Evolutionary Dynamics for Designing Multi-Period Auctions (Short Paper)

Tomas Klos*
Center for Mathematics and Computer Science
P.O. Box 94079
1090 GB Amsterdam
The Netherlands
tomas.klos@cwi.nl

Gerrit Jan van Ahee
Delft University of Technology
P.O. Box 5031, The Netherlands
2600 GA Delft
The Netherlands
van.ahee@yes2web.nl

ABSTRACT

Mechanism design (MD) has recently become a very popular approach in the design of distributed systems of autonomous agents. A key assumption required for the application of MD is that agents behave rationally in the mechanism or game, since this provides the predictability of agent behavior required for optimal design of the mechanism. In many cases, however, we are confronted with the intractability both of establishing rational equilibrium behavior, as well as of designing optimal mechanisms even if rational agent behavior can be assumed.

In this paper, we study both sides of the problem simultaneously by designing and analyzing a 'meta-game' involving both the designer of the mechanism (game, multi-agent system) and the agents interacting in the system. We use coupled replicator dynamics to investigate equilibrium outcomes in this game. In addition, we present an algorithm for determining the expected payoffs required for our analysis, thus sidestepping the need for extensive simulations as in previous work. Our results show the validity of the algorithm, some interesting conclusions about multi-period auction design, and the general feasibility of our approach.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Intelligent Agents, Multiagent systems*

General Terms

Design, Economics, Experimentation

Keywords

Mechanism Design, Auctions, Evolutionary Game Theory, Replicator Dynamics

1. INTRODUCTION

Mechanism Design has recently grown to become a very popular approach in the design of distributed systems of autonomous agents. Also sometimes called 'inverse game theory' [10], MD is concerned with designing the games or systems in which agents interact, and to do this in such a way that rational agent behavior leads to certain desirable properties for the system as a whole.

This undertaking depends crucially on the standard game theoretic assumption of agent rationality, since such rationality provides the predictability of agents' behavior required for optimizing the mechanism's design. In many practical circumstances, however, agents don't behave rationally. They may lack necessary information about the game or about other agents' beliefs, or the problem of optimizing their behavior may be too hard. Typically, therefore, the mechanism designer will want to implement the mechanism in dominant strategies, giving the agents a strategy which is optimal irrespective of the other agents' strategic choices. Making the agents' desired behavior dominant, however, may be very costly indeed.

An alternative is to implement a mechanism in Nash equilibria, but this requires the agents to find Nash equilibrium behavior in the game set up by the mechanism designer. Unfortunately, in general, not only finding Nash equilibria (in games of 2 or more players) [2], but also designing a mechanism implemented in Nash equilibria is computationally infeasible [4].

Because of these negative results, many have resorted to heuristic approaches to these two problems. For example, evolutionary methods are quite popular, for evolving both games as well as strategies, or even both [3, 14]. Here, we propose studying the interaction between the mechanism designer and the game participants as a higher level, meta-game. The designer chooses among alternative mechanism designs, while the agents choose among alternative strategies. We solve this game 'heuristically' using evolutionary game theory techniques, specifically, the replicator dynamics [6]. To illustrate, we adopt the multi-period auction scenario developed by Pardoe and Stone (henceforth PS) [11].

2. PARDOE AND STONE'S SCENARIO

Pardoe and Stone [11] studied a setting in which a seller wants to sell a large number of identical goods, and chooses to do so by repeatedly auctioning off batches of 60 items in a sequence of uniform-price sealed-bid auctions. In each auction (of 60 goods), all winning bidders pay the price of the highest losing bid. Bidders are not assumed to bid strictly rationally, but to probabilistically choose one of a limited set of 'heuristic' bidding strategies. Given such bidders, the

---

*Now at Delft University of Technology.

Cite as: Evolutionary Dynamics for Designing Multi-Period Auctions (Short Paper), Tomas Klos and Gerrit Jan van Ahee, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008,Estoril,Portugal,pp.1589-1592.

seller in turn, is unsure whether it is most profitable to auction off all 60 items at once, or to distribute them evenly over $2, 3$, or 4 'periods,' thereby revealing the winning price in between periods, and allowing the remaining bidders to update their bids.

There are 120 bidders for each auction, and each receives a 'signal,' indicating the value of each of the (identical) goods, drawn uniformly at random from $[0, 1]$. Each bidder may use any one of 5 different heuristic bidding strategies, as considered by PS:

equilibrium (EQ) (see [7, sec. 15.2]) This strategy bids the bidder's valuation (signal) times the fraction

$$\frac{\text{numberBidders} - \text{numberItems}}{\text{numberBidders} - \text{itemsSold}},$$

where itemsSold is the number of items sold up to and including the current period.

overestimate (OE) This strategy bids the same fraction as EQ, but overestimates numberBidders by 40.

underestimate (UE) This strategy similarly bids like EQ, but underestimates numberBidders by 40.

dropout (DO) These bidders use EQ, but leave the auction after each period with a probability of 0.15.

affiliated values (AV) These bidders bid $\frac{1}{4}$ of their signal in period 1, and thereafter value the item at the previous period's winning price, using EQ for bidding.

In what they call Adaptive Mechanism Design [12], PS study whether the seller can learn to adaptively optimize her auction's design (choice of number of periods), in the face of an unknown bidder population. They equip the seller with an ϵ-greedy reinforcement learning (RL) algorithm [17], and find that the seller can indeed approximate optimal revenue with a static bidding strategy distribution. Even in the case of a dynamically changing distribution (following a random walk), the average revenue obtained by the adaptive approach is higher than the revenue from sticking to a particular fixed number of periods.

3. EVOLUTIONARY DYNAMICS

We study this game in more detail, as a 'meta game' in which the seller chooses among alternative mechanism designs, while the agents choose among alternative strategies. We approximate the seller's stochastic RL algorithm used by PS, by a deterministic evolutionary process, the replicator dynamics (RD) [6], which turns out to be equivalent to RL under certain conditions [1, 15, 16, 18]. Also, instead of following a random walk, we assume new bidders to adopt heuristic strategies with probabilities proportional to their success in earlier auctions. This also leads to an RD, giving rise to connected replicator equations [8, 15, 16]. Our contribution lies not exactly in the knowledge this yields about designing multi-period auctions given the set of strategies used by Pardoe and Stone (also see our conclusions), but rather in the perspective on and analysis of designing mechanisms as a higher level game involving the mechanism designer and the agents participating in the game.

Various previous studies have investigated buyer-side evolution using RD, mainly in the context of the Continuous Double Auction (CDA) [13, 19]. This typically follows the

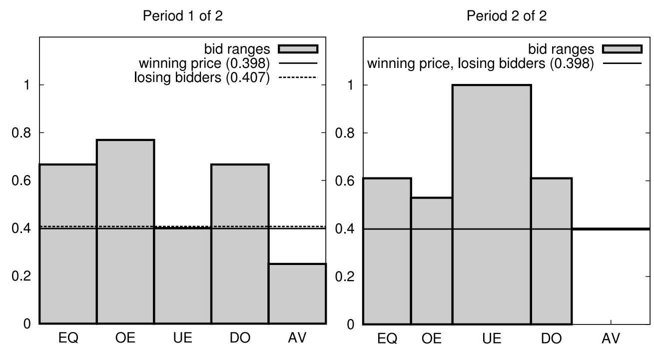

Figure 1: Our expected payoffs algorithm.

approach proposed by Walsh et al. [20], who suggest that in the case of heuristic strategies, (evolutionary) game theoretic analysis can proceed on the basis of a 'heuristic payoff matrix' for an N-person game. The entries of this matrix specify the payoffs to anyone employing a particular strategy when certain *numbers of* other players employ the other strategies. Assuming such symmetry reduces the size of the payoff matrix dramatically, but huge amounts of simulated games are typically still necessary to estimate all entries, and those are then valid only for the given numbers of players.

3.1 Expected Payoffs

We now present an efficient algorithm to determine (1) expected utilities for bidders using different strategies and (2) expected revenues for the seller choosing different numbers of periods, thus sidestepping the need for extensive computer simulations as in previous work [21]. Our algoritm is based on *proportions of bidders* using different strategies, which means we can handle arbitrary numbers of bidders. The method proposed by Walsh et al. [20] also works with probabilities of heuristic strategies being adopted, but their heuristic payoff matrix is given for specific *number of bidders* using each strategy, and quickly becomes too large, having $\binom{A+S-1}{A}$ entries for A agents and S strategies.

The operation of the algorithm is illustrated in Figure 1 for two periods and an initially uniform distribution over the bidding strategies. The width of each graph is 100% (of the remaining bidders), with the width of each bar indicating the proportion of each strategy in the population: for example, with a uniform distribution in the first period (left figure), each bar is equally wide. Vertically, each bar represents the range of the uniformly distributed bids placed by the bidders using each of the strategies. This range is given by the fractions in Sec. 2 applied to the range of valuations (signals). For example, in period 1, bidders using the strategies EQ and DO bid $(120 - 60)/(120 - 30) = 60/90 = 2/3$ times their valuation (uniform in $[0, 1]$), leading to uniformly distributed bids in $[0, \frac{2}{3}]$. Similarly, bidders using the AV strategy generate bids uniformly distributed in $[0, .25]$.

In the first of the 2 periods in this example, the 30 highest bidders win an item, and pay the 31st highest bidding agent's price. As percentages of the remaining number of bidders (120 in the first period), the lowest 75% of the bid densities determines who lose, while the lowest 74.16% determines the winning price. We draw horizontal lines cutting through the total (shaded) area covered by the bid ranges, while leaving the required percentages below them. For the first period, this yields a winning price of 0.398, and an ex-

Table 1: Example heuristic payoff table.

number	seller	average utility per strategy				
periods	revenue	EQ	OE	UE	DO	AV
1	21.875	0.202	0.202	0.202	0.202	0.0
2	23.897	0.181	0.181	0.181	0.178	0.0
3	24.348	0.174	0.166	0.190	0.162	0.0
4	24.130	0.176	0.163	0.197	0.157	0.0

pected revenue of $30 \cdot 0.398 = 11.94$.

After each period, we update the range of valuations remaining per strategy, after winning bidders (the ones with the highest valuations) have been removed. We also update the remaining proportions of the DO (-15%) as well as the other strategies. For example, none of the UE bidders are expected to have won in the first period, so 20% of the original population, i.e. approximately $20\%/75\% = 26.67\%$ of the second period population remains (actually 27.33% because of the dropout bidders), their valuations still distributed uniformly in $[0, 1]$. On the other hand, the highest valuing $(.667 - .407)/.667 = 39\%$ of the EQ bidders are expected to have won, leaving $[0 : .61]$ as the remaining range of valuations. In the second period (on the right), the algorithm works the same, so for example, the UE bidders now bid $((120 - 40) - 60)/((120 - 40) - 60) = 1$ times their valuation, yielding uniformly distributed (valuations and) bids in the range $[0, 1]$. Bidders using the AV strategy all have exactly the same valuation of 0.398—the previous period's winning price—and therefore also the same bid. The horizontal line (not a rectangle!) in the AV column thus completely captures its remaining proportion (also 27.33%), and in this case it is also the cutoff for both the losing bidders and the winning price, again giving an expected revenue of 11.94, yielding a total of 23.88 for 2 periods.

Expected revenues for all numbers of periods can thus efficiently be calculated for any distribution of bidding strategies. The corresponding expected utility for each bidding strategy can similarly be determined by subtracting the winning price from the mean of the range of winning bidders' valuations and renormalizing. Overall, we obtain the payoff matrix in Table 1 for a uniform distribution of bidding strategies. We immediately notice that the AV strategy generates no utility, and this happens in most if not all distributions. Also, the various strategies only generate different utilities for higher numbers of periods.

3.2 Replicator Dynamics

The RD are a popular and intuitive way of modeling deterministic evolutionary dynamics in games [6], highlighting the role of selection rather than mutation. With RD, the state of a population is represented as an n-dimensional vector $x = (x_1, \ldots, x_n)$ of relative frequencies x_i.

In each of a sequence of generations, the state of the population is evolved using the replicator equation:

$$\frac{dx_i}{dt} = (u(e_i, x) - u(x, x))x_i,$$

where e_i is pure (heuristic) strategy i, $u(e_i, x)$ is the expected payoff of strategy i when the population is in state x, and $u(x, x)$ is the expected payoff to an individual drawn at random from the population: $u(x, x) = \sum_{i=1}^{n} x_i u(e_i, x)$.

Applications of the replicator dynamics have often focused on settings involving 3 pure strategies, probably because this enables one to visualize the dynamics' direction fields on the 2-simplex (a triangle) graphically in 2 dimensions. Popular candidate games are Rock-Paper-Scissors (RPS) [6, 16], and the continuous double auction (CDA) [13, 19, 20], with 3 strategies selected from a set of published high-level heuristic strategies (e.g., ZIP, GD, Kaplan, Truth-Telling, RE, or PvT). Having borrowed the scenario from PS, however, we now have a meta game with 5 strategies for the bidders and 4 for the seller, which we analyze using coupled replicator equations.

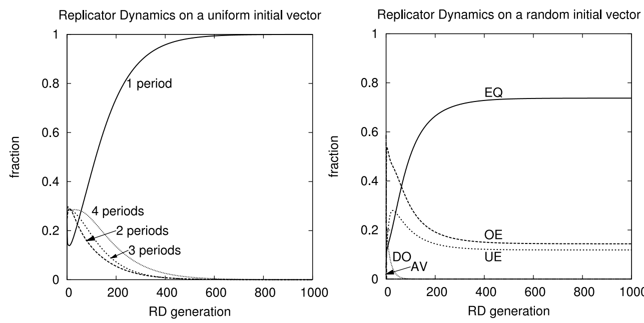

Figure 2: Coupled replicator dynamics.

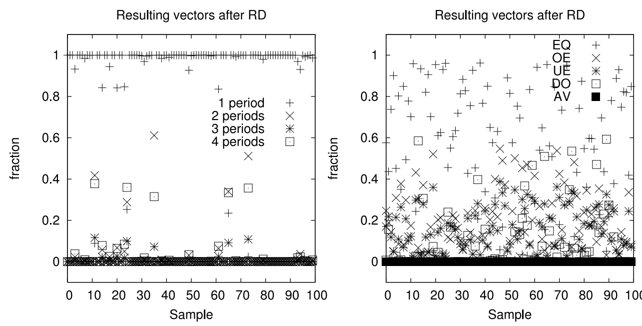

Figure 3: Vectors after RD.

4. EXPERIMENTS

We performed experiments in which the bidding strategy distribution is evolved using RD based on the expected average utility generated by each bidding strategy, weighted by the seller's probabilities of choosing each number of periods. At the same time, another, connected replicator equation evolves a probability distribution over the range of choices for the number of periods of the auction, based on different choices' expected revenues (as in Table 1). The seller's initial probability distribution is always uniform.

Figure 2 shows an example where the initial bidding strategy distribution is drawn from the 4-simplex uniformly at random, see [5]: draw 5 samples $y_i \sim U[0, 1]$, set $y_i \leftarrow -\ln y_i$ and renormalize. We ran experiments using $10,000$ such randomly generated strategy distributions. In Figure 3, we show the period- and strategy-distribution vectors resulting after $1,000$ generations of replication in the first 100 samples. Each of the 100 samples shows 4 (5) values in the graph on the left (right), one for each choice, with the values for each sample adding to 1. The tendency for the number of periods (left) is clear: 1 period results as the most potent value, defeating higher values in most evolutionary circumstances. As we have seen (cf. Table 1), this often leaves

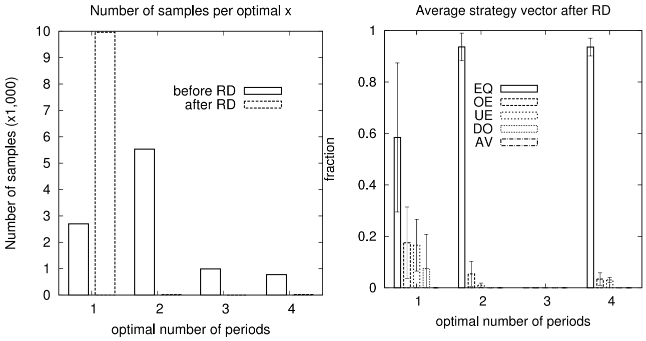

Figure 4: Vectors after RD.

little difference between the first 4 bidding strategies' average utilities. This means the bidding strategy-RD has trouble differentiating among those strategies (right), although the AV strategy is always eradicated quickly, and the EQ strategy is most often the best.

In Figure 4 (left) we plot the number of samples in which each of the numbers of periods maximizes revenue—before and after RD. Choosing 1 or 2 periods generates the highest performance in 82% of all randomly generated distributions (before RD). After RD, choosing 1 period is optimal virtually always: 2 (4) periods is optimal in only 19 (21) out of $10,000$ samples. The graph on the right shows the average strategy distributions occuring in those cases, showing a high prominence of the EQ bidding strategy, although when 1 period is optimal, the high standard deviation (errorbar) suggests maybe we're averaging over several qualitatively distinct attractors. Further research is in order here.

5. CONCLUSIONS AND FUTURE WORK

Notwithstanding the obvious generality of our approach, the usefulness in the PS scenario is limited by our working with the original set of bidding strategies, which have their problems, as described above. Of course, the efforts of PS were not focused on the bidder side of the game, but rather on the seller's adaptive design, and their method clearly works for the selected set of bidding strategies. On the other hand, our method shows one can go further, and even without costly learning.

The choice of 3 heuristic strategies in much previous work not only provides an aesthetically pleasing way of presenting results (by plotting the dynamics' direction fields in the 2-simplex, and shading the simplex' interior with the magnitude of the $\dot{x}_i$'s), but this also gives one a convenient way of finding starting points for the search for Nash equilibria in these games (by the Folk theorem of Evolutionary Game Theory, see [6]). Working in the 3- and 4-simplices, we have less insight in where the Nash equilibria might be located. Still, in future work, we would like to investigate this further—though probably in a game involving a different set of more 'interesting' bidding strategies than the current ones we borrowed from PS—for example by using the amoeba algorithm (see [20]) for finding equilibria. The multi-period auction set up by PS is not so widely studied in the literature. If auctions have multiple rounds, then these usually serve to allow bidders to update their (standing) bids based on newly available information, for example in iterative combinatorial auctions [9].

6. ACKNOWLEDGMENTS

This research has been partially funded by the Netherlands Organisation for Scientific Research (NWO) under FOCUS/BRICKS grant number 642.066.604. We thank the anonymous reviewers for helpful comments.

7. REFERENCES

[1] T. Börgers and R. Sarin. Learning through reinforcement and replicator dynamics. *Journal of Economic Theory*, 77, 1997.

[2] L.-C. Chen and X. Deng. Settling the complexity of two-player Nash equilibrium. In *FOCS*. IEEE, 2006.

[3] D. Cliff. Explorations in evolutionary design of online auction market mechanisms. *ECR&A*, 2, 2003.

[4] V. Conitzer and T. W. Sandholm. Complexity of mechanism design. In *UAI*, 2002.

[5] L. Devroye. *Non–Uniform Random Variate Generation*. Springer, 1986.

[6] J. Hofbauer and K. Sigmund. Evolutionary game dynamics. *Bulletin of the AMS*, 40, 2003.

[7] V. Krishna. *Auction Theory*. Academic Press, 2002.

[8] M. Kunigami and T. Terano. Analyzing and taming collective learning of a multiagent system with connected replicator dynamics. In *JSAI 2003/2004*. Springer, 2007.

[9] B. Mendoza and J. M. Vidal. Bidding algorithms for a distributed combinatorial auction. In *AAMAS*, 2007.

[10] C. H. Papadimitriou. Algorithms, games and the internet. In *STOC*. ACM, 2001.

[11] D. Pardoe and P. Stone. Developing adaptive auction mechanisms. *ACM SIGecom Exchanges*, 5, 2005.

[12] D. Pardoe, P. Stone, M. Saar-Tsechansky, and K. Tomak. Adaptive mechanism design: A metalearning approach. In *ICEC*. ACM, 2006.

[13] S. Phelps, M. Marcinkiewicz, S. Parsons, and P. J. McBurney. A novel method for automatic strategy acquisition in N-player non-zero-sum games. In *AAMAS*, 2006.

[14] S. Phelps, P. J. McBurney, S. Parsons, and E. Sklar. Co-evolutionary auction mechanism design: A preliminary report. In *AMEC IV*. Springer, 2002.

[15] Y. Sato, E. Akiyama, and J. D. Farmer. Chaos in learning a simple two-person game. *PNAS*, 99, 2002.

[16] Y. Sato and J. P. Crutchfield. Coupled replicator equations for the dynamics of learning in multiagent systems. *Phys. Rev. E*, 67, 2003.

[17] R. S. Sutton and A. G. Barto. *Reinforcement Learning*. MIT Press, 1998.

[18] K. Tuyls, A. Nowé, T. Lenaerts, and B. Manderick. An evolutionary game theoretic perspective on learning in multi-agent systems. *Synthese*, 139, 2004.

[19] P. Vytelingum, D. Cliff, and N. R. Jennings. Evolutionary stability of behavioral types in the continuous double auction. In *TADA/AMEC 2007*. Springer, 2007.

[20] W. E. Walsh, R. Das, G. Tesauro, and J. O. Kephart. Analyzing complex strategic interactions in multi-agent systems. In *AAAI GTDT Workshop*, 2002.

[21] M. P. Wellman, J. K. MacKie-Mason, D. M. Reeves, and S. Swaminathan. Exploring bidding strategies for market-based scheduling. In *ACM EC*. ACM, 2003.