Using Population-based Search and Evolutionary Game Theory to Acquire Better-response Strategies for the Double-Auction Market

Steve Phelps and Peter McBurney
Department of Computer Science
University of Liverpool
Chadwick Building
Liverpool L69 7ZF
United Kingdom

Marek Marcinkiewicz and Simon Parsons
Department of Computer Science
Graduate Center
City University of New York
New York 10016
USA

Abstract

We present a novel method for automatically acquiring strategies for the double-auction by combining evolutionary optimization together with a principled game-theoretic analysis. Previous studies in this domain have used standard co-evolutionary algorithms, often with the goal of searching for the "best" trading strategy. However, we argue that such algorithms are often ineffective for this type of game because they fail to embody an appropriate game-theoretic *solution-concept*, and it is unclear, what, if anything, they are optimizing. In this paper, we adopt a more appropriate criteria for success from evolutionary game-theory based on the likely adoption-rate of a given strategy in a large population of traders, and accordingly we are able to demonstrate that our evolved strategy performs well.

1 Introduction

The automatic discovery of game-playing strategies has long been considered a central problem in Artificial Intelligence. In Evolutionary Computing, the standard technique for discovering new strategies is *co*-evolution, in which the fitness of each individual in an evolving population of strategies is assessed relative to other individuals in that population by computing the payoffs obtained when the selected individuals interact. Co-evolution can sometimes result in *arms-races*, in which the complexity and robustness of strategies in the population increases as they counter-adapt to adaptations in their opponents.

Often, however, co-evolutionary learning can fail to converge on robust strategies. The reasons for this are many and varied; for example, the population may enter a limit cycle if strategies learnt in earlier generations are able to exploit current opponents and current opponents have "forgotten" how to beat the revived living fossil. Whilst many effective techniques have been to developed to overcome these problems, there remains, however, a deeper problem which is only beginning to be addressed successfully. In some games, such as chess, we can safely bet that if player A consistently beats player B, and player B consistently beats player C, then player A is likely to beat player C. Since the dominance relationship is transitive, we can build meaningful *rating systems* for objectively ranking players in terms of ability, and the use of such ranking systems can be used to assess the "external" fitness of strategies evolved using a co-evolutionary process and ensure that the population is evolving toward better and better strategies. In many other games, however, the dominance graph is highly intransitive, making it impossible to rank strategies on a single scale. In such games, it makes little sense to talk about "best", or even "good", strategies since even if a given strategy beats a large number of opponent strategies there will always be many opponents that are able to beat it. The best strategy to play in such a game is always dependent on the strategies adopted by one's opponents.

Game theory provides us with a powerful concept for reasoning about the best strategy to adopt in such circumstances: the notion of a *Nash equilibrium*. A set of strategies for a given game is a Nash equilibrium if, and only if, no player can improve their payoff by unilaterally switching to an alternative strategy.

If there is no dominant strategy[1] for the game, then we should play the strategy that gives us the best payoff based on what we believe our opponents will play. If we assume our opponents are payoff maximisers, then we know that they will play a Nash strategy set by *reductio ad absurdum*; if they did not play Nash then by definition at least one of them could do better by changing their strategy, and hence they would not be maximising their payoff. This is very powerful concept, since although not every game has a dominant strategy, every finite game posseses at least one *equilibrium* solution [Nash, 1950]. Additionally, if we know the entire set of strategies and payoffs, we can deterministically compute the Nash strategies. If only a single equilibrium exists for a given game, this means that, in theory at least, we can always compute the "appropriate" strategy for a given game.

Note, however, that the Nash strategy is not always the *best* strategy to play in all circumstances. For 2-player zero-sum games, one can show that the Nash strategy is not exploitable. However, if our opponents do not play their Nash strategy, then there may be other non-Nash strategies that are better at exploiting off-equilibrium players. Additionally, many equi-

[1] A strategy which is always the best one to adopt no matter what any opponent does.

libria may exist and in n-player non-constant-sum games it may be necessary for agents to *coordinate* on the same equilibrium if their strategy is to remain secure against exploitation; if we were to play a Nash strategy from one equilibrium whilst our opponents play a strategy from an alternative equilibrium we may well find that our payoff is significantly lower than if we had coordinated on the same equilibrium as our opponents.

2 Beyond Nash equilibrium

Standard games theory does not tell us which of the many possible Nash strategies our opponents are likely to play. *Evolutionary* game theory [Smith, 1982] and its variants attack this problem by positing that, rather than computing the Nash strategies for a game using brute-force and then selecting one of these to play, our opponents are more likely to gradually adjust their strategy over time in response to to repeated observations of their own and others' payoffs. One approach to evolutionary game-theory uses the *replicator dynamics* equation to specify the frequency with which different pure strategies should be played depending on our opponent's strategy:

$$\dot{m}_j = [u(e_j, \vec{m}) - u(\vec{m}, \vec{m})]\, m_j \tag{1}$$

where $\vec{m}$ is a mixed-strategy vector, $u(\vec{m}, \vec{m})$ is the mean payoff when all players play $\vec{m}$, and $u(e_j, \vec{m})$ is the average payoff to pure strategy j when all players play $\vec{m}$, 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.

For any initial mixed-strategy we can find the eventual outcome from this co-evolutionary process by solving $\dot{m}_j = 0$ for all j to find the final mixed-strategy of the converged population. This model has the attractive properties that: (i) all Nash equilibria of the game are stationary points under the replicator dynamics; and (ii) 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 1. 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).

This is all well and good in theory, but the model is of limited practical use since many interesting real-world games are *multi-state*[2]. Such games can be transformed into normal-form games, but only by introducing an intractably large number of pure strategies, making the payoff matrix impossible to compute. But what if we were to approximate the replicator dynamics by using an evolutionary search over the strategy space?

[2]The payoff for a given move at any stage of the game depends on the history of play.

Rather than considering an infinite population consisting of a mixture of all possible pure strategies, we use a small finite population of randomly sampled strategies to approximate the game. By introducing mutation and cross-over, we can search hitherto unexplored regions of the strategy space. Might such a process converge to some kind of approximation of a true Nash equilibrium? Indeed, this is one way of interpreting existing co-evolutionary algorithms; fitness-proportionate selection plays a similar role to the replicator dynamics equation in ensuring that successful strategies propagate, and genetic operators allow them to search over novel sets of strategies. There are a number of problems with such approaches from a game-theoretic perspective, however, which we shall discuss in turn.

Firstly, the proportion of the population playing different strategies serves a dual role in a co-evolutionary algorithm [Ficici and Pollack, 2003]. On the one hand, the proportion of the population playing a given strategy represents the probability of playing that pure strategy in a mixed-strategy Nash equilibrium. On the other hand, evolutionary search requires diversity in the population in order to be effective. This suggests that if we are searching for Nash equilibria involving mixed-strategies where one of the pure strategy components has a high frequency, corresponding to a co-evolutionary search where a high percentage of the population is adopting the same strategy, then we may be in danger of over-constraining our search as we approach a solution.

Secondly and relatedly, although the final set of strategies in the converged population may be best responses to each other, there is no guarantee that the final mix of strategies is not invadable by other yet-to-be-countered strategies in the search space, or strategies that became extinct in earlier generations because they performed poorly against an earlier strategy mix that differed from the final converged strategy mix. Genetic operators such as mutation or cross-over will be poor at searching for novel strategies that could potentially invade the newly established equilibrium because of the above problem. If these conditions hold, then the final mix of strategies is implausible as a true Nash equilibrium or ESS, since there will be unsearched strategies that could potentially break the equilibrium by obtaining better payoffs for certain players. We might, nevertheless, be satisfied with the final mix of strategies as an approximation to a true Nash equilibrium on the grounds that if our co-evolutionary algorithm is unable to find equilibrium-breaking strategies, then no other algorithm will be able to do so. However, as discussed above, we expect *a priori* that co-evolutionary algorithms will be particularly *poor* at searching for novel strategies once they have discovered a (partial) equilibrium.

Thirdly, in the case where there are multiple equilibria, the particular one to which our population converges will be highly sensitive to the initial configuration of the population, that is the particular mix of random strategies that we start with, and certain equilibrium solutions may only be obtainable if we start with a given mix of initial strategies. In evolutionary game theory, we can simply take many samples of initial mixed-strategy vectors and for each of them solve the replicator dynamics equation in order to find stationary points. However, such brute-force approaches require the sampling of many thousands of initial mixed strategies in order to accurately assess the population dynamics of a three-strategy game. If we translate this into a co-evolutionary algorithm with a large strategy space, it necessitates running the co-evolutionary process hundreds of thousands of times with different randomly initialised populations in order to discover *robust* equilibria, which is computationally impractical in most cases.

Finally, co-evolutionary algorithms employ a number of different selection methods, not all of which yield population dynamics that converge on game-theoretic equilibria [Ficici and Pollack, 2000].

These problems have led researchers in co-evolutionary computing to design new algorithms employing game-theoretic solution concepts [Ficici, 2004]. In particular, [Ficici and Pollack, 2003] describe a game-theoretic search technique for acquiring approximations of Nash strategies in large symmetric 2-player constant-sum games with type-independent payoffs. In this paper, we address n-player non-constant-sum multi-state games with type-dependent payoffs. In such games, playing our Nash strategy (or an approximation thereof) does not guarantee us security against exploitation, thus if there are multiple equilibria, it may be more appropriate to play a *best-response* to the strategies that we infer are in play.

3 Heuristic-strategy approximation

[Walsh *et al.*, 2002] obviate many of the problems of standard co-evolutionary algorithms by restricting attention to small representative sample of "heuristic" strategies that are known to be commonly played in a given multi-state game. For many games, unsurprisingly none of the strategies commonly in use is dominant over the others. Given the lack of a dominant strategy, it is then natural to ask if there are mixtures of these "pure" strategies that constitute game-theoretic equilibria.

For small numbers of players and heuristic strategies, we can construct a relatively small normal-form payoff matrix which is amenable to game-theoretic analysis. This *heuristic* payoff matrix is calibrated by running many iterations of the game; variations in payoffs due to different player-types (eg black or white, buyer or seller) or stochastic environmental factors (eg PRNG seed) are averaged over many samples of type information resulting in a single mean payoff to each player for each entry in the payoff matrix. Players' types are assumed to be drawn independently from the same distribution, and an agent's choice of strategy is assumed to be independent of its type, which allows the payoff matrix to be further compressed, since we simply need to specify the number of agents playing each strategy to determine the expected payoff to each agent. Thus for a game with k strategies, we represent entries in the heuristic payoff matrix as vectors of the form

$$\vec{p} = (p_1, \ldots p_k)$$

where p_i specifies the number of agents who are playing the ith strategy. Each entry $p \in P$ is mapped onto an outcome vector $q \in Q$ of the form

$$\vec{q} = (q_1, \ldots q_k)$$

where q_i specifies the expected payoff to the ith strategy. For a game with n agents, the number of entries in the payoff matrix is given by

$$s = \frac{n^k - 1}{(k-1)!} \quad (2)$$

For small n and small k this results in payoff matrices of manageable size; for $k = 3$ and $n = 6$, 8, and 10 we have $s = 28$, 45, and 66 respectively. Once the payoff matrix has been computed we can subject it to a rigorous game-theoretic analysis and search for Nash equilibria solutions and apply different models of learning and evolution, such as the replicator dynamics model, in order to analyse the dynamics of adjustment to equilibrium.

The equilibria solutions that are thus obtained are not rigorous Nash equilibria; there is always the possibility that an unconsidered strategy could invade the equilibrium. Nevertheless, heuristic-strategy equilibria are often more plausible as models of real-world game playing than those obtained using a standard co-evolutionary search precisely because they *restrict* attention to strategies that are commonly known and are in common use. We can therefore be confident that no commonly known strategy for the game at hand will break our equilibrium, and thus the equilibrium stands at least some chance of persisting in the short term future.

Of course, once an equilibrium is established, the designers of a particular strategy may not be satisfied with their strategy's adoption-rate in the game-playing population at large. As [Walsh *et al.*, 2002] suggest, the designers of, for example, a particular trading strategy in a market game may have financial incentives such as patent rights to increase their "market-share" – that is, the proportion of players using their strategy, or, in game-theoretic terms, the probability of their pure strategy being played in a mixed-strategy equilibrium with a large basin of attraction. They go on to propose a simple methodology for performing such optimization using manual design methods. A promising-looking candidate strategy is chosen for perturbation analysis; a new, perturbed, version of the original heuristic payoff matrix is computed in which the payoffs of the candidate strategy are increased by a small fixed percentage, thus modelling a hypothetical tweak to the strategy that yields in a small increase in payoffs. The replicator-dynamics direction field is then replotted to establish whether the hypothetically-optimized strategy is able to to achieve a high adoption rate in the population. Strategy designers can then concentrate their efforts on improving those strategies that become strong attractors with a small increase in payoffs.

In this paper, we extend this technique by using a genetic-algorithm (GA) to *automatically* optimize candidate strategies by searching for a hitherto-unknown best-response – or, to use more appropriate nomenclature, a *better-response* – to an existing mix of heuristic strategies. Rather than using a standard co-evolutionary algorithm to perform the optimization, we use a single-population GA where the fitness of an individual strategy is computed from the heuristic-strategy payoff matrix according to its expected payoff when it is played against the existing mixed strategy.

4 An HSA analysis of a double-auction

We apply our method to the acquisition of strategies for the *double-auction* [Friedman and Rust, 1991], a generalisation of more widely-known single-sided auctions such as the English ascending auction. Single-sided auctions involve a single seller trading with multiple buyers. In double auctions, we have multiple traders on both sides of the market. As well as soliciting offers from buyers, that is bids, we also solicit offers to sell, so called *asks*. Variants of the double-auction are used in many real-world markets, such as stock exchanges, in scenarios where supply and demand are highly dynamic. Whilst single-sided auctions are well-understood from a game theoretic perspective, double-sided auctions remain intractable to a full game-theoretic analysis when there are relatively few traders on each side of the market. Thus much analysis of this game has focused on using *agent-based computational economics* (ACE) [Tesfatsion, 2002] to explore viable bidding strategies.

[Phelps *et al.*, 2004] used a heuristic-strategy analysis to analyse two variants of the double-auction market mechanism populated with a mix of heuristic strategies, and were able to find approximated game-theoretic equilibrium solutions thereof. In this paper, we use the same basic framework, but we focus on the *clearing-house* double-auction (CH) [Friedman and Rust, 1991] with uniform pricing, in which all agents are polled for their offers before transactions take place, and all transactions are then executed at the same market-clearing price. We consider the following heuristic-strategies:

- The truth-telling strategy (TT), whereby agents submit offers equal to their valuation for the resource being traded (in a strategy-proof market, TT will be a dominant strategy);
- The Roth-Erev strategy (RE) – a strategy based on reinforcement-learning, described in [Erev and Roth, 1998] and calibrated with the parameters specified in [Nicolaisen *et al.*, 2001]; and
- The Gjerstad-Dickhaut strategy (GD) [Gjerstad and Dickhaut, 1998], whereby agents estimate the probability of any bid being accepted based on historical market data and then bid to maximize expected profit.

Since all mixed-strategy vectors lie in the unit-simplex, for $k = 3$ strategies we can project the unit-simplex onto a two dimensional space and plot the switching between strategies predicted by 2. Figure 1 shows the direction-field of the replicator-dynamics equation for these three heuristic strategies, showing that we have two equilibrium solutions—these are the points from which no arrows lead away. Firstly, since there are no arrows leading away from the bottom-right corner, we see that 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*; most randomly selected initial configurations will end up in the GD corner. Additionally, there is a second mixed-strategy equilibria at the coordinates (0.88, 0.12, 0) in the simplex 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 vs 94% for pure GD.

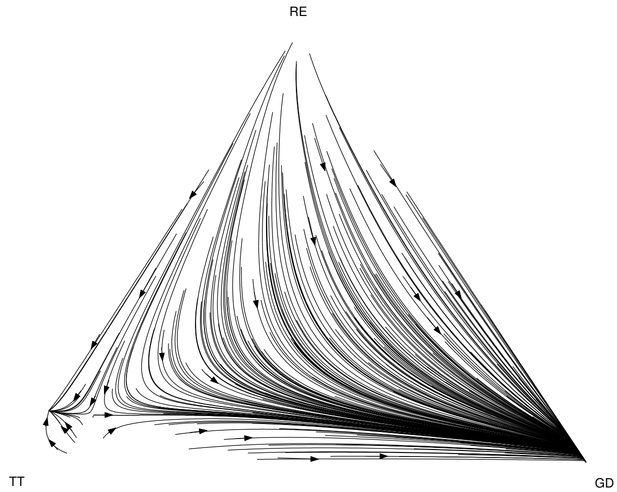

Figure 1: The original replicator dynamics direction field for a 12-agent clearing-house auction with the original unoptimized Roth-Erev strategy (labeled RE).

Hence, according to this analysis, we expect most of the population of traders to adopt the GD strategy.

How much confidence can we give to this analysis given that the payoffs used to construct the direction-field plot were estimated based on only 2000 samples of each game? One approach to answering this question is to conduct a sensitivity analysis; we perturb the mean payoffs for each strategy in the matrix by a small percentage to see if our equilibria analysis is robust to errors in the payoff estimates. Figure 2 shows the direction-field plot after we perform a perturbation where we remove 5% of the payoffs from the TT and GD strategies and assign +5% payoffs to the RE strategy. This results in a qualitatively different set of equilibria; the RE strategy becomes a best-response to itself with a large basin of attraction (61%), and we see a mixed-strategy equilibrium at (0, 0.45, 0.55) corresponding to a 45% mix of RE and a 55% mix of GD. Thus we conclude that our equilibrium analysis is sensitive to small errors in payoff estimates, and that our original prediction of widespread adoption of GD may not occur if we have underestimated the payoffs to RE.

If we observe a mixture of all three strategies in actual play, however, the sensitivity analysis suggests we could bring about widespread defection to RE if were able to tweak the strategy by improving its payoff slightly — *it points to RE as a candidate for potential optimization*.

5 Optimizing RE

The key question we need to answer in order to perform this optimization is *"What are we trying to optimize?"*. The obvious answer is that we should attempt optimize the payoff to our new version of the strategy. However, since we are unlikely to find a dominant strategy, we know that we are unlikely to be able to optimize our candidate strategy so that it obtains the maximal payoff no matter what our opponents do. We should therefore, attempt to optimize our strategy so that it gives us the best payoff against the strategy that we believe

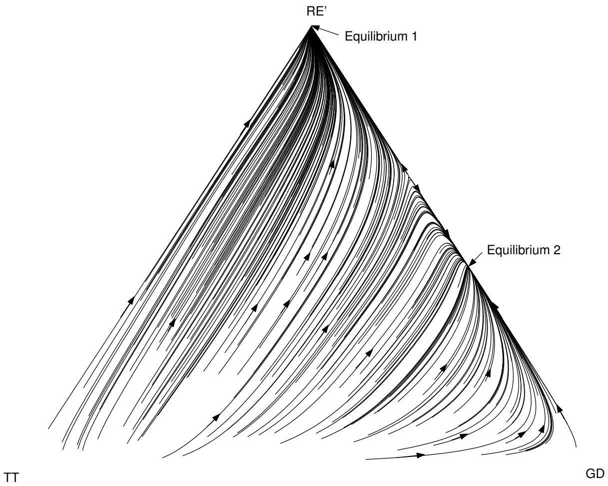

Figure 2: Replicator dynamics direction field for a 12-agent clearing-house auction perturbed with +5% payoffs to the Roth-Erev strategy (labeled RE')

will be actually played. The solution to this maximisation problem is likely to vary accordingly; one solution may do extremely well when faced with truth-tellers (TT), but not so well when faced with agents playing GD. The best-response strategy will depend on the precise *mix* of strategies we expect to encounter, and thus the solution space consists of an infinite pareto-frontier of possible best-responses parameterized by the mixed-strategy in play.

There are many possible techniques for inferring the likely mixed-strategy in use in an actual game. We might, for example, infer a particular mix based on an observation of play to date, in a manner akin to fictitious play. Alternatively, if we have no empirical data on play to date, or we wish to optimize a strategy for a market mechanism which has not yet been deployed, we might assume that players are likely to eventually adopt a heuristic-equilibrium strategy with a large basin of attraction under the replicator dynamics, and optimize accordingly. Such an analysis is beyond the scope of this paper. Instead, we pick an arbitrary mixed-strategy and use a genetic algorithm to optimize for this particular mixture. In the rest of the paper, the mixed-strategy that we optimize for is $(0.25, 0.25, 0.25, 0.25)$ where the corresponding pure strategies are (GD, TT, RE, OS), and OS refers to the strategy currently being evaluated.

5.1 Searching for a better-response

The RE strategy uses reinforcement learning (RL) to choose from n possible markups over the agent's limit price based on a reward signal computed as a function of profits earned in the previous round of bidding. Agents bid or ask at price p

$$p = l \pm mo \tag{3}$$

where l is the agent's limit price, o is the output from the learning algorithm and m is a scaling parameter. Additionally, the Roth-Erev learning algorithm itself has several free parameters: the recency parameter r, the experimentation parameter x, and an initialisation parameter $s1$. In addition to the original Roth-Erev (RE) algorithm, there are several other learning-algorithms that that have successfully been used for RL strategies in ACE. We search over three additional possibilities: stateless Q-learning (SQ), modifications to RE used by [Nicolaisen *et al.*, 2001] (NPT) and a control algorithm which selects a uniformly random action regardless of reward signal (DR). SQ has free parameters: the discount-rate g, epsilon e, and a learning-rate p.

Individuals in the search space were represented as a 50-bit genome, where:

- bits 1–8 coded for parameter m in the range $(1, 10)$;
- bits 9–16 coded for the parameters e or x in the range $(0, 1)$;
- bits 17–24 coded for parameter n in the range $(2, 258)$;
- bits 25–32 coded for parameters g or r in the range $(0, 1)$;
- bits 33–40 coded for parameter $s1$ in the range $(1, 15000)$;
- bits 41–42 coded for the choice of learning algorithm amongst RE, NPT, SQ or DR; and
- bits 43–50 coded for parameter p in the range $(0, 1)$.

This space was searched using a GA with a population size of 100, with single-point cross-over, a cross-over rate of 1, a mutation-rate of 10^{-4} and fitness-proportionate selection. The expected payoff to our candidate strategy was computed from the heuristic-strategy payoff matrix according to our benchmark mixed-strategy $(0.25, 0.25, 0.25, 0.25)$. This necessitated recomputing the payoff matrix for each individual that was evaluated. We used only 10 samples of the game in order to populate each entry in the payoff matrix in the expectation that the GA would be robust to the additional noise that this would introduce into the payoffs.

6 Results

Figure 3 shows the mean fitness of the evolving population per generation. By generation 50, the mean fitness had plateaued to 0.94 with a standard deviation of 0.03, and the fitness of the best individual was 0.99. The best individual coded for a strategy using the stateless Q-learning algorithm with parameters $n = 5$, $m = 5.39453125$, $e = 0.0234375$, $g = 0.1484375$ and $p = 0.1484375$.

The goal of this exercise was to see if we could find a replacement strategy for RE that would likely be adopted under replicator-dynamics learning given a population starting near the centre of the simplex. Figure 4 shows the direction-field of the replicator dynamics when we replace RE with our optimized strategy OS using 2000 samples of the game for each entry in the payoff matrix. Although our optimized strategy does not capture the greatest number of starting points, it is taken up in an equilibrium with a large basin of attraction where it is played with a high probability; equilibrium 1 is arrived at in 32% of cases, in this equilibrium OS is played with 98% probability. This gives us a total *expected* "market share" of approximately 31%, meaning that nearly a third of traders who start by randomly selecting a mixed strategy will

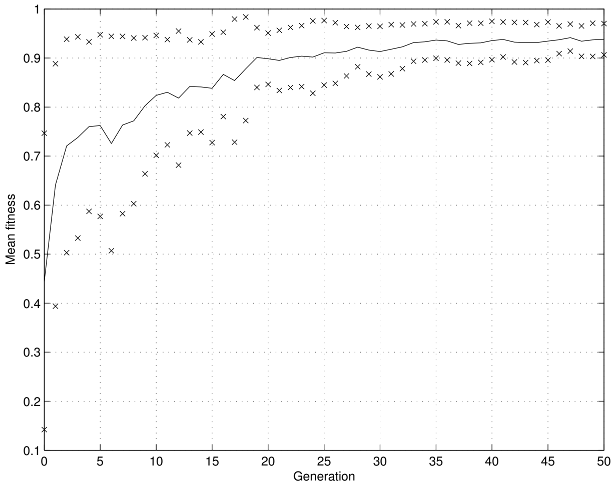

Figure 3: Mean fitness of the GA population with one standard deviation when optimizing for payoffs against the mixed strategy $(0.25, 0.25, 0.25, 0.25)$

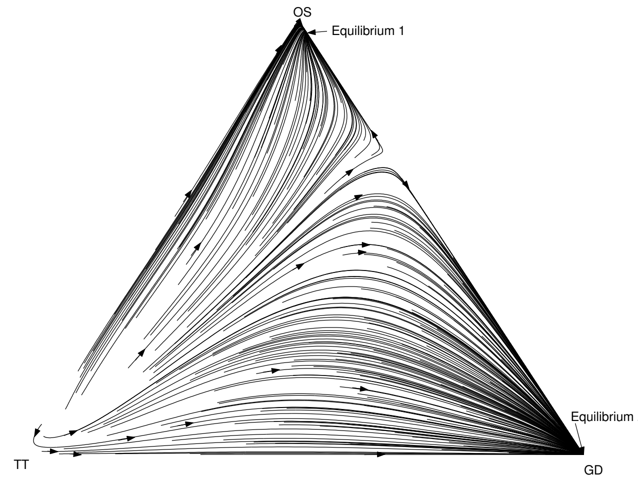

Figure 4: Replicator dynamics direction field for a 12-agent clearing-house auction using the GA-optimized strategy (labeled OS)

ned up playing OS in a given game. Although our optimized strategy is not predominant, it performs significantly better than the original RE strategy. Additionally, it succeeds in capturing defectors to our strategy when the population starts near the middle of the simplex – that is, when all strategies are being played with equal probability, which corresponds to our original design objective of maximising for payoffs against the mixed strategy $(0.25, 0.25, 0.25, 0.25)$.

7 Conclusion

In this paper we have applied a novel method combining evolutionary search together with a principled game-theoretic analysis in order to automatically acquire a trading strategy for the double-auction market. We defined an appropriate measure of success in this game based on evolutionary game-theory, and we were able to demonstrate that our evolved strategy performed robustly according to this criterion. This is a first step towards a principled approach to automatically creating new bidding strategies and auction mechanisms—no small matter given the current volume of auction-based trade. Now that we are able to create new bidding mechanisms, we can extending Cliff's work on automatically generating new auction mechanisms [Cliff, 2001], knowing that we can evolve bidding strategies that can work within such mechanisms, and so provide a meaningful comparison of those mechanisms using the technique of [Phelps *et al.*, 2004].

References

[Cliff, 2001] D. Cliff. Evolution of market mechanism through a continuous space of auction-types. Technical Report HPL-2001-326, HP Labs, 2001.

[Erev and Roth, 1998] I. Erev and A. E. Roth. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review*, 88(4):848–881, 1998.

[Ficici and Pollack, 2000] S. G. Ficici and J. B. Pollack. A game-theoretic approach to the simple coevolutionary algorithm. In *PPSN VI*. Springer Verlag, 16-20 2000.

[Ficici and Pollack, 2003] S.G. Ficici and J.B. Pollack. A game-theoretic memory mechanism for coevolution. In *LNCS 2723*, pages 286–297. Springer-Verlag, 2003.

[Ficici, 2004] S. Ficici. *Solution Concepts in Coevolutionary Algorithms*. PhD thesis, Brandeis University, May 2004.

[Friedman and Rust, 1991] D. Friedman and J. Rust. *The Double Auction Market*. Westview, 1991.

[Gjerstad and Dickhaut, 1998] S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behaviour*, 22:1–29, 1998.

[Nash, 1950] J. Nash. Equilibrium points in n-person games. In *Proc. of the National Academy of Sciences*, volume 36, pages 48–49, 1950.

[Nicolaisen *et al.*, 2001] J. Nicolaisen, V. Petrov, and L. Tesfatsion. Market power and efficiency in a computational electricity market with discriminatory double-auction pricing. *IEEE Trans. on Evol. Computation*, 5(5):504–523, October 2001.

[Phelps *et al.*, 2004] S. Phelps, S. Parsons, and P. McBurney. An evolutionary game-theoretic comparision of two double auction market designs. In *Agent-Mediated Electronic Commerce VI*. Springer-Verlag, 2004.

[Smith, 1982] J.M. Smith. *Evolution and the theory of games*. Cambridge University Press, 1982.

[Tesfatsion, 2002] L. Tesfatsion. Agent-based computational economics: growing economies from the bottom up. *Artificial Life*, 8(1):55–82, 2002.

[Walsh *et al.*, 2002] W.E. Walsh, R. Das, G. Tesauro, and J.O. Kephart. Analyzing complex strategic interactions in multi-agent games. In *AAAI-02 Workshop on Game Theoretic and Decision Theoretic Agents*, 2002.