Auctions, evolution, and multi-agent learning

Steve Phelps[1], Kai Cai[2], Peter McBurney[1], Jinzhong Niu[2], Simon Parsons[1,2,3], and Elizabeth Sklar[2,3]

[1] Department of Computer Science, University of Liverpool,
Ashton Building, Ashton Street, Liverpool L69, 3BX
phelps.sg@googlemail.com, mcburney@liverpool.ac.uk
[2] Department of Computer Science, Graduate Center, City University of New York,
365 5th Avenue, New York, NY 10016, USA
{kcai, jniu}@gc.cuny.edu
[3] Department of Computer and Information Science, Brooklyn College,
City University of New York, 2900 Bedford Avenue, Brooklyn, NY 11210, USA
{parsons, sklar}@sci.brooklyn.cuny.edu

Abstract. For a number of years we have been working towards the goal of automatically creating auction mechanisms, using a range of techniques from evolutionary and multi-agent learning. This paper gives an overview of this work. The paper presents results from several experiments that we have carried out, and tries to place these in the context of the overall task that we are engaged in.

1 Introduction

The allocation of resources between a set of agents is a challenging problem, and one that has been much studied in artificial intelligence. Resource allocation problems are especially difficult to solve efficiently in an open system if the values that 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 allocate the resources to those who value them most highly is that participating agents cannot necessarily be relied upon to report those values truthfully — there is nothing to prevent "greedy" agents from exaggerating their resource requirements.

To overcome this problem, it has been suggested that resource allocation be solved using market mechanisms [4, 50] in which agents support their value-claims with hard cash. This has two advantages. First it punishes greedy agents by making them pay for the resources that they have oversubscribed to. (Alternatively one can think of this as preventing agents from oversubscribing by forcing them to pay a higher price than they would have to pay for the resources they actually need.) Second, it allocates resources to the agents who pay the most, which should be the agents who value the resources most highly. *Auctions* are a subclass of market mechanisms that have received particular attention. This is due to the fact that auctions, when well designed, can achieve desired economic outcomes like high *allocative efficiency*.

Designing mechanisms to achieve specific economic requirements, such high efficiency or maximal 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 Europe [22] and the US [10, 27]. The economic theory of mechanism design [19] approaches the task of designing efficient resource allocation mechanisms by studying the formal, analytical properties of alternative mechanisms. Mechanism design views auctions as form of game, and applies traditional analytic methods from game theory to some kinds of auctions [25], for example the second-price sealed-bid auctions or Vickrey auctions [48].

The high complexity of the dynamics of some other auction types, especially *double-sided auctions* [13] or DAs, however makes it difficult to go further in this direction [24, 43, 49]. As a result, researchers turned to experimental approaches. Smith pioneered the experimental approach [45] conducting auctions involving human traders that revealed many of the properties of double auctions. For example, his work showed that in *continuous double auctions* or CDAs, even a handful of traders can lead to high overall efficiency, and transaction prices can quickly converge to the theoretical equilibrium. More recently has come the suggestion that economists should take an "engineering approach" [38, 40] to problems in microeconomics in general, building computational models of auctions, testing them experimentally, and refining them to create robust markets. We see our work as being part of this engineering approach to market design.

One approach to the computational design of markets 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 [30, 38] and establish the ways in which price-setting behaviour can affect consumer markets [47]. Our work is in this line. However, we differ from much of the existing work on machine learning in computational market design by using machine learning to design the auction rules themselves, rather than just in the service of exploring their behaviour. We refer to this line of work as *automated mechanism design*, and the idea behind this paper is to summarise the work that we have been doing over the past few years on automated mechanism design. It does not provide any new results, but instead sketches the relationship between the series of experiments that we have carried out, describes the results that we have obtained, and tries to explain how all we have done fits into the overall scope of our work.

We should stress that we are not trying to evolve entire auction mechanisms. The computational complexity of doing so places this out of our reach at the moment. Instead we concentrate on parts of an existing mechanism, the continuous double auction, and look to automatically tune them for specific situations. Our work is experimental, and so comes with no formal guarantees. It thus stands in stark contrast to the work of Contizer and Sandholm [8, 9], which looks to create entire mechanisms subject to absolute guarantees on their performance. However, our work addresses much more complex mechanisms, and we see it as addressing the same problem from a different direction.

2 Background

2.1 Auctions, briefly

To frame our work, we borrow from Friedman's [13] attempt to standardise terminology in which *exchange* is the free reallocation of goods and money between a set of traders.

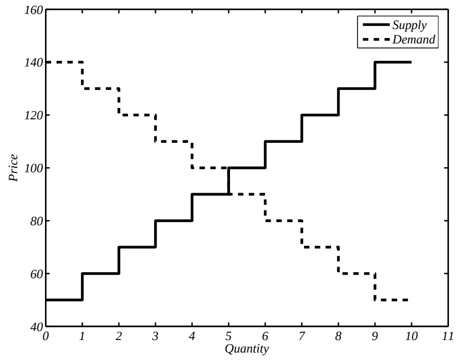

Fig. 1. Supply and demand curves

A market institution lays down the rules under which this exchange takes place, and an *auction* is a specific kind of market institution. A given institution defines what information traders can exchange, and how the reallocation of goods and money will occur, a process known as *clearing* the market. In an auction, the only information that traders can exchange are offers to buy at a given price, called *bids*, and offers to sell at a given price, called *ask*s, and an auction gives priority to higher bids and lower asks. An auction can allow only buyers or only sellers to make offers, in which case it is *one-sided*, or it can allow both, in which case it is *two-sided*. A double auction is a two-sided auction, and from here on we will only deal with double auctions. Here the aim of the mechanism is to pair buyers and sellers, matching pairs such that the buyer is prepared to pay a higher price than the seller wants. We are most interested in two kinds of double auction. The *clearing house* (CH) auction matches traders by collecting offers over a period and, at the end of that period, identifying the matching pairs. The *continuous double auction* (CDA), in contrast, constantly looks for matches, identifying one as soon as it has some bid that is greater than some ask. Once matches have been found, a *transaction price* is set, somewhere in the interval between bid and ask.

In common with most work on double auctions, we only deal with the auction of a single kind of good, and we assume that every trader has a *private value* for the good — the price that the good is really worth to the agent. A rational buyer will not bid above its private value, and a rational seller will not ask below that value. If we know the private values of a set of traders, we can construct supply and demand curves for the market they are involved with, as in Figure 2.1. Here the heavy line, the supply curve, indicates that one seller has a private value of 50 — below that value no goods will be sold, and once the price rises to 50 exactly one good will be sold. The second trader has a private value of 60 and at a price of 60, exactly two goods will be sold. Similarly, there is one buyer who is willing to pay 140, and at that price one good will be bought, but as soon as the price falls to 130, two goods will be bought.

The intersection of the supply and demand curve indicates the point at which supply and demand are in balance — here any price between 90 and 100 will see exactly five goods bought and sold. Economic theory predicts that this *equilibrium* situation is what will hold if 20 traders with the indicated private values get together to trade. However the theory offers no clues as to how the traders will figure out which of them should trade, an at what price, and it is clear that it is not in the traders' interest to make offers that are truthful and indicate their private value — a trade which shades their offer, stating a lower price than their private value if they are a buyer, will make a profit if that offer is accepted.

If we know the private values of the traders, then, as described above, we can compute the equilibrium. Combining information about the equilibrium with information about what actually happens in the market, we can compute metrics that summarise the performance of the market. The *actual overall profit*, pr^a, of an auction is the sum of the actual profits of buyers and sellers:

$$pr^a = pr^a_b + pr^a_s$$

and these are computed as:

$$pr^a_b = \sum_i v_i - p_i$$
$$pr^a_s = \sum_j p_j - v_j$$

where p_i is the price of a trade made by *buyer* i and v_i is the private value of buyer i for all buyers who trade and p_j is the price of a trade made by *seller* j and v_j is the private value of buyer j for all sellers who trade.

The *theoretical* or *equilibrium profit*, pr_e, is:

$$pr^e = pr^e_b + pr^e_s \tag{1}$$

the sum of the *equilibrium profits* of buyers and sellers, the profit that they would make if all trades took place at the *equilibrium price* p_0, the price predicted by theory. These can be computed as:

$$pr^e_b = \sum_i v_i - p_0$$
$$pr^e_s = \sum_j p_0 - v_j$$

The *allocative efficiency* of an auction is then:

$$e_a = \frac{pr^a}{pr^e} \tag{2}$$

which, of course, is the same as:

$$e_a = \frac{pr^a_b + pr^a_s}{pr^e_b + pr^e_s}$$

The allocative efficiency measures how close the market is to the equilibrium that theory predicts in terms of profit for traders. All other things being equal, economists prefer markets with high efficiency snce this indicates that the market is transfering goods to the buyers that value them most from the sellers that value them least. This mmaximises *social welfare*, making the traders as happy as possible. Allocative efficiency is maximal if just the traders to the left of the intersection between supply and demand curves in Figure 2.1 end up trading. While measuring allocative efficiency is useful, it says nothing about price. An auction that trades at equilibrium will be efficient, but high efficiency does not indicate that the market is trading near the equilibrium price [15]. The convergence coefficient, α, was introduced by Smith [45] to measure how far an active auction is away from the equilibrium point. It measures the rms deviation of transaction prices from the equilibrium price:

$$\alpha = \frac{100}{p_0}\sqrt{\frac{1}{n}\sum_{i=1}^{n}(p_i - p_0)^2} \tag{3}$$

These are the measures that we will make most use of in this paper.

Our experimental work follows the usual pattern for work on automated trading agents. We run each auction for a number of trading *days*, with each day being broken up into a series of *rounds*. A round is an opportunity for agents to make offers, and we distinguish different days because at the end of a day, agents have their inventories replenished. As a result, every buyer can buy goods every day, and every seller can sell every day. Days are not identical because agents are aware of what happened the previous day. Thus it is possible for traders to learn, over the course of several days, the optimal way to trade. Following [30], we use a k*-double-auction* transaction pricing rule [43], in which the transaction price for each matched bid-ask pair is set according to the following function:

$$p_t = kp_a + (1-k)p_b \tag{4}$$

where p_t is the transaction price, p_a is the ask price, p_b is the bid price and k is a parameter that can be adjusted by the auction designer. This is a *discriminatory* pricing rule since the price may be different for each transaction. In constrast, a *uniform* pricing rule ensures all transactions take place at the same price. In [30] and in much of our work, k is taken to be 0.5. To run most of the simulations described here we used JASA [20][4], which supports a wide range of auction types and trading strategies, and which matches bids and asks using the 4-heap algorithm [51].

2.2 Related work

Much of the computational work on analysing markets has been concerned with algorithms that can be used to decide what price to trade at. From the economics side, this work has often been motivated by the lack of an adequate theory of price formation — a theory that says how individuals decide what offers to make (though as Smith [45] demonstrated, this doesn't stop individuals being good at making these decisions)

[4] More accurately, JASA was developed as a result of the need to write software to run the simulations. The initial version of JASA was designed and written by Steve Phelps, and more recently has been contributed to by Jinzhong Niu and Kai Cai.

— and the desire to understand what makes markets work. From the computer science side, the motivation has usually been to find algorithms that can trade profitably, finding a price that the market will support while ensuring that this is at a profit.

Gode and Sunder [15, 16] were among the first to address this question, claiming that no intelligence is necessary for the goal of achieving high efficiency — so the outcome is due to the auction mechanism itself. They introduced two trading strategies: *zero intelligence without constraint* (ZI-U) and *zero intelligence with constraint* (ZI-C), and showed that ZI-U, the more naïve version, which shouts an offer at a random price without considering whether it is losing money or not, performs poorly. In contrast, ZI-C, which lacks the motivation of maximizing profit just like ZI-U but guarantees no loss, generates high efficiency solutions [15]. These results were however questioned by Cliff and Bruten [4, 6], who thought Gode and Sunder's conclusion was not convincing because the scenarios considered were not as comprehensive as in Smith's experiments, and showed that in different scenarios the ZI-C agents performed poorly, especially in terms of convergence to the theoretical equilibrium.

Cliff and Bruten further [4, 5] designed a simple adaptive trading strategy called *zero intelligence plus* or ZIP, and showed ZIP worked better than ZI-C, generating high efficiency outcomes and converging to the equilibrium price. This led Cliff and Bruten to suggest that ZIP embodied the minimum intelligence required by traders. Subsequent work has led to the development of many further trading strategies, the best known of which include Roth and Erev's [11, 39] reinforcement learning strategy, which we call RE, Gjerstad and Dickhaut's [14] approach, commonly referred to as GD, which uses the past history of accepted bids and asks to compute the expected value of every offer a trader might make, and the simplification of ZIP introduced by Preist and van Tol [37].

This work on trading strategies is only one facet of the research on auctions. Gode and Sunder's results suggest that auction mechanisms play an important role in determining the outcome of an auction, and this is further bourne out by the work of Walsh *et al.* [49] (which also points out that results hinge on both auction design and the mix of trading strategies used). For example, if an auction is *strategy-proof*, traders need not bother to conceal their private values and in such auctions complex trading agents are not required. While typical double auctions are not strategy-proof, McAfee [26] has derived a form of double auction that is strategy-proof (though this strategy-proofness comes at the cost of lower efficiency).

3 Evolving the whole system

Our initial approach to automated mechanism design was to use techniques from evolutionary computing. Inspired by the biological metaphor of evolution, genetic algorithms (GAs) [18] 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 bits of the strings from different individuals) and mutation (flipping individual bits). Genetic programming (GP) [23] 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 sub-

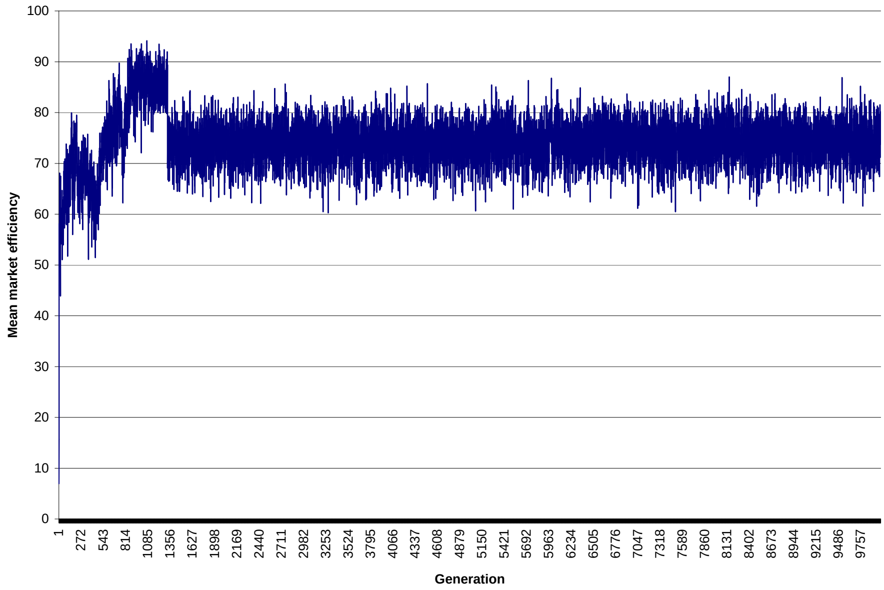

Fig. 2. Evolving traders: efficiency by generation

ject 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.

Both genetic algorithms and genetic programming perform a search through the space of possible 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 machine learning techniques. 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, some researchers have turned to *co-evolution*, for example [1, 17, 29].

In co-evolution, simultaneously evolving populations of agents interact, providing each other with a fitness measure that changes as the agents evolve. In successful applications, an "arms race" spiral develops wherein each population spurs the other 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 optimum 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, 7, 12]. In the context of auction design, it is possible to look at a number of different forms of co-evolution. First, different traders co-evolve against one another, with different offer-

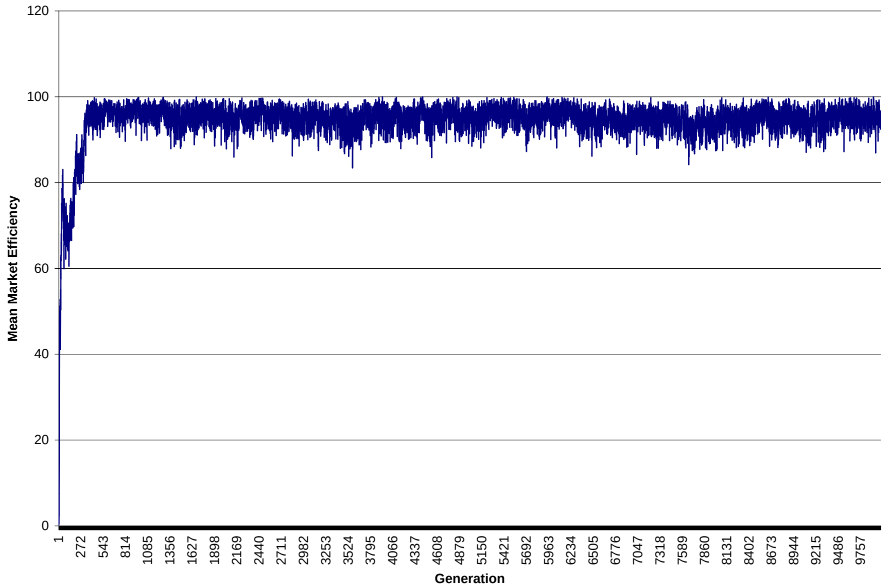

Fig. 3. Evolving traders and auctioneer: efficiency by generation

making strategies being the co-evolving populations, each attempting to gain an advantage over the others. Since all traders are looking to maximise their profits, they are to some extent in competition, although it is possible for a number of successful traders to coexist. Second, traders co-evolve against the auction mechanism itself — the auctioneer if you will — as the co-evolving populations. The traders' aim is to achieve high profits while the auction(eer)'s aim is to provide an efficient market. While these aims ned not be mutually exclusive, they may also be in conflict.

In [34], we explored a simple approach to co-evolving mechanisms, attempting to evolve the rules by which traders decided how to make offers, and the rules by which the auctioneer decides to set trade prices based upon those offers. When evolving rules for the traders alone, setting prices using a standard rule, we obtained the results in Figure 3. When we simultaneously evolved rules for traders and rules for the auctioneer, we obtained the results in Figure 3. While the efficiency of the whole system is not particularly high when we only evolve traders[5], when we evolve both traders and auctioneer, we obtain quite respectable efficiencies of around 95%.

There is a problem with these results, however. The problem is that it appears that the systems that we managed to evolve were systems that had fallen into the trap of a mediocre stable state. If we look at the kinds of rule that the traders were learning to use in these experiments, they are of the form:

---

[5] An average efficiency of around 75% compares poorly with the results reported in the literature for automated trading strategies, and with our own work [3].

```
if(not(QuoteBidPrice < (PrivateValue * 0.081675285))
 {
    PrivateValue
 }
else
 {
    PrivateValue * 0.081675285
 }
```

where `QuoteBidPrice` is the highest unmatched bid (this is a rule for a buyer). In other words, the traders were learning to make a constant markup, but nothing more sophisticated than that. While such a strategy can be quite successful when competing against traders doing the same — as discussed by [52] — we know that it does not compete well with more sophisticated strategies [3, 41][6]. Even more worrying, the auctioneer was clearly not learning meaningful strategies — a typical evolved pricing rule was:

```
BidPrice - constant
```

which, once again, is not a terribly sophisticated strategy, and one that it is possible to imagine traders, more sophisticated than the ones we were able to co-evolve, learning to exploit.

4 Evolving traders

One of the problems we identified with our attempt to evolve both traders and auctioneer from scratch was that this approach makes it too hard to learn sophisticated strategies for making offers. Starting, as is standard in genetic programming, from random strategies[7] means that the traders have to make huge strides to reach even the same level of sophistication as, for example, ZIP. Since traders can achieve reasonable levels of profit with the fixed margin rules we were discovering, there is little evolutionary pressure for them to continue to evolve, and lots of competitive strategies to drown out any mutations that aren't immediately successful. These observations led us to try to learn new trading stratgies by starting from existing strategies.

As described in [33], we adopted the *heuristic strategy analysis* of Walsh *et al.* [49]. In its original form, the aim of this approach was to be able to compute plausible equilibria of the double auction. While performing a game theoretic analysis of the auction is infeasible (as discussed above) because of the number of players and the large number of possible actions at each of the many stages, it is possible to analyse double auctions at higher level of abstraction. The idea is to reduce the game to that of picking the best trading strategy from the literature. Thus, if you are interested in auctions with 10 participants, you pick a range of strategies for those participants, run

[6] The fixed margin strategy "Gamer" was not competitive in the Santa Fe tournament [41], and the fixed-markup strategy PS is one of the weakest strategies of those analysed in [3].

[7] That is strategies composed of randomly selected functions, not strategies that bid at random — the latter perform surprisingly well [15].

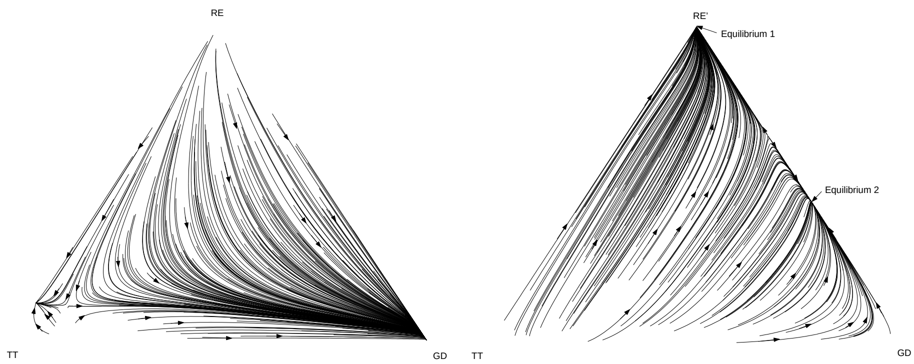

(a) (b)

Fig. 4. The replicator dynamics direction field for a 12-agent clearing-house with trading strategies RE, TT and GD, (a) with the payoffs obtained in experiments and (b) with +5% payoffs to RE.

a number of iterations of the auction, and that allows you to establish certain properties of the auction.

Now, it is clear that such a static analysis will not tell us much about the auction. Why should the participants in the auction pick the strategies that you choose, particularly if those strategies aren't very successful? To deal with this problem, Walsh *et al.* used *evolutionary game theory* [46] to compute Nash equilibria. The idea can be glossed as follows — rather than always selecting one strategy, traders are more likely to gradually adjust their strategy over time in response to to repeated observations of their own and others' payoffs. The adjustment can be modelled using the following *replicator dynamics* equation to specify the frequency with which different trading strategies should be used depending on our opponent's strategy:

$$\dot{m}_j = [u(e_j, \boldsymbol{m}) - u(\boldsymbol{m}, \boldsymbol{m})]\, m_j \tag{5}$$

where $\boldsymbol{m}$ is a mixed-strategy vector, $u(\boldsymbol{m}, \boldsymbol{m})$ is the mean payoff when all players play $\boldsymbol{m}$, and $u(e_j, \boldsymbol{m})$ is the average payoff to pure strategy j when all players play $\boldsymbol{m}$, and $\dot{m}_j$ is the first derivative of m_j with respect to time. Strategies that gain above-average payoff become more likely to be played, and this equation models a simple *co-evolutionary* process of mimicry learning, in which agents switch to strategies that appear to be more successful.

Now, for any initial mixed-strategy we can find the eventual outcome of this co-evolutionary process by solving $\dot{m}_j = 0$ for all j. This tells us the points at which the mixed strategy no longers changes — the stationary points of the replicator dynamics — and allows us to discover the final mixed-strategy that corresponds to the mixed strategy we started with. Repeating this for a range of initial mixed strategies allows us to discover all the stationary points that might develop. This model has the attractive properties that:

1. all Nash equilibria of the game are stationary points under the replicator dynamics; and

2. all focal points of the replicator dynamics are Nash equilibria of the evolutionary game.

Thus the Nash equilibrium solutions are embedded in the stationary points of the direction field of the dynamics specified by equation 5, and the replicator dynamics allows us to identify the Nash equilibria. Although not all stationary points are Nash equilibria, by overlaying a dynamic model of learning on the equilibria we can see which solutions are more likely to be discovered by *boundedly-rational* agents. Those Nash equilibria that are stationary points at which a larger range of initial states will end up, are equilibria that are more likely to be reached (assuming an initial distribution that is uniform).

Figure 4 (a) gives the direction field for a 12-agent clearing-house with traders allowed to pick between the RE, TT and GD strategies. This is a standard 2-simplex where the coordinates of any point represent a mixture of trading strategies. Each vertex denotes a situation in which all traders are using a single trading strategy. Any point on an edge of the simplex denotes a situation in which all traders using one of the two strategies denoted by the vertices joined by the side. Thus every point on the bottom of the simplex in Figure 4 (a) denotes a mixture of strategies such that some traders us TT and some use GD.

We can see that in Figure 4 (a) GD is a best-response to itself, and hence is a pure-strategy equilibrium. We also see it has a very large *basin of attraction* — for any randomly-sampled initial configuration of the population, most of the flows end up in the bottom-right-hand-corner. Additionally, there is a second mixed-strategy equilibria at the coordinates $(0.88, 0.12, 0)$ in the field, corresponding to an 88% mix of TT and a 12% mix of RE, however the attractor for this equilibrium is much smaller than the pure-strategy GD equilibrium; only 6% of random starts terminate here as against 94% for pure GD. Hence, according to this analysis, we would expect most of the population of traders to adopt the GD strategy.

From the point of view of evolving new trading strategies, the interesting thing is that GD is not as dominant as it might appear from Figure 4 (a). If we perform a sensitivity analysis to assess the robustness of GD's performance, by removing 2.5% of its payoffs and assigning them to RE, along with 2.5% of the payoffs from TT, then we get the direction field in Figure 4 (b). This second direction field gives us a qualitatively different set of equilibria — the RE strategy becomes a best-response to itself with a large basin of attraction (61%) — and allows us to conclude that a slightly improved version of RE can compete well against GD.

To test this conclusion, as described in [33], we used a genetic algorithm to search for such an improved version of RE, searching through parameter settings for a combination of four strategies — the original version of RE, a variation on RE introduced in [30], stateless Q-learning, and a strategy that randomly selects offers — evaluating the evolved strategies by the size of the basin of attraction they attain under the replicator dynamics. The GA converged on a version of stateless Q-learning, and Figure 5 shows how this optimised strategy OS performs against TT, GD, and the original version of RE. Our conclusion is that it is possible to evolve trading strategies to compete with the best hand-coded strategies provided that one has the hand-coded strategies to evolve against.

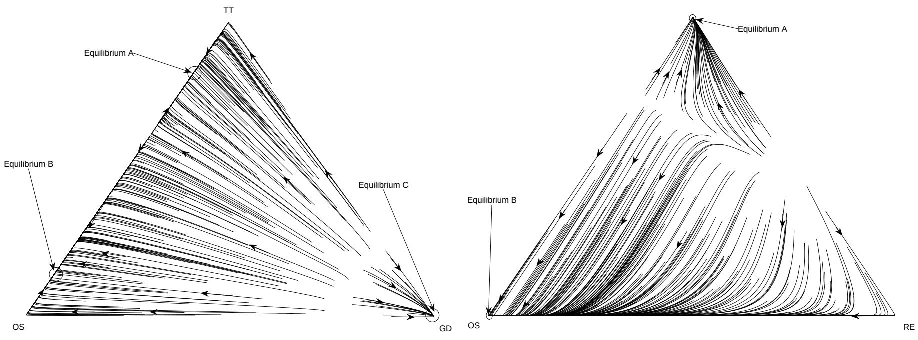

(a) (b)

Fig. 5. Replicator dynamics direction field for a 12-agent clearing-house auction showing interaction between the GA optimised strategy OS and (a) TT and GD, and (b) TT and RE.

5 Evolving mechanisms

One lesson to draw from Sections3 and 4 is that one can evolve better traders if that evolution takes place in a more structured way. Rather than evolving the OS strategy from scratch, we structured it as a search through the parameters of a set of existing trading strategies, and rather than evolving the auction pricing rule at the same time, we fixed the pricing rule to one that is commonly adopted. This section describes two experiments that we carried out to explore if it possible to do the reverse — evolve aspects of the auction mechanism given traders using a known trading strategy.

5.1 Evolving a pricing rule

The first experiment that we carried out in evolving parts of an auction mechanism separately from the traders is described in detail in [35], and considered the evolution of the rule for setting trade prices given the prices bid by buyers p_b and asked by sellers p_s. This work used a continuous double auction with 30 buyers and 30 sellers, all of them using the RE strategy to pick offers. To evaluate the rules we evolved, we used the measure F:

$$F = \frac{e_a}{2} + \frac{\widehat{mp_b} + \widehat{mp_s}}{4} \tag{6}$$

where e_a is as defined in (2), and mp_b and mp_s measure the market power of the buyers and sellers respectively, that is the extent to which the profits made by those groups differ from what they would be at theoretical equilibrium:

$$mp_b = \frac{pr_b - pr_b^e}{pr_b^e}$$

$$mp_s = \frac{pr_s - pr_s^e}{pr_s^e}$$

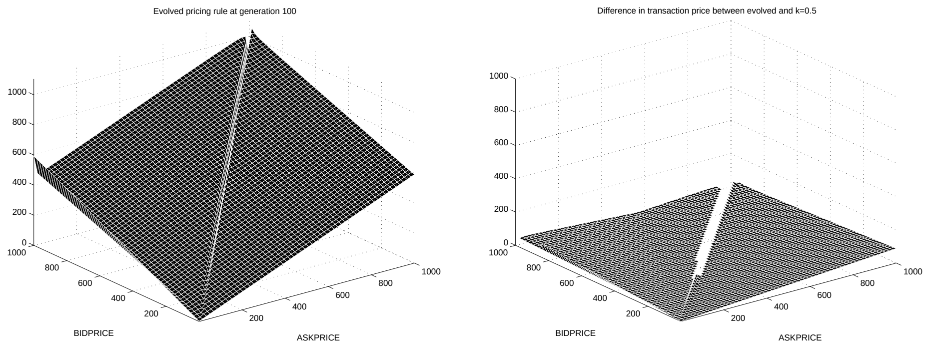

(a) (b)

Fig. 6. The results of evolving a pricing rule, (a) the rule itself, (b) the difference between the rule and the $k = 0.5$ rule.

and $\widehat{mp_b}$ and $\widehat{mp_s}$ are normalized versions of these measures:

$$\widehat{mp_b} = \frac{1}{1 + mp_b}$$
$$\widehat{mp_s} = \frac{1}{1 + mp_s}$$

We used genetic programming to evolve the rules. The functions that could be used in the genetic program consisted of the terminals $ASKPRICE$ and $BIDPRICE$, representing p_a and p_b respectively, together with the standard set of arithmetic functions $\{+, -, \times, \div\}$, and a function representing a random constant in the range $[0, 1]$. Thus all we assumed about the pricing function is that it was an arithmetic function of the bid and ask prices.

Pricing rules were evaluated using the measure F from (6). The following:

```
((0.6250385(0.93977016(ASKPRICE+0.76238054)))
 + (((((-0.19079465)/(ASKPRICE-(((BIDPRICE +BIDPRICE)/
 (((((ASKPRICE-1)+1.6088724)/(((1-ASKPRICE) -(ASKPRICE/
 ASKPRICE))+(2.5486426+(BIDPRICE + 0.000012302072))))
 +((BIDPRICE/ASKPRICE)+((BIDPRICE+BIDPRICE)+(1.430315)/
 (BIDPRICE . ASKPRICE)))))ASKPRICE)) ...
```

are the first few terms of the actual pricing rule that was evolved after 90 generations. It has been algebraically simplified, but as can be seen it is still far from straightforward, something that is not surprising given the way that standard genetic programming approaches handle the evolution of a program. Plotting the surface of the transaction price as a function of p_b and p_a, given in Figure 6 (a), and comparing it with the surface for:

$$0.5p_a + 0.5p_b$$

shows — the difference between the two rules is given in Figure 6 (b) — that these two functions are approximately equal apart from a slight variation when the ask price is very small or when the ask price is equal to the bid price. Thus the experiment effectively evolved a pricing rule for a discriminatory-price k double auction with $k = 0.5$ from the space of all arithmetic functions of ask and bid price. Our main conclusion from this is that our approach is able to evolve an eminently sensible rule, since the rule it came up with is virtually indistinguishable from one that has been widely used in practice[8].

5.2 Minimising price fluctuation

The work described in the previous section looked at optimizing one very specific part of the continuous double auction, the rule for setting trade prices, with respect to one specific measure, that in (6). We can apply the same kind of optimization to different aspects of the auction mechanism, and with different measures in mind. [31] describes some experiments with some alternatives.

In particular, [31] is concerned with minimising Smith's measure α (3), and thus fluctuations in the transaction price of the auction. The choice to minimise α was partly in order to see if it was possible to minimise this metric while keeping the efficiency of the auction high — testing the extent to which performance of the auction could be optimised — but one can imagine that this is also an attractive feature of an auction. If the auction has a low α, then transactions are, by definition, close to the theoretical equilibrium point. If this can be achieved for a range of trading strategies, then there is some guarantee that, no matter how a trader bids, the price that trader pays will be in some sense fair, saving the trader the burden of needing to bid cleverly.

To minimise α, we looked at learning a new pricing rule, a rule between that often used in a continuous double auction — where the price is the average of the bid and the ask — and the usual rule for a clearing house auction — where the price is the price that clears the market[9]. In essence, this new rule looks at the n most recent matching bid/ask pairs, and averages over them to obtain the transaction price. Figure 7 (a) compares the value of α for a continuous double auction with 10 buyers and 10 sellers all of which trade using the ZIC strategy and the $k = 0.5$ pricing rule with that of the value of α for the same auction that sets prices using the average of the last 4 matched sets of bid and ask[10]. We only considered auctions involving ZIC traders in order to make the problem of minimising price fluctuation as hard as possible — ZIC, making offers randomly, has leads to high values of α compared with other trading strategies.

Clearly the moving average rule is effective in reducing α, but the value it attains is still high compared with the levels attained using different trading strategies. Auctions

[8] It is also possible to argue in the other direction — that since we came up with the $k = 0.5$ rule, the rule makes sense for scenarios like the one that we were investigating.

[9] The price that would be the theoretical equilibrium if the bids and asks were truthful.

[10] Note that the rule uses *at most* four sets of matched bids and asks. Since the auction is continuous, the price of the trade between the first matched pair of offers is exactly that of the $k = 0.5$ rule since there is only one matched bid and ask pair to use, the price of the second trade is the average the first two matched bids and the first two matched asks and the price of the third trade is the average of the first three matched sets.

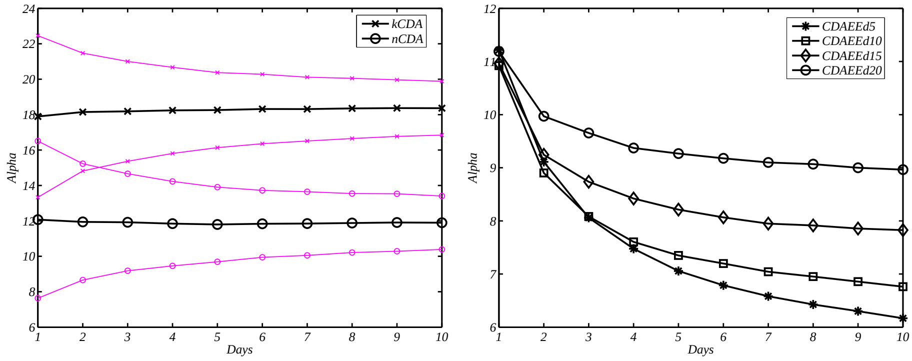

(a) (b)

Fig. 7. The value of α for a double auctions with ZIC traders (a) comparing the standard continuous double auction price rule with the sliding windowpricing rule, and (b) comparing α for different versions of the shout improvement rule. The grey lines denote one standard deviation above and below the average value over 400 iterations.

with traders that only use GD attain α values of around 4. To try to reduce fluctuation even more, we examined another aspect of the auction, the rule the auctioneer uses for accepting shouts as being valid. The idea is to generalize the "New York Stock Exchange rule", the rule used in that market, among others, which insists that successive bids and asks for the same good improve on each other. In other words, sucessive bids must increase, and successive asks must decrease. The generalization we adopted makes a running estimate of the equilibrium price for the market, and the shout acceptance rule (which we call the "shout improvement" rule) requires that bids are above this estimate and asks are below it. Note that our rule, unlike the NYSE rule, continues to apply after an individual good has been traded — indeed, as Figure 7 (b) shows, the effect of the rule on α improves over time.

In fact, it turns out that there is one last tweak to the shout improvement rule that it behooves us to add. If the rule is applied strictly as described, the estimate of the equilibrium price can get badly thrown off by errant offers at the start of the auction (and errant offers are common with ZIC traders). To ameliorate this situation, we introduce a parameter δ, an increment that is applied to the estimated equilibrium price to relax the improvement rule — bids above the estimate minus δ and asks below the estimate plus δ are considered valid. Figure 7 (b) shows the effect of different values of δ.

Overall, the combination of these measures can reduce α for all-ZIC markets to a value around 6 with little or no loss of efficiency. Indeed, for some values of δ, efficiency of the all ZIC market is greater than that of an all ZIC market under the usual CDA mechanism. In addition, it seems that these new market rules do not hurt the performance of markets consisting of more sophisticated traders. We tested the same market rules when the traders all used GD, and found that, if anything, the new rules reduced α and increased efficiency.

6 Evaluating mechanisms

The work on mechanisms that we have described so far optimized one specific aspect of an auction, and showed that this was achievable. However, the kind of evaluation of auctions that we have used in this work, focussing on a single measure when agents are all of the same type — in the sense of which bidding strategy they used — seems a bit narrow, and so we have experimented with alternative forms of evaluation of mechanism.

6.1 Comparing markets

In [36] we experimented with using heuristc strategy analysis to compute metrics for different types of auction. The motivation for doing this is as follows. Most of the properties that we might use to rate auctions, whether efficiency or Smith's α as above, or metrics like price dispersion [15], differ for the same auction as the traders use different trading strategies. They are not properties of the traders, since the same traders generate different efficiencies, αs and price dispersions in different auctions, but they are not entirely properties of the auctions either. Thus it is difficult to say with authority that a given auction has a given property. What we can do, however, is to use a heuristic strategy analysis to establish what mixtures of trading strategies will hold at equilibrium, and use this to compute an estimate of the properties that we are interested in.

Figure 8 shows the results of a heuristic strategy analysis for the continuous double auction and the clearing house auction with different numbers of traders. For all of these analyses we used three trading stategies, truth telling TT, the Roth-Erev strategy RE that we used in the first pricing rule experiment, and the modification of ZIP proposed by Preist and van Tol (PVT) [37]. Our choice of stratgies was intended to examine the relative performance of the human-like RE strategy[11], the simple "program trader" provided by PVT, with the performance of TT measuring how far the markets are from being strategy-proof (in a strategy-proof market there is no advantage to not telling the truth about one's valuation for a good).

There are a number of conclusions that one can draw from the plots in Figure 8[12]. First, there is a significant difference between the direction fields of the continuous double auction and the clearing house auction for any number of traders. While each strategy is pure strategy equilibrium in all cases, the basins of attraction are rather different as are the locations, along the edges and in the middle of the direction field, of the mixed equilibria. Second, the difference becomes more marked the larger the number of agents — the basin of attraction of TT shrinks as the CDA includes more traders, and grows as the CH includes more traders. The latter is in accordance with theoretical results [42] which predict that the disadvantages of truth-telling decline as the number of traders grows. Third, truth telling is not dominant in any of the markets, so none of them are strategy proof.

[11] Roth and Erev originally introduced their approach as a way of replicating human behavior in games [39].

[12] Note that we have not indicated the direction of the field on the plots in the interests of readability — the direction of flow is from the middle of the simplex towards the edges.

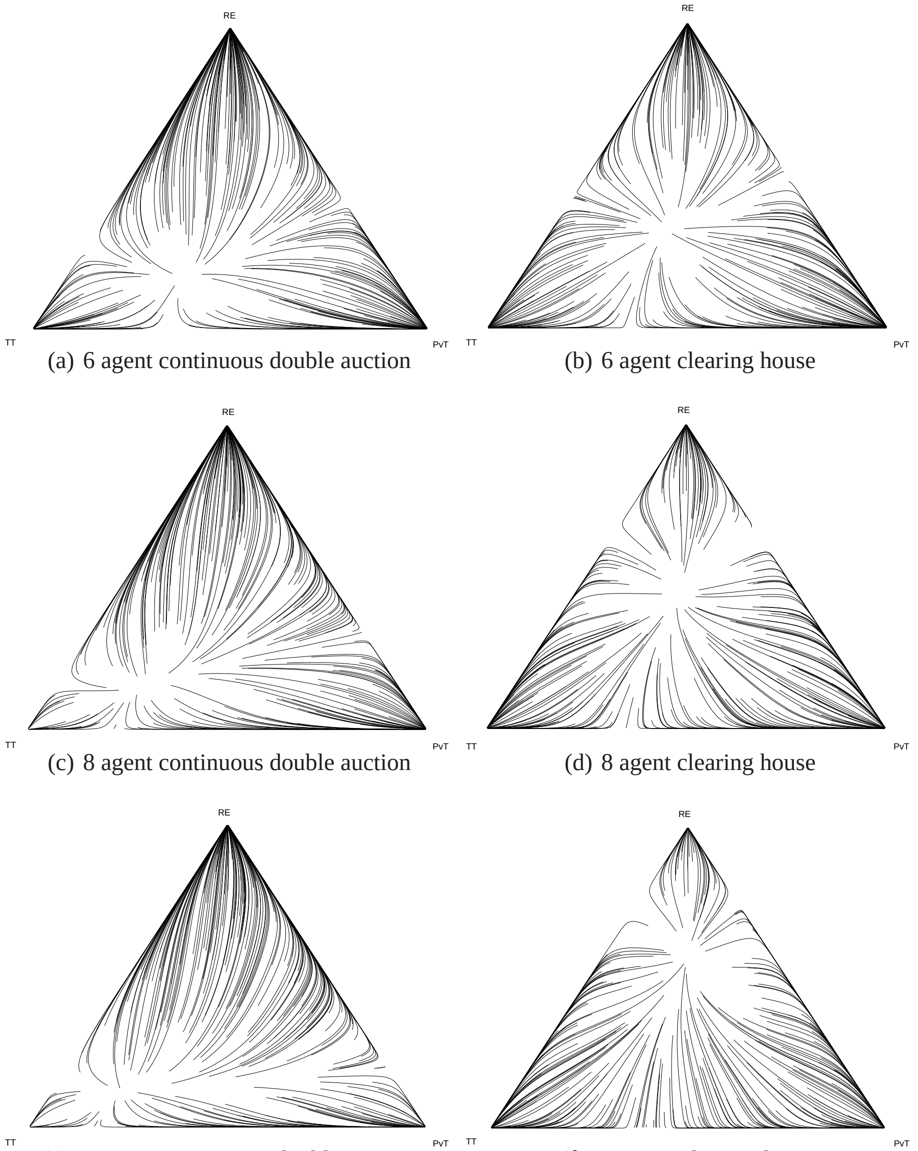

(a) 6 agent continuous double auction

(b) 6 agent clearing house

(c) 8 agent continuous double auction

(d) 8 agent clearing house

(e) 10 agent continuous double auction

(f) 10 agent clearing house

Fig. 8. The replicator dynamics direction field for double auctions with trading strategies TT, RE and PVT.

Equilibrium	CH probability	payoff	CDA probability	payoff
TT	0.38	1.00	0.05	0.86
RE	0.11	0.99	0.70	0.97
PvT	0.51	0.99	0.25	0.94

Table 1. Probabilities of equilibria for 10 agent markets

It is also possible to draw more quantitative conclusions. Taking 1000 random starting points within the direction fields for each of the 10 agent clearing house and continously double auctions, we established which of the pure strategy equilibria these starting points led to. Assuming that the starting points, each of which represents a mix of trading strategies, are equally likely, we could then compute the relative frequency of occurrence of the pure strategies — these are given in Table 1. Now, since we can easily establish whatever metrics we want for the equilibrium points (again these are given in Table 1), we can use the probabilities of reaching these equilibria to determine the expected value of the metrics. For example for the 10 trader CDA we can compute the expected efficiency as:

$$0.05 \times 0.86 + 0.70 \times 0.97 + 0.25 \times 0.94 = 0.96$$

compared with

$$0.38 \times 1.00 + 0.11 \times 0.99 + 0.51 \times 0.99 = 0.99$$

for the 10 trader CH.

Note that the assumption of equally likely start points is not the only assumption involved in this computation. Since the probability of arriving at a particular equilibrium is a function of the replicator dynamics, we are also assuming that the replicator dynamics is an accurate description of trader behavior. One can argue this either way — the only guarantee that the replicator dynamics give is that the stationary points in the field are Nash equilibria.

6.2 Direct competition between markets

The comparison between markets described above is useful, but indirect. It compares markets while still thinking of the markets as operating in isolation — it tells us nothing about how the markets would fare if they were in running in parallel, as markets often do in the real world[13]. In [32], we looked at the relative performance of markets when they are in competition with one another for traders.

To this end, we ran a series of experiments[14] where traders were offered a choice of markets at the start of every trading day, making this choice using simple reinforcement

[13] For example, Shah and Thomas [44] describe the competition between India's National Stock Exchange and the established Bombay Stock Exchange for trade in the stock of Indian companies when the National Stock Exchange opened.

[14] These experiments were run using JCAT [21], an extension of JASA [20] that allows multiple markets and provides both a mechanism for markets to charge traders, and for traders to decide which market provides them with the best profit.

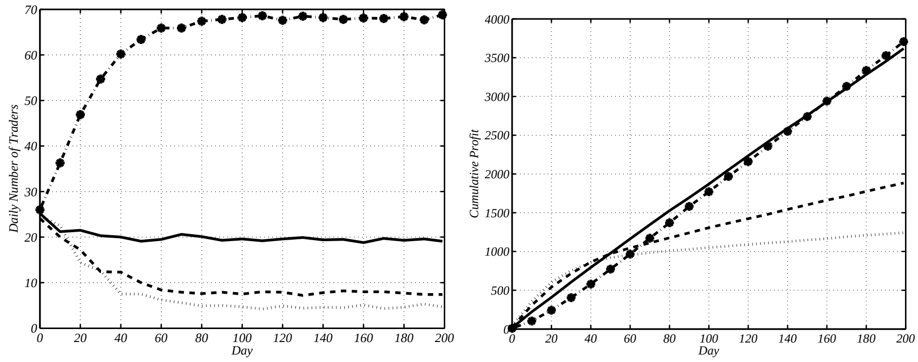

(a) Number of traders each day

(b) Cumulative profits

Fig. 9. Four markets that compete for traders.

learning based on the profit that they made in the markets (learning which markets were profitable over time). The profit that a given trader makes in a market is the profit from trade — the difference between the trade price and the private value of the trader — minus any charges imposed by the markets. We allowed markets to charge because this is a feature of real markets, and because the profit made by markets is a natural basis of comparison.

Figure 9 shows some typical results from [32]. Figure 9 (a), which gives the number of traders in each market at the start of every trading day, shows how, as the agents learn, the markets stratify by charges. As one might expect, the lowest charging market attracts the largest number of traders and the highest charging market attracts the smallest number of traders. Note that even the highest charging market continues to attract some traders — those that make a good profit even with the charges. Figure 9 (b), which gives the cumulative profit of each market on each day, shows how the lowest charging market catches the higher charging markets over time. These results are for markets with a simple, fixed, policy for charging. [32] also considers adaptive charging policies — one that undercuts all other markets, one that cuts prices until it has a large market share and then increases prices, and one that works like ZIP — showing that the relationship between such policies has some of the complexity of the relationship between trading strategies.

7 Conclusion

Auctions are a powerful mechanism for resource allocation in multiagent systems and elsewhere, and there are many situations in which one might make use of them. However, it is not advisable to employ auctions “off-the-peg” — as some expensive failures have demonstrated [28] — instead, it is necessary to carefully tailor auction mechanisms for the particular niche that they are required to fill. Our work is intended automate this tailoring process. Using a combination of evolutionary computation, reinforcement learning, and evolutionary game theory, we have successfully tailored variants of the double auction for different purposes, and traders to operate in these auctions, and our future work aims to extend the scope of this automated generation. In particular, we can easily imagine combining the techniques we have described here into a high-level process for co-evolving markets and traders. For a fixed mechanism we could evolve traders, as in Section 4, and then fix the equilibrium set of traders and evolve parts of the mechanism as in Section 5, evaluating evolved mechanisms just as we did in Section 6.1. Repeating this process will then allow us to create traders that can operate in the new mechanism. Demonstrating this co-evolutionart spiral is the focus of our current work.

Acknowledgments

This work was supported by the National Science Foundation under grant IIS-0329037 and by the European Union under grant IST-1999-10948. We are grateful to Marek Marcinkiewicz, Jarrett Chaiken and Ravi Cohen for their work on JASA, Jamari Robinson and Victoria Au for their work on JCAT, and to Jeff Mackie-Mason and Michael Wellman for helpful comments.

References

1. P. J. Angeline and J. B. Pollack. Competitive environments evolve better solutions for complex tasks. In S. Forrest, editor, *Proceedings of the Fifth International Conference on Genetic Algorithms*, pages 264–270, 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. K. Cai, J. Niu, and S. Parsons. An empirical game-theoretic analysis of trading strategies in a continuous double auction market. In *Proceedings of the Symposium on Adaptive Learning Agents and Multiagent Systems*, 2007. (this volume).
4. D. Cliff. Minimal-intelligence agents for bargaining behaviours in market-based environments. Technical Report HP-97-91, Hewlett-Packard Research Laboratories, Bristol, England, 1997.
5. D. Cliff and J. Bruten. Less than human: Simple adaptive trading agents for CDA markets. Technical Report HP-97-155, Hewlett-Packard Research Laboratories, Bristol, England, 1997.
6. D. Cliff and J. Bruten. Zero is not enough: On the lower limit of agent intelligencefor continuous double auction markets. Technical Report HP-97-155, Hewlett-Packard Research Laboratories, Bristol, England, 1997.
7. D. 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.
8. V. Conitzer and T. Sandholm. Automated mechanism design: Complexity results stemming from the single-agent setting. In *Proceedings of the International Conference on Electronic Commerce*, pages 17–24, Pittsburgh, September 2003.
9. V. Conitzer and T. Sandholm. An algorithm for automatically designing deterministic mechanisms without payments. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 128–135, New York, July 2004.

10. P. Cramton and J. Schwartz. Collusive bidding: Lessons from the FCC spectrum auctions. *Journal of Regulatory Economics*, 17(3):229–252, May 2000.
11. I. Erev and A. E. Roth. Predicting how people play games with unique, mixed-strategy eqilibria. *American Economic Review*, 88:848–881, 1998.
12. S. G. Ficici and J. B. Pollack. Challenges in coevolutionary learning: Arms-race dynamics, open-endedness, and mediocre stable states. In *Proceedings of the 6th International Conference on Artificial Life"*, pages 238–247, 1998.
13. D. Friedman. The double auction institution: A survey. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, pages 3–25. Perseus Publishing, Cambridge, MA, 1993.
14. S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behaviour*, 22:1–29, 1998.
15. D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial sustitute for individual rationality. *The Journal of Political Economy*, 101(1):119–137, February 1993.
16. D. K. Gode and S. Sunder. Lower bounds for efficiency of surplus extraction in double auctions. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, pages 199–219. Perseus Publishing, Cambridge, MA, 1993.
17. W. D. Hillis. Co-evolving parasites improve simulated evolution as an optimization procedure. In *Proceedings of the Second International Conference on Artificial Life*, pages 313–324, 1992.
18. J. Holland. *Adaptation in natural and artificial systems*. University of Michigan Press, Ann Arbor, MI, 1975.
19. M. O. Jackson. Mechanism theory. In U. Devigs, editor, *Optimization and Operations Research*, The Encyclopedia of Life Support Science. EOLSS Publishers, Oxford, UK, 2003. The working paper version of this article includes a more comprehensive bibliography and some bibliographic notes.
20. `http://sourceforge.net/projects/jasa/`.
21. `http://jcat.sourceforge.net/`.
22. P. Klemperer. How (not) to run auctions: The European 3G telecom auctions. *European Economic Review*, 46:829–845, 2002.
23. J. R. Koza. *Genetic programming: On the programming of computers by means of natural selection*. MIT Press, Cambridge, MA, 1992.
24. A. Madhavan. Trading mechanisms in securities markets. *The Journal of Finance*, 47(2):607–641, June 1992.
25. E. S. Maskin and J. G. Riley. Auction theory with private values. *American Economic Review*, 75:150–155, 1985.
26. R. P. McAfee. A dominant strategy double auction. *Journal of Economic Theory*, 56:343–450, 1992.
27. R. P. McAfee and J. McMillan. Analyzing the airwaves auction. *Journal of Economic Perspectives*, 10:159–176, 1996.
28. J. McMillan. Selling spectrum rights. *Journal of Economic Perspectives*, 8:191–199, 1994.
29. 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.
30. 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, 2001.
31. J. Niu, K. Cai, S. Parsons, and E. Sklar. Reducing price fluctuation in continuous double auctions through pricing policy and shout improvement. In *Proceedings of the 5th International Conference on Autonomous Agents and Multi-Agent Systems*, Hakodate, Japan, 2006.

32. J. Niu, K. Cai, S. Parsons, and E. Sklar. Some preliminary results on competition between markets for automated traders. In *Proceedings of the Workshop on Trading Agent Design and Analysis*, Vancouver, BC, 2007.
33. S. Phelps, M. Marcinkiewicz, S. Parsons, and P. McBurney. A novel method for automated strategy acquisition in n-player non-zero-sum games. In *Proceedings of the 5th International Conference on Autonomous Agents and Multi-Agent Systems*, Hakodate, Japan, 2006.
34. S. Phelps, P. McBurney, S. Parsons, and E. Sklar. Co-evolutionary mechanism design: A preliminary report. In J. Padget, O. Shehory, D. Parkes, N. Sadeh, and W. E. Walsh, editors, *Agent Mediated Electronic Commerce IV: Designing Mechanisms and Systems*, number 2531 in Lecture Notes in Artificial Intelligence, pages 123–142. Springer Verlag, 2002.
35. S. Phelps, P. McBurney, E. Sklar, and S. Parsons. Using genetic programming to optimise pricing rules for a double auction market. In *Proceedings of the Workshop on Agents and Electronic Commerce*, Pittsburgh, PA, 2003.
36. S. Phelps, S. Parsons, and P. McBurney. An evolutionary game-theoretic comparison of two double auction markets. In P. Faratin and J. A. Rodríguez-Aguilar, editors, *Agent Mediated Electronic Commerce VI: Theories for and Engineering of Distributed Mechanisms and Systems*, number 3435 in Lecture Notes in Artificial Intelligence, pages 101–114. Springer Verlag, 2004.
37. C. Preist and M. van Tol. Adaptative agents in a persistent shout double auction. In *Proceedings of the 1st International Conference on the Internet, Computing and Economics*, pages 11–18. ACM Press, 1998.
38. A. E. Roth. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica*, 70:1341–1378, 2002.
39. 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.
40. J. Rust, J. H. Miller, and R. Palmer. Behaviour of trading automata in a computerized double auction market. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, pages 155–199. Perseus Publishing, Cambridge, MA, 1993.
41. J. Rust, J. H. Miller, and R. Palmer. Characterizing effective trading strategies. *Journal of Economic Dynamics and Control*, 18:61–96, 1994.
42. A. Rustichini, M. A. Satterthwaite, and S. R. Williams. Convergence to efficiency in a simple market with incomplete information. *Econometrica*, 62(5):1041–1063, September 1994.
43. M. A. Satterthwaite and S. R. Williams. The Bayesian theory of the *k*-double auction. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, pages 99–123. Perseus Publishing, Cambridge, MA, 1993.
44. A. Shah and S. Thomas. David and Goliath: Displacing a primary market. *Global Financial Markets*, 1(1):14–21, 18th June 2000.
45. V. L. Smith. An experimental study of competitive market behaviour. *The Journal of Political Economy*, 70(2):111–137, April 1962.
46. V. L. Smith. Microeconomics as a experimental science. *The American Economic Review*, 72(5):923–955, December 1982.
47. 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.
48. W. Vickrey. Counterspeculation, auctions, and competitive sealed bids. *Journal of Finance*, 16(1):8–37, 1961.
49. W. E. Walsh, R. Das, G. Tesauro, and J. O. Kephart. Analyzing complex strategic interactions in multi-agent systems. In P. Gymtrasiwicz and S. Parsons, editors, *Proceedings of the 4th Workshop on Game Theoretic and Decision Theoretic Agents*, 2001.
50. M. P. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, 1:1–23, 1993.

51. P. R. Wurman, W. E. Walsh, and M. P. Wellman. Flexible double auctions for electronic commerce: Theory and applications. *Decision Support Systems*, 24:17–27, 1998.
52. W. Zhan and D. Friedman. Markups in double auction markets. Technical report, LEEPS, Department of Economics, University of Santa Cruz, 2005.