

Auction Analysis by Normal Form Game Approximation

Michael Kaisers, Karl Tuyls
Eindhoven University of Technology
Eindhoven, The Netherlands
{m.kaisers, k.p.tuyls}@tue.nl

Frank Thuijsman
Maastricht University
Maastricht, The Netherlands
frank@micc.unimaas.nl

Simon Parsons
Brooklyn College, City University of New York
Brooklyn NY 11210 USA
parsons@sci.brooklyn.cuny.edu

Abstract

Auctions are pervasive in today's society and provide a variety of real markets. This article facilitates a strategic choice between a set of available trading strategies by introducing a methodology to approximate heuristic payoff tables by normal form games. An example from the auction domain is transformed by this means and an evolutionary game theory analysis is applied subsequently. The information loss in the normal form approximation is shown to be reasonably small such that the concise normal form representation can be leveraged in order to make strategic decisions in auctions. In particular, a mix of trading strategies that guarantees a certain profit is computed and further applications are indicated.

1. Introduction

Auctions are deployed in a variety of real markets to foster highly efficient trading. They range from consumer-to-consumer markets like eBay to government-to-business auctions for mineral rights [3]. Buyers and sellers place offers to indicate their intention to trade at a certain price. The here considered *clearing house auction* proceeds in rounds and polls offers from each trader each round. When all offers are collected, an equilibrium price is established based on the available offers such that demand meets supply at this price. It is set to the average of the two offers that define the range of possible equilibrium prices, i.e. the lowest bid and the highest ask that can be matched in the equilibrium. The good to be traded may have a different *private value* for each trader. The difference between the transaction price and the private value of the trading agent determines the agent's *profit*, assuming that buyers will not buy above and sellers will not sell below their private value.

A multitude of trading strategies has been devised to derive the next offer, possibly exploiting the knowledge about offers and transactions that were observed in previous rounds. The experiment of this article considers three

trading strategies: Modified Roth-Erev (MRE), Zero Intelligence Plus (ZIP) and Gjerstad and Dickhaut (GD). A survey of these strategies can be found in [2]. Given a set of available trading strategies, it is of high interest which strategy is *best* in the sense that it yields the highest expected payoff. However, this question cannot be answered in general as the performance of a trading strategy is highly dependent on the competition it faces [5].

A *heuristic payoff table* has been proposed to capture the average profit of each type of trading strategy for all possible mixtures of strategies in the competition [7]. Although the heuristic payoff table provides the basis for analyzing the dynamics in auctions, it is unintuitive and lacks information about the payoffs for strategies that are not yet present in a population. The normal form game on the other hand is more intuitive, e.g. allowing to compute optimal strategies, best replies and Nash equilibria more easily. This article shows how a normal form approximation of a heuristic payoff table can be found using a linear least squares algorithm or linear programming and illustrates the method by an example from the auction domain. The results presented below show a reasonably small error such that the approximation can be leveraged for strategic considerations and an intuitive grasp of the game in auctions.

The remainder of this article is structured as follows: Section 2 introduces the game theoretical background that is required by the methodology presented in Section 3. Subsequently, Section 4 illustrates the method by applying it to an example from the auction domain. The results are discussed in Section 5 which also concludes this article.

2. Game theoretical background

Classical game theory is the mathematical study of strategic conflicts of rational agents. A game comprises a set of players $I = \{1, 2, \dots, n\}$ and a set of available pure strategies $S_i = \{1, 2, \dots, k_i\}$ for each player i , for n and k_i some finite integer. In the context of auctions, each pure strategy corresponds to a trading strategy.

2.1. Normal form games

The players of normal form games are assumed to choose their pure strategies simultaneously and independently and receive a payoff that is dependent on the joint strategy profile $s \in S_1 \times \dots \times S_n$.

The payoff for symmetric two-player normal form games can be described by matrix A , where for any joint strategy (i, j) , A_{ij} denotes the payoff to player one and A_{ji} describes the payoff to player two. Both players seek to maximize their expected payoff and optimal mixed strategies can be derived such that a highest least profit v against any opponent is guaranteed by randomizing over the pure strategies. This profit v is also called the value of a game.

Within a matrix game, the player optimizes his expected payoff against an opponent that plays according to a certain probability distribution. Similarly, he faces a field of traders that are distributed over the strategies in reality. It does not actually matter which opponent plays which strategy but rather how many opponents deploy which strategy. Assuming he encounters a random individual from the competition, his opponent's strategy will be drawn from the distribution of strategies in the competition. The opponent in normal form games therefore resembles the competition in which the agent is situated in reality.

2.2. Heuristic payoff tables

A heuristic payoff table may also be used to capture the payoffs of a symmetric game. However, it requires a finite population of traders such that all possible combinations of strategies can be evaluated. For an individual trader it is only important to know how many of his opponents are playing each of the different strategies. The distribution of n agents on k pure strategies is a combination with repetition, hence a heuristic payoff table requires $\binom{n+k-1}{n}$ rows. Each row yields a *discrete profile* $N = (N_1, \dots, N_k)$ telling exactly how many agents play each strategy.

The payoffs of these discrete profiles can be measured in many domains, e.g. in simulated auctions. However, measurements do not allow to capture the payoff to strategies that are not present, i.e. whenever $N_i = 0$ then $U_i(N)$ is unknown for that discrete profile. Table 1 shows a heuristic payoff table obtained from an auction simulation, indicating unknown payoffs with a dash.

2.3. Replicator dynamics

Evolutionary game theory assumes an infinitely large population of individuals. In this population, each trader may apply one of the available trading strategies. The pure strategy i is played with probability x_i , according to the vector $x = (x_1, \dots, x_k)$. The profit of each trader depends on

x and may be interpreted as Darwinian fitness. From auction to auction, traders may switch their strategies, adopting successful strategies and dropping less successful ones. This adaptation can be modeled by the *replicator dynamics* from evolutionary game theory which are formally connected to reinforcement learning [6]. Let $\pi_i(x)$ denote the expected payoff to strategy i in the population playing according to x . The replicator dynamics formally define the population change over time:

$$\dot{x}_i = x_i \cdot [\pi_i(x) - x\pi_i(x)]$$

For the matrix game A , the expected payoff function $\pi(x) = Ax$ is linear. The heuristic payoff table may define more complex expected payoff functions. Let the set of all discrete profiles be denoted as $\nu = \{(n, 0, \dots, 0), \dots, (0, \dots, 0, n)\}$ and let $\mu_i = \{N \in \nu \mid N_i = 0\}$ be the set of profiles where strategy i is not played. Furthermore, let $\bar{\mu}_i = \{N \in \nu \mid N_i \neq 0\}$ be the complement of μ_i . The expected payoff can be computed from the heuristic payoff table:

$$\pi_i(x) = \frac{\sum_{N \in \bar{\mu}_i} U_i(N) \cdot \binom{n}{N_1, \dots, N_k} \cdot x_1^{N_1} \cdot \dots \cdot x_k^{N_k}}{1 - \sum_{N \in \mu_i} \binom{n}{N_1, \dots, N_k} \cdot x_1^{N_1} \cdot \dots \cdot x_k^{N_k}}$$

The resulting dynamics can be visualized in a force field plot as in Figure 2, where the arrows indicate the direction and strength of change. The replicator dynamics give rise to a dynamical system which may feature repellers and attractors of which the latter are of particular importance to the analysis of asymptotic behavior. Each attractor consumes a certain amount of the strategy space that eventually converges to it — this space is also called the basin of attraction [1]. A good approximation of the game dynamics should have minimal impact on the basins of attraction.

3. Methodology

This section presents the conversion from normal form games to heuristic payoff tables and the newly proposed method for finding a suitable normal form game approximation.

Table 1. Heuristic payoff table of a clearing house auction

N_{ZIP}	N_{MRE}	N_{GD}	U_{ZIP}	U_{MRE}	U_{GD}
6	0	0	99	-	-
5	1	0	97	100	-
	\vdots			\vdots	
0	1	5	-	43	79
0	0	6	-	-	79

From normal form games to heuristic payoff tables

The heuristic payoff table lists all possible discrete profiles with the average payoff of playing against a finite population that mixes accordingly. The payoff vector against the mixed strategy p can be computed from the matrix game M as Mp . Let D be the matrix where each row corresponds to a discrete profile N of n agents and let $P = \frac{1}{n} \cdot D$ map the discrete profiles to probabilities. The matrix

$$U = P \cdot M^T \quad (1)$$

yields the corresponding payoff vectors $U(N)$ as rows. The heuristic payoff table $H = (D, U)$ is the composition of the discrete profiles and the corresponding payoffs.

Unfortunately, Equation (1) cannot simply be solved for M as the values in the heuristic payoff table may be noise-prone due to stochasticity in the experiments and may also feature non-linear dynamics which leads to an over-constrained system of equations. Therefore, it needs to be approximated, e.g. by minimizing the mean squared error or the maximal absolute deviation.

Minimizing mean squared error

A normal form game M that approximates the heuristic payoff table $H = (D, U)$ can be determined incrementally for each row M_i by finding a least mean squared error fit between the i 'th column of U , denoted as U_i , and the reconstructed payoff vector $\tilde{U}_i = P \cdot M_i^T$ from the normal form game, where $P = \frac{1}{n} \cdot D$ as above, by solving the minimization problem:

$$\min_{M_i} \left\| U_i - \tilde{U}_i \right\|_2$$

A standard linear least square fitting algorithm can be used to solve this system for each row and compose the normal form game matrix.

Minimizing maximal absolute deviation

Linear programming optimizes a linear goal function subject to a system of linear inequalities. Using the same definitions of the profile matrix D , the probability matrix P , the game M and the payoff matrix U as above, the following program can be formulated.

$$\begin{aligned} &\text{minimize } \epsilon \\ &\text{variables } \epsilon, M_{ij}, \text{ for } i, j \in \{1, \dots, k\} \\ &\text{subject to } P \cdot M^T \leq U + \epsilon \\ &\quad \quad \quad P \cdot M^T \geq U - \epsilon \end{aligned}$$

Least mean squared error			
	ZIP	MRE	GD
ZIP	97.4	98.8	52.3
MRE	96.8	98.6	42.6
GD	64.8	59.1	83.4

However, this program needs to be transformed to standard notation in order to apply common algorithms from linear programming. For sake of convenience, each row M_i is determined separately. Let $c = (1, 0, \dots, 0)$ and $x = (\epsilon, M_i)$ such that the goal function minimizes epsilon. Furthermore,

$$\text{let } A = \begin{pmatrix} -1 & & P \\ & \vdots & -P \\ -1 & & \end{pmatrix} \text{ and } b = \begin{pmatrix} U_i \\ -U_i \end{pmatrix} \text{ where } U_i$$

is the i 'th column of the payoff matrix. Then, this linear program can be solved in standard notation:

$$\min_x c \cdot x^T \quad \text{subject to } A \cdot x^T \leq b, x \geq 0$$

In order to approximate the heuristic payoff table, we need to solve k linear programs to compute the complete normal form matrix.

4. Experiments

This section presents the experimental setup and results of measuring the information loss in the normal form game approximation of a heuristic payoff table from the auction domain. The heuristic payoff table given in Table 1 is obtained by simulating auctions with the *Java Auction Simulator API* (JASA) [4]. This empirical platform contains the trading strategies ZIP, MRE and GD which were setup with default parameters. The heuristic payoff table is obtained from an average of 2000 iterations of clearing house auctions. On the start of each auction, all traders are initialized without knowledge of previous auctions and with a private value drawn from the same distribution as in [7].

The heuristic payoff table is approximated as described in Section 3, which leads to the normal form game representations given in Figure 1. The replicator dynamics are derived from the heuristic payoff table and the normal form game representations and compared in Figure 2. There is a clear qualitative correspondence of the dynamics that arise from the three models.

A mixed attractor can be found at (0.81, 0.18, 0.0) for the heuristic payoff table, at (1, 0, 0) in least mean squared error fitting and at (0.72, 0.27, 0.0) in minimized maximal absolute deviation. However, there are very similar basins of attraction. The pure attractor at (0, 0, 1) is present in all dynamics and is estimated to consume 26.0% of the strategy space in the heuristic payoff table in comparison to 26.2%

Least maximal absolute deviation			
	ZIP	MRE	GD
ZIP	93.8	102.7	52.9
MRE	94.9	100.0	38.3
GD	66.2	60.5	81.8

Figure 1. Symmetric two-player normal form game approximations of the heuristic payoff table




Figure 2. Comparison of replicator dynamics showing directional and force field plots

and 27.3% in the approximations, based on the analysis of convergence of 1000 trajectories with uniformly sampled starting points. The attractors are evolutionary stable in the normal form game approximations and predict the attractors that are observed in the auction game dynamics with a small error. The maximal absolute deviation is 9.76% and 6.64% while the square root of the mean squared error is 3.03% and 3.42% respectively.

5. Discussion and conclusions

The results show that heuristic payoff tables in the domain of auctions may be approximated by normal form games with a reasonably small error, using a case study of an auction with $n = 6$ agents and $k = 3$ strategies. The full representation of $k^n = 729$ payoff entries was reduced to a heuristic payoff table with $k \binom{n+k-1}{n} = 84$ payoff entries, and then to a normal form game with $k^2 = 9$ payoff entries. Consider the normal form representation of the auction game obtained from minimizing the maximal absolute deviation as given in Figure 1. The optimal strategy $\pi^* = (0.3, 0, 0.7)$ guarantees an expected payoff of at least 73.1 against any combination of GD, ZIP and MRE. This strategy could be used for risk averse trading, given that traders stick to the given set of trading strategies.

The contributions can be summarized as follows: A methodology to approximate heuristic payoff tables by normal form games has been introduced. This smaller representation is more intuitive, computationally less expensive to analyze and adds information about missing payoffs in the heuristic payoff table. It can even be constructed from partial heuristic payoff tables. Rather than merely participating myopically, a rational agent can now inspect the game strategically. With further refinements, this method

may become the basis of a more general framework for the analysis of complex multi-agent games.

Acknowledgments

The authors wish to express their gratitude to Jinzhong Niu and Steve Phelps for their support and helpful discussions. This research was partially sponsored by a TopTalent 2008 grant of the Netherlands Organisation for Scientific Research (NWO).

References

- [1] M. W. Hirsch, S. Smale, and R. Devaney. *Differential Equations, Dynamical Systems, and an Introduction to Chaos*. Academic Press, 2004.
- [2] S. Parsons, M. Marcinkiewicz, J. Niu, and S. Phelps. Everything you wanted to know about double auctions, but were afraid to (bid or) ask. Technical report, Brooklyn College, City University of New York, 2900 Bedford Avenue, Brooklyn, NY 11210, USA, 2006.
- [3] S. Parsons, J. Rodriguez-Aguilar, and M. Klein. A bluffer's guide to auctions. Technical report, Center for Coordination Science, MIT, 2004.
- [4] S. Phelps. Java auction simulator api. <http://www.csc.liv.ac.uk/~sphelps/jasa/>, 2005.
- [5] J. Rust, J. Miller, and R. Palmer. Behavior of trading automata in a computerized double auction market. In *The Double Auction Market: Institutions, Theories, and Evidence*. Addison-Wesley, 1993.
- [6] K. Tuyls and S. Parsons. What evolutionary game theory tells us about multiagent learning. *Artificial Intelligence*, 171(7):406–416, 2007.
- [7] W. E. Walsh, R. Das, G. Tesauro, and J. O. Kephart. Analyzing complex strategic interactions in multi-agent systems. In *Proceedings of the Workshop on Game Theoretic and Decision Theoretic Agents*, pages 109–118, 2002.