Co-evolutionary auction mechanism design: a preliminary report

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

[1] Department of Computer Science, University of Liverpool
Chadwick Building, Liverpool L69 7ZF, UK.
{sphelps,p.j.mcburney}@csc.liv.ac.uk

[2] Center for Coordination Science, Sloan School of Management,
M.I.T., 3 Cambridge Center NE20-336, Cambridge, MA 02142, USA
sparsons@mit.edu

[3] Department of Computer Science, Columbia University
1214 Amsterdam Avenue, Mailcode 0401, New York, NY 10027, USA
sklar@cs.columbia.edu

Abstract. Auctions can be thought of as resource allocation mechanisms. The economic theory behind such systems is *mechanism design*. 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 an *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.

1 Introduction

Much recent work in the field of Multi-Agent Systems (MAS) has focused on resource allocation problems, for example [8, 14]. 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 encouraging 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 [15] and the US [7, 17]. The economic theory of mechanism design approaches the task of designing efficient resource allocation mechanisms by studying the formal, analytical properties of alternative mechanisms [13, 26]. However, for some kinds of mechanisms, including continuous double auctions [10], the mechanisms are too complex to admit analytical solutions. Because of these complexities, economists are increasingly turning to computational methods in an attempt to take an engineering approach to "microeconomic design" [23, 25]. We follow such an approach in this paper.

2 Co-evolution

One approach to computational microeconomic design is to use techniques from machine learning to explore the space of possible ways in which agents might act in particular markets. For example, reinforcement learning has been used to explore bidding patterns in auctions [20, 23] and establish the ways in which price-setting behaviour can affect consumer markets [27]. Another approach is to use techniques from evolutionary computing, that is from genetic algorithms [12] and genetic programming [16].

Inspired by the biological metaphor of evolution, genetic algorithms code aspects of a solution to a problem in an artificial "chromosome" (typically a binary string) and then breed a population of chromosomes using techniques like crossover (combining different bits of the strings from different individuals) and mutation (flipping individual bits). Genetic programming extends this approach by evolving not a bit-string-encoded solution to a problem, but an actual program to solve the problem itself. Programs are encoded as s-expressions and modelled as trees (nodes are function names and branches arguments of those functions); and these trees are subject to crossover (swapping subtrees from different programs) and mutation (replacing subtrees with random subtrees). Whichever approach is used, the best individuals, evaluated using a fitness function, are kept and "bred"; and bad individuals are rejected. However, deciding which individuals are the best is a hard problem.

Evolutionary approaches perform a search through a space of solutions with the theoretical advantage that random jumps around the search space — created by crossover and mutation — can prevent the system from getting stuck in local optima, unlike other techniques like hill climbing. Unfortunately, in practice this is not always the case at least partly because what constitutes the best fitness measure can change over time. To overcome this problem, researchers turned to *co-evolution* [1, 11, 18], and the aim of our work is to apply co-evolution to economic mechanism design.

In successful applications of co-evolution, simultaneously evolving populations of agents interact with each other, each providing a fitness measure for the other which changes as both populations evolve. When this works, 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 [2, 6, 9, 21]. The following aspects are of particular importance (some of which are relevant for both evolutionary and co-evolutionary techniques):

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. operators and proportion of population(s) used for reproduction;
4. selection of learning experiences for individuals (i.e., who interacts with whom, how many times and how frequently);
5. size of population and number of populations;
6. 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
7. a clearly defined vision of the fitness 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. Thus it is our long term aim to understand the above aspects for the evolution of trading strategies and auction rules.

In our work, we are using genetic programming (GP) [16] 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 auctioneer population will adapt defenses, and that strategy-proof, incentive-compatible mechanisms would evolve. Investigation of 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 (though [4] describes similar work—this is discussed in more detail in Section 5).

Here, we report our initial work towards this aim. In Section 3, we describe the scenario we are studying. Section 4 then describes our use of genetic programming to co-evolve trading strategies for buyers and sellers in these auctions, and some of our preliminary results in using genetic programming to evolve auction pricing rules. Section 5 discusses how these results fit into our overall plan

[1] Note that this is not necessarily the same as the notion of collusion in auction theory. Collusion in co-evolution is where members of the co-evolving populations help each other to score high fitness, the by-product being that the populations as a whole settle into a local optimum. Collusion in auction theory is where several bidders work together to buy goods for less than would have been paid were they not working together.

of work, and describes some future lines of work. Finally, Section 6 concludes with a brief summary.

3 The Experimental Scenario

To provide a multi-agent test-bed for such an approach we have adopted the wholesale electricity market auction simulation model of [20]. In this scenario, 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.

In [20] agents use a myopic reinforcement learning algorithm which is a modification of the Roth-Erev algorithm [24]; 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.

The scenario is investigated in terms of 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.

The main results from [20] 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 seed is tested by running the experiment 100 times with different seeds on each iteration. For each

	Relative Capacity								
	1/2			1.00			2.00		
			stdev			stdev			stdev
	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)
2									
	Efficiency	99.81	(0.02)	Efficiency	96.30	(0.05)	Efficiency	99.88	(0.06)
Relative	Buyer MP	-0.22	(0.12)	Buyer MP	-0.13	(0.10)	Buyer MP	0.13	(0.33)
Concentration	Seller MP	0.80	(0.53)	Seller MP	0.28	(0.35)	Seller MP	-0.10	(0.26)
1									
	Efficiency	92.13	(0.09)	Efficiency	94.59	(0.07)	Efficiency	100.00	(0.00)
	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)
1/2									
	Efficiency	91.84	(0.09)	Efficiency	94.24	(0.07)	Efficiency	100.00	(0.00)

Table 1. 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 [20] for a detailed description of the MRE parameters: r the recency parameter; e the experimentation parameter and $s(1)$ the scaling parameter.

variable we reproduce 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. [20] concludes that the inherent market-structure is responsible for failure of this hypothesis.

4 Co-evolution using Genetic Programming

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.

4.1 Co-evolution of Trading Strategies

In our initial work, we evolved 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. The scenario is basically that described above, the only difference in our approach being that instead of using the modified Roth-Erev algorithm to select prices, buyers and sellers select prices by evaluating a function that was evolved using genetic programming.

The heart of our simulation was a Java implementation of the 4-heap algorithm [29] which 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. This software is available under an open-source

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 2. GP functions common to all function-sets

license at `http://jasa.sourceforge.net/`. For the genetic programming part of the experiments, we made use of a Java-based evolutionary computation system called ECJ.[2] ECJ implements a strongly-typed GP [19] version of Koza's [16] 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 2 and 3 were used as the GP function-set, and the initial populations were 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. The fitness function for each population corresponded to the profits earned by the corresponding trader. The use of a separate co-evolving population for each trader allowed us to explore the potential emergence of collusive tactics between self-interested traders; each population attempts to maximise its own profits, but in certain situations populations may be able to increase their profits by co-operating with rival populations. This could not be modeled, by, for example, representing all of the buyers as a single population, since this optimization problem would not account for self-interested behaviour of individual traders.

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 [16]. 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 this set of 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 at the end of the rounds. This process 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 excep-

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

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 3. Additional GP functions used in evolving trading strategies

tion, the price was set to 0 and this was used as the requisite bid or ask[3]. These fitness values, established by competition between populations, are then used, as described above, to select which individuals from a single population will be permitted to reproduce (both in terms of being copied to the next generation and undergoing crossover).

Initially, we were interested in whether high-efficiency outcomes are sustained in this experiment. As with the original 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. However, if overall market efficiency were to 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 were to successfully adapt to new market conditions, we would expect to see mean market efficiency remain stable at around 100% since mean market efficiency measures the performance of the different varieties of buyer and seller as a whole. 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 relatively 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 [22]. Our interest in co-evolving strategies was mainly to verify that such an approach worked for this scenario, which it appears to. The work described in this section was also 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. This is the main focus of our research, and our preliminary work towards doing this will be the subject of the next section.

[3] In double auction terminiology, buyers make "bids" and sellers make "asks". Bids are offers to buy at a given price, asks are offers to sell at a given price.

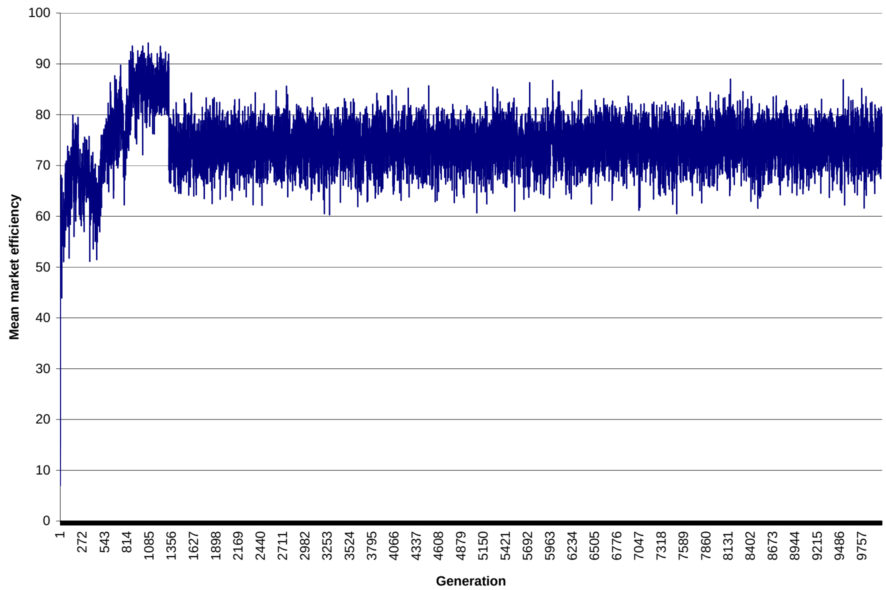

Fig. 1. Evolution of mean efficiency for RCON=1 and RCAP=1 over 10,000 generations using a fixed discriminatory-pricing auctioneer, and 6 sub-populations of co-evolving strategies each of size 100.

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 4. Additional GP functions used in evolving auctioneer pricing rules

4.2 Co-evolution of Auction Pricing Rules

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 pricing rule are those functions in Tables 2 and 4. The space of possible pricing rules thus encompasses, but is not restricted to, both uniform-price and discriminatory-price versions of the k-double auction; pricing rules making use of the `AskPrice` or `BidPrice` functions[4]

[4] `AskPrice` and `BidPrice` are the current prices that the auctioneer is considering as possible matches.

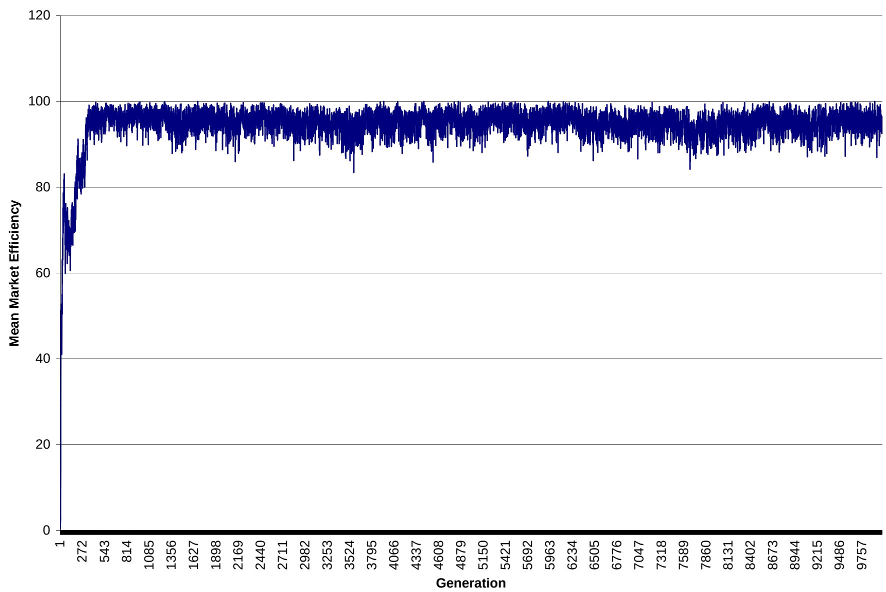

Fig. 2. Evolution of mean efficiency for RCON=1 and RCAP=1 over 10,000 generations using an auctioneer with a GP-evolved pricing rule, and 6 additional populations of co-evolving strategies.

correspond to discriminatory-price auctions, whereas pricing rules not making use of these functions correspond to uniform-price auctions. Whereas the trading-strategy populations' fitness was proportional to their individual profits, 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 buying and selling 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 — it should discover the equilibrium price for the market. It should do this 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, by considering the hypothesis that the most robust pricing rule is the one that sets prices at the equilibrium level, we will be able to assess the validity of our un-

		RCAP $\frac{1}{2}$	1	2
	2	$QuoteBidPrice - 0.39$	$QuoteBidPrice$	$QuoteBidPrice$
RCON	1	$\approx QuoteAskPrice$	$\approx QuoteAskPrice$	$\approx QuoteAskPrice$
	$\frac{1}{2}$	$35.47 - 35.47AskPrice$	$BidPrice$	$QuoteBidPrice$

Table 5. GP-evolved auction pricing rules at generation 1000 for different market conditions, i.e. different values of $RCON$ and $RCAP$

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 at generation 1000 for $RCON = 1$, $RCAP = 1$

derlying assumption. 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%, as compared to 74.3% for the case where the auctioneer does not evolve. In addition, the market reaches stability more quickly, after only 500 generations.

Table 5 shows the stable pricing function evolved for the auctioneers' pricing rule under different market conditions. In all cases the pricing rule is a linear function of the Bid and Ask Prices and the function only uses either Bid Price or Ask Price[5]. When the number of buyers and sellers is equal, the pricing rule

[5] Some of the pricing rules also use `QuoteBidPrice` or `QuoteAskPrice` which are the values made public by the auctioneer to give buyers and sellers an idea of what the

is only determined by the Ask Price, suggesting that the sellers control the market whatever the relative capacity. Table 6 shows the trading strategy-set for the auction after 1000 generations in the case $RCON = RCAP = 1$. These expressions are simplifications of the s-expressions generated by the GP. Most reduce exactly to the expressions given, but several seem to resist simplification— these were plotted against `QuoteAskPrice` and were found to be approximately equal to it. They are thus given as $\approx QuoteAskPrice$.

5 Discussion

In terms of the seven aspects of any application of co-evolution that were raised in Section 2, we believe that the work described here has provided an adequate start to dealing with the first four — how to represent individuals, what counts as fitness, how to carry out reproduction, and how to perform selection. We can claim this because the choices we have explained above lead to the evolution of reasonable individuals as evidenced by the high level of market efficiency obtained in our experiments. However, there is still much work to be done. Even the results obtained so far have raised some interesting questions, such as how to interpret the different auction rules that can be evolved for each of the combinations of RCAP and RCON, and how to incorporate market-power metrics into the fitness function for auction rules. Clearly we also have to address the final three issues mentioned in Section 2 as well — correct population size, how to detect and avoid collusion, and how to measure progress.

This latter is a particularly important question since we need to be able 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 [6], in order to gain insights into the co-evolutionary dynamics of these experiments, and using pareto co-evolution [28] in order to ensure that auction designs are robust in the face of a diverse range of strategies.

Finally we should discuss the relation of our work with that of Cliff [4], which is the only other work that we are aware of in which the auction mechanism itself evolves. Cliff's work in this area builds on his Zero-Intelligence-Plus [3] traders, and first used genetic algorithms to determine the parameters that control the bidding behaviour of the agents [5]. This work is analagous to our use of genetic programming to decide how buyers and sellers bid. The next stage of the work, which was undertaken concurrently with, but independently of, ours was to add an extra parameter into the genetic algorithm representing the probability with which a buyer or seller is selected to make a Bid or Ask[6]. This, then, has the same kind of aim as our work, but uses a genetic algorithm rather than genetic programming, and, as a result, is only concerned with tuning

current trading price is. It is the equivalent of the prices displayed on stock exchange tickerboards.

[6] The experiment thus explores a continuum between auctions in which only buyers act, like an English auction, and auctions in which only sellers act.

one, admittedly important, parameter rather than constructing the auction rules from scratch. Furthermore, since Cliff's work involves just a single population of chromosomes—which capture the parameters which determine buyers, sellers and auctioneers—it is an evolutionary but not a co-evolutionary approach.

6 Summary

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. We first described work in which buyer and seller strategies are co-evolved using genetic programming. The genetic programming approach was able to produce reasonably high efficiency outcomes in this case. Next we presented some of our preliminary work on evolving *auction designs* using genetic-programming which again was able to produce relatively high efficiency outcomes and was able to reach stability quicker than when the buyer and seller strategies evolved alone. We believe that this is the first attempt to evolve auction mechanisms, and, though far from complete, makes it possible to frame further research in this area.

Acknowledgements This work was partially supported by the EU IST Programme through the SLIE Project (IST-1999-10948). We also thank Larry Bull, Sam Hough, Michael Phelps, Alvin Roth, Leigh Tesfatsion and Peter Wurman for their assistance, and the anonymous referees for their helpful comments.

References

1. P. J. Angeline and J. B. Pollack. Competitive environments evolve better solutions for complex tasks. In S. Forrest, editor, *Genetic Algorithms: Proceedings of the Fifth International Conference (GA93)*, 1993.
2. A. D. Blair, E. Sklar, and P. Funes. Co-evolution, determinism and robustness. In *Proceedings of the Second Asia-Pacific Conference on Simulated Evolution and Learning*, pages 389–396. Springer-Verlag, 1999.
3. D. Cliff. Minimal-intelligence agents for bargaining behaviours in market environments. Technical Report HPL-97-91, HP Labs, 1997.
4. D. Cliff. Evolution of market mechanism through a continuous space of auction-types. Technical Report HPL-2001-326, HP Labs, 2001.
5. D. Cliff. Evolutionary optimization of parameter sets for adaptive software-agent traders in continuous double auction markets. Technical Report HPL-2001-99, HP Labs, 2001.
6. Dave Cliff and Geoffrey F. Miller. Tracking the red queen: Measurements of adaptive progress in co-evolutionary simulations. In *European Conference on Artificial Life*, pages 200–218, 1995.
7. P. Cramton and J. Schwartz. Collusive bidding: Lessons from the FCC spectrum auctions. *Journal of Regulatory Economics*, 17, 2000.

8. S. S. Fatima and M. Wooldridge. Adaptive task and resource allocation in multi-agent systems. In *Proceedings of the 5th International Conference on Autonomous Agents*, pages 537–544. ACM Press, 2001.
9. S. G. Ficici and J. B. Pollack. Challenges in coevolutionary learning: Arms-race dynamics, open-endedness, and mediocre stable states. In *Proceedings of ALIFE-6*, 1998.
10. D. Friedman and J. Rust, editors. *The Double Auction Market: Institutions, Theories and Evidence*. Santa Fe Institute Studies in the Sciences of Complexity. Perseus Publishing, Cambridge, MA, USA, 1993.
11. W. D. Hillis. Co-evolving parasites improve simulated evolution as an optimization procedure. In *Proceedings of ALIFE-2*, pages 313–324. Addison Wesley, 1992.
12. J. Holland. *Adaptation in Natural and Artificial Systems.* University of Michigan Press, 1975.
13. M. O. Jackson. Mechanism theory. In *The Encyclopedia of Life Support Systems.* EOLSS Publishers, 2000.
14. N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, M. Wooldridge, and C. Sierra. Automated negotiation: prospects, methods and challenges. *Group Decision and Negotiation*, 10(2):199–215, 2001.
15. P. Klemperer. How (not) to run auctions: the European 3G telecom auctions. *European Economic Review (forthcoming)*, 2002.
16. J. R. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection.* MIT Press, 1992.
17. R. P. McAfee and J. McMillan. Analyzing the airwaves auction. *Journal of Economic Perspectives*, 10:159–176, 1996.
18. G. F. Miller and D. Cliff. 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*, pages 411–420, 1994.
19. David J. Montana. Strongly typed genetic programming. Technical Report #7866, Bolt Beranek and Newman, 10 Moulton Street, Cambridge, MA 02138, USA, July 1993.
20. J. Nicolaisen, V. Petrov, and L. Tesfatsion. Market power and efficiency in a computational electricity market with discriminatory double-auction pricing. *IEEE Transactions on Evolutionary Computation*, 5(5):504–523, October 2001.
21. J. B. Pollack and A. D. Blair. Co-evolution in the successful learning of backgammon strategy. *Machine Learning*, 32:225–240, 1998.
22. T. C. Price. Using co-evolutionary programming to simulate strategic behaviour in markets. *Journal of Evolutionary Economics*, 7:219–254, 1997.
23. A. E. Roth. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica (forthcoming)*, 2001.
24. A. E. Roth and I. Erev. Learning in extensive form games: experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior*, 8:164–212, 1995.
25. J. Rust, J. Miller, and R. Palmer. Behavior of trading automata in a computerized double auction market. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, pages 155–198. Perseus Publishing, Cambridge, MA, USA, 1993.
26. T. W. Sandholm. Distributed rational decision making. In G. Weiss, editor, *Multiagent Systems: A Modern Introduction to Distributed Artificial Intelligence*, pages 201–258. MIT Press, Cambridge, MA, USA, 1999.

27. G. Tesauro and J. O. Kephart. Pricing in agent economies using multi-agent Q-learning. *Autonomous Agents and Multi-Agent Systems*, 5(3):289–304, 2002.
28. R. A. Watson and J. B. Pollack. Symbiotic combination as an alternative to sexual recombination in genetic algorithms. In *6th International Conference on Parallel Problem Solving from Nature*. Springer Verlag, 16-20 2000.
29. P. R. Wurman, W. E. Walsh, and M.P. Wellman. Flexible double auctions for electronic commerce: theory and implementation. *Decision Support Systems*, 24:17–27, 1998.