. We believe that this is the first attempt to evolve auction mechanisms.

Future work in this area will focus on: (i) an analysis of the different auction rules evolved for each of the combinations of RCAP and RCON that were originally discussed in NPT and (ii) incorporation of market-power metrics into the fitness function for auction rules.

A key question concerning this work is how to track the *adaptive progress*, as opposed to the instantaneous fitness, of the auctioneers verses the trading strategies. We are currently investigating the possibility of using CIAO (Current Individual vs. Ancestral Opponents) metrics as proposed in (Cliff & Miller 1995), in order to gain insights into the co-evolutionary dynamics of these experiments. We are also thinking about the possibility of using pareto co-evolution (Watson & Pollack 2000) in order to ensure that auction designs are robust in the face of a diverse range of strategies.

Our research work is part of a larger, European-wide effort, the Sustainable Lifecycles in Information Ecosystems Project[3] exploring the use of biological paradigms in the study of multi-agent systems. In particular, recent work by our project partners (Sierra *et al.* 2002) has shown how generic MAS systems may be designed by evolutionary processes. In this context, our work focuses specifically on the design of electronic institutions for multi-agent trading systems.

Acknowledgements

This work was partially supported by the EU IST Programme through the SLIE Project (IST-1999-10948). We also thank Michael Phelps, Leigh Tesfatsion and Peter Wurman for their assistance.

[3] `http://www.dai.ed.ac.uk/groups/ssp/slie/`

Seller 1	QuoteAskPrice
Seller 2	QuoteBidPrice
Seller 3	QuoteBidPrice
Buyer 1	QuoteAskPrice
Buyer 2	QuoteAskPrice
Buyer 3	QuoteBidPrice

Table 6: The set of trading strategies taking part in the final auction corresponding to Figure 3

References

Angeline, P. J., and Pollack, J. B. 1993. Competitive environments evolve better solutions for complex tasks. In Forrest, S., ed., *Genetic Algorithms: Proceedings of the Fifth International Conference (GA93).*

Blair, A. D.; Sklar, E.; and Funes, P. 1999. Coevolution, determinism and robustness. In *Proceedings of the Second Asia-Pacific Conference on Simulated Evolution and Learning*, 389–396. Springer-Verlag.

Cliff, D., and Miller, G. F. 1995. Tracking the red queen: Measurements of adaptive progress in coevolutionary simulations. In *European Conference on Artificial Life*, 200–218.

Fatima, S. S., and Wooldridge, M. 2001. Adaptive task and resource allocation in multi-agent systems. In *Proceedings of the 5th International Conference on Autonomous Agents*, 537–544. ACM Press.

Ficici, S. G., and Pollack, J. B. 1998. Challenges in coevolutionary learning: Arms-race dynamics, openendedness, and mediocre stable states. In *Proceedings of of ALIFE-6.*

Hillis, W. D. 1992. Co-evolving parasites improve simulated evolution as an optimization procedure. In *Proceedings of ALIFE-2*, 313–324. Addison Wesley.

Jackson, M. O. 2000. Mechanism theory. In *The Encyclopedia of Life Support Systems.* EOLSS Publishers.

Jennings, N. R.; Faratin, P.; Lomuscio, A. R.; Parsons, S.; Wooldridge, M.; and Sierra, C. 2001. Automated negotiation: prospects, methods and challenges. *Group Decision and Negotiation* 10(2):199–215.

Klemperer, P. 2002. How (not) to run auctions: the European 3G telecom auctions. *European Economic Review (forthcoming).*

Koza, J. R. 1992. *Genetic Programming: On the Programming of Computers by means of Natural Selection.* MIT Press.

Miller, G. F., and Cliff, D. 1994. Protean behavior in dynamic games: Arguments for the co-evolution of pursuit-evasion tactics. In *Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, 411–420.

Montana, D. J. 1993. Strongly typed genetic programming. Technical Report #7866, 10 Moulton Street, Cambridge, MA 02138, USA.

Nicolaisen, J.; Petrov, V.; and Tesfatsion, L. 2001. Market power and efficiency in a computational electricity market with discriminatory double-auction pricing. *IEEE Transactions on Evolutionary Computation* 5(5):504–523.

Pollack, J. B., and Blair, A. D. 1998. Co-evolution in the successful learning of backgammon strategy. *Machine Learning* 32:225–240.

Price, T. C. 1997. Using co-evolutionary programming to simulate strategic behaviour in markets. *Journal of Evolutionary Economics* 7:219–254.

Roth, A. E., and Erev, I. 1995. Learning in extensive form games: experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior* 8:164–212.

Roth, A. E. 2001. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica (forthcoming).*

Sandholm, T. W. 1999. Distributed rational decision making. In Weiss, G., ed., *Multiagent Systems: A Modern Introduction to Distributed Artificial Intelligence.* Cambridge, MA, USA: MIT Press. 201–258.

Sierra, C.; Sabater, J.; Agusti, J.; and Garcia, P. 2002. Evolutionary computation in MAS design. In *Proceedings of the 15th European Conference on Artificial Intelligence.*

Watson, R. A., and Pollack, J. B. 2000. Symbiotic combination as an alternative to sexual recombination in genetic algorithms. In *6th International Conference on Parallel Problem Solving from Nature.* Springer Verlag.

Wurman, P. R.; Walsh, W. E.; and Wellman, M. 1998. Flexible double auctions for electronic commerce: theory and implementation. *Decision Support Systems* 24:17–27.