

Discovering the Game in Auctions

Michael Kaisers^{ab} Karl Tuyls^a Frank Thuijsman^b Simon Parsons^c

^a *Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven*

^b *Maastricht University, P.O. Box 616, 6200 MD Maastricht*

^c *Brooklyn College, City University of New York, Brooklyn NY 11210 USA*

Abstract

Auctions are pervasive in today's society. They provide a variety of markets, ranging from consumer-to-consumer online auctions to government-to-business auctions for telecommunications spectrum licenses. Starting from a set of trading strategies, this article enables a strategic choice by introducing the use of linear programming as a methodology to approximate heuristic payoff tables by normal form games. This method is evaluated on data from auction simulation by applying an evolutionary game theory analysis. 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 against any population of traders is computed and further applications are indicated.

Keywords: Evolutionary game theory, Auction theory, Multi-agent games

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 via business-to-business stock exchanges to government-to-business auctions for mineral rights or government licenses for the telecommunications spectrum [7, 9]. Furthermore, auction mechanisms have been transferred successfully to solve other resource allocation problems, e.g. in the domain of efficient internet traffic routing [12]. Despite this diversity, even single markets of enormous scale may have profound impact on society, e.g. the New York Stock Exchange reported a trading volume of 11,060 billion USD in 2000 [3]. This motivates researching how to run auctions and how to extract profit by trading within them as auctions are pervasive in today's society.

The traders that participate in an auction agree to subject to a set of market rules in order to exchange goods for money. Within the scope of this article only commodity markets are considered, i.e. a single type of an abstract good is traded. Each trader is assumed to have a *private valuation* of the good which is only known to himself. In double auctions, buyers and sellers place offers to indicate their intention to trade at a certain price [9]. 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 commonly 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. Each buyer with an offer above that price is matched with a seller having an offer below that price. The *profit* of a transaction can be computed as the difference between the transaction price and the private value, assuming that buyers will not buy above their private value 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 most trivial one is *Truth Telling* (TT) which just reveals the private value by placing offers exactly at that value. The experiment of this article instead considers three more sophisticated trading strategies which give raise to non-trivial evolutionary dynamics. Roth and Erev devised a reinforcement learning model of human trading behavior in [2] which is modified to perform in a clearing house auction as *Modified Roth-Erev* (MRE) [8]. MRE is

evaluated in competition to *Gjerstad and Dickhaut* (GD) and *Zero Intelligence Plus* (ZIP). GD maximizes the expected profit by computing the profit and probability of leading to a transaction for a set of relevant prices [5]. ZIP places stochastic bids within a certain profit margin, which is lowered when a more competitive offer was rejected and increased when a less competitive offer was accepted [1]. It can be noted that most trading strategies are adaptive to the progress of the auction, the internal processes however are outside the scope of this article.

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 [13]. Let us therefore assume an auction that is populated by traders, each deploying one of the trading strategies above. The profit of each trader is dependent on the overall mix of strategies and traders may choose to change their strategy in the course of time. A *heuristic payoff table* is proposed in [14] and is adopted by several authors to capture the average profit of each type of trading strategy under all possible mixtures of strategies in a finite population [6, 10]. This table is a first step towards revealing the dynamics of adopted trading strategies in auctions and can for example be used to analyze which trading strategy yields the highest potential for improvements [10].

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. However, exactly these payoffs would provide information about whether it is profitable or not to be the first one to adopt this strategy. The normal form game on the other hand enables an individual trader to calculate his expected profit for each of his possible choices against any mix of strategies he faces. It is more intuitive and allows inspecting the strategic situation with means from game theory, e.g. allowing to compute optimal strategies, best replies and Nash equilibria. This suggests the question whether a heuristic payoff table can be approximated by a normal form game in order to open up these opportunities. Answering this question affirmatively, this article demonstrates how an approximation can be found using linear programming, a common technique to optimize a linear goal function subject to a set of linear inequalities. The methodology is illustrated by approximating a heuristic payoff table from the auction domain and 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 and presents the resulting performance. These results are discussed in Section 5 and the paper is concluded with future directions in Section 6.

2 Game theoretical background

Classical game theory is the mathematical study of strategic conflicts of rational agents. Each individual i chooses a pure strategy s_i from the set of available strategies S_i and has a preference relation over all possible outcomes. The players are assumed to choose their actions simultaneously and independently. This implies that the preference relation can be captured by a numerical *payoff function* τ_i which is public knowledge and assigns a value of desirability to each possible joint strategy $s = (s_1, \dots, s_n)$, where n is the total number of agents.

$$\tau_i : S_1 \times \dots \times S_n \rightarrow \mathbb{R}$$

In the context of auctions, each pure strategy corresponds to a trading strategy and the preference relation is proportional to the profit that an agent can make given the set of opponents' trading strategies. This section introduces two different means to capture payoff functions for multi-agent games and auctions in particular: The normal form game and the heuristic payoff table. Furthermore, advantages and disadvantages of both representations are discussed. Subsequently, the evolutionary perspective is compared to the classical view from game theory and the concepts of replicator dynamics and basins of attraction are presented.

2.1 Normal form games

A normal form game commonly describes the payoff to each agent in matrix notation. The matrix given in Figure 1 describes a symmetric two-player normal form game. The first player may choose a row r , the second player chooses a column c and the joint choice (r, c) determines the payoff which the matrix gives for the first player. However, the payoff matrix for the second player equals the transposed of the first player's

	<i>Rock</i>	<i>Paper</i>	<i>Scissors</i>
<i>Rock</i>	0	-1	1
<i>Paper</i>	1	0	-1
<i>Scissors</i>	-1	1	0

Figure 1: Payoff matrix for the symmetric two-player normal form game Rock-Paper-Scissors.

payoffs in symmetric games. Hence, it can be derived from the same table by consulting the entry (c, r) . Given k strategies, the matrix yields k^2 entries which compares favorably to the size of heuristic payoff tables given below. Both players seek to maximize their expected payoff; optimal strategies and the value of this game are derived in Section 5.2.

Within a normal form game, the player optimizes his expected payoff against an opponent that mixes 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. The opponent in normal form games therefore resembles the population 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 game [14]. However, it requires a finite population of traders such that all possible combinations of strategies can be evaluated. If each agent $i \in \{1, 2, \dots, n\}$ has to choose a strategy $s_i \in \{1, 2, \dots, k\}$, this leads to a joint strategy (s_1, \dots, s_n) . However, for an individual trader it is only important to know how many of his opponents are playing each of the different strategies. So, given (s_1, \dots, s_n) the individual trader could derive that there are N_1 agents playing strategy 1, N_2 agents playing strategy 2, etc.. This would yield a *discrete profile* $N = (N_1, \dots, N_k)$ telling exactly how many agents play each strategy. The average profit for playing a strategy can then be denoted by a payoff vector $U(N) = (U_1(N), \dots, U_k(N))$ indicating that strategy $s \in \{1, 2, \dots, k\}$ would yield an average payoff of $U_s(N)$ for the discrete profile N . The distribution of n agents on k pure strategies is a combination with repetition, hence the number of profiles in a heuristic payoff table is given by:

$$\binom{n + k - 1}{n}$$

The payoffs of these discrete profiles can be measured in many domains, e.g. in auctions. However, measurements do not allow to capture the payoff to strategies that are not present, i.e. whenever $N_s = 0$ then $U_s(N)$ is unknown for that discrete profile. Table 1 shows the heuristic payoff table obtained from the experiments described in Section 4, indicating unknown payoffs with a dash.

The heuristic payoff table only delivers a snapshot of the game payoffs under fixed discrete profiles and does not yet assume adaptive agents. It is a merely descriptive representation of the conditions and the game dynamics need to be deduced using models on top of that. One opportunity to do so is the approximation of

Table 1: The first six rows give the transposed heuristic payoff table of a clearing house auction with 6 agents and the three strategies ZIP, MRE and GD. Each column gives a discrete profile N over the trading strategies and the corresponding payoff vector $U(N)$. Below, the deviation for the reconstructed payoff table from the normal form game representation is given which is discussed in Section 4. It features a maximal absolute deviation of 6.64% and a root mean squared error of 2.96%.

N_{ZIP}	6	5	5	4	4	4	3	3	3	3	2	2	2	2	2	1	1	1	1	1	1	0	0	0	0	0	0	0
N_{MRE}	0	1	0	2	1	0	3	2	1	0	4	3	2	1	0	5	4	3	2	1	0	6	5	4	3	2	1	0
N_{GD}	0	0	1	0	1	2	0	1	2	3	0	1	2	3	4	0	1	2	3	4	5	0	1	2	3	4	5	6
U_{ZIP}	99	97	89	96	90	85	97	87	85	76	97	91	84	78	62	97	93	86	73	73	56	-	-	-	-	-	-	-
U_{MRE}	-	100	-	94	88	-	92	90	80	-	96	91	83	70	-	97	89	84	71	57	-	94	91	84	75	65	43	-
U_{GD}	-	-	69	-	65	69	-	64	73	73	-	66	67	76	80	-	62	69	75	77	80	-	62	67	71	76	79	79
ΔU_{ZIP}	6	3	3	1	3	6	0	-2	4	3	-2	1	2	4	-4	-3	2	2	-3	6	-3	-	-	-	-	-	-	-
ΔU_{MRE}	-	5	-	-1	3	-	-4	4	4	-	-1	4	7	3	-	-1	1	7	4	0	-	-4	2	6	6	6	-5	-
ΔU_{GD}	-	-	1	-	-2	-1	-	-2	3	0	-	1	-1	4	4	-	-2	1	4	3	1	-	-1	0	1	2	2	-2

the heuristic payoff table by a normal form game which allows to draw on the classical means from game theory to derive strategic choices (see Sections 3.2 and 5.2).

2.3 Replicator dynamics

Replicator dynamics describe game dynamics from an evolutionary perspective. Evolutionary game theory assumes an infinitely large population of individuals that choose their pure strategy according to some probability distribution. It assumes this population to evolve such that successful strategies with higher payoffs grow while less successful ones decay and it suggests to analyze the asymptotic behavior [4].

Evolutionary game theory takes a rather descriptive perspective replacing hyper-rationality from classical game theory by the concept of natural selection from biology. The evolutionary pressure by natural selection can be modeled by the replicator equations. Symmetric games only require the single-population replicator dynamics that define the growth of a strategy proportional to the fraction of the population that already uses this strategy and the difference between the payoff to this strategy and the average payoff.

The game dynamics for the normal form game with payoff matrix A can be calculated for strategy i given that the opponent mixes over the pure strategies according to the probability vector x :

$$\dot{x}_i = x_i \cdot [(Ax)_i - xAx] \quad (1)$$

It is also possible to construct the replicator dynamics from the heuristic payoff table immediately. To achieve this, the probability of each discrete profile for a finite population where each agent independently chooses its strategy according to some probability distribution needs to be calculated. This is a multinomial process for which the probability of the discrete profile N given the mixed strategy p can be computed as:

$$Pr(N|p) = \binom{n}{N_1, \dots, N_k} \cdot p_1^{N_1} \cdot \dots \cdot p_k^{N_k} \quad (2)$$

The payoff for each strategy can then be computed as the weighted average over the payoffs received in all profiles. However, a correction term is required if the payoffs for non-occurring strategies are unknown in the heuristic payoff table.

$$U_{average,i} = \frac{\sum_N Pr(N|p) \cdot U_i(N)}{1 - Pr(unknown|i)} \quad (3)$$

The resulting dynamics can be visualized in a vector field plot as in Figure 2 where the arrows indicate the direction of change and the length of an arrow is proportional to $|\dot{x}|$.

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. Assuming that an evolutionary process may start uniformly at any point in the strategy space, the size of the basin of attraction may be used to estimate the practical importance of an attractor. This can be achieved by uniform sampling of the strategy space and analysis of trajectory convergence, or under the assumption of continuous dynamics by grid sampling of the strategy space and a graph labeling algorithm. The latter applied to the given replicator dynamics leads to the basins of attraction depicted in Figure 4. A good approximation of the game dynamics should naturally have minimal impact on the basins of attraction.

3 Methodology

The conversion of a normal form game to a heuristic payoff table is straight forward. This gives raise to the question whether the inverse is also possible. However, the inverse transformation is over constrained and a heuristic payoff table can only be approximated by a normal form game. This section presents the newly proposed method to use linear programming for finding a suitable normal form game approximation.

3.1 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 payoff matrix A as A_p . Given a matrix D where each row corresponds to a discrete profile N we can compute the matrix U that then yields the corresponding payoff vectors $U(N)$ as rows:

$$U = \frac{1}{n} \cdot D \cdot A^T \quad (4)$$

The heuristic payoff table is the compound of the discrete profiles D and the corresponding average payoffs given by U .

3.2 From heuristic payoff tables to normal form games

While the previous section has shown the transition from a normal form game to a heuristic payoff table, this section will reverse this step. However, the equation cannot simply be reversed 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. Although this leads to an over-constrained system of equations, an approximation with minimal maximum absolute deviation can be found using linear programming.

Linear programming optimizes a linear goal function subject to a system of linear inequalities. The following program can be formulated in order to approximate a heuristic payoff table, where D is a matrix with all discrete profiles as rows, U is a matrix that yields the payoff vectors corresponding to D , $P = \frac{1}{n} \cdot D$ maps the discrete profiles to probabilities and M is the payoff matrix of the game that we search:

$$\begin{aligned} \min. \quad & \epsilon \\ \text{s.t.} \quad & P \cdot M^T \leq U + \epsilon \\ & P \cdot M^T \geq U - \epsilon \end{aligned} \quad (5)$$

However, this notation needs to be transformed to standard notation in order to apply common algorithms

from the linear programming literature. Let $c = (1, 0, 0, \dots, 0)$, $x = (\epsilon, M_i)$, $A = \begin{pmatrix} -1 & P \\ \vdots & -P \\ -1 & \end{pmatrix}$

and $b = \begin{pmatrix} U_i \\ -U_i \end{pmatrix}$ where M_i is the i 'th row of the payoff matrix to find and U_i is the i 'th column of the payoff matrix. Then, this linear program can be solved in standard notation:

$$\begin{aligned} \min. \quad & c \cdot x^T \\ \text{s.t.} \quad & A \cdot x^T \leq b \end{aligned} \quad (6)$$

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) [11]. This empirical platform contains the trading strategies ZIP, MRE and GD which were set up with the following parameters, according to [1, 5, 8]: ZIP uses a learning rate of 0.3, a momentum of 0.05 and a JASA specific scaling of 0.2, MRE chooses between 40 discrete prices using a recency parameter of 0.1, an exploration of 0.2 and scaling of 9 and GD evaluates prices up to 360.

The heuristic payoff table is obtained from an average over 2000 iterations of clearing house auctions with 6 traders. 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 [14], i.e. an integer lower bound b is drawn uniformly from $[61, 160]$ and the upper bound from $[b + 60, b + 209]$ for each buyer. The sellers' private values are initialized similarly. These private values then remain fixed over the course of the auction which runs 300 rounds on each of 5 trading days where each trader is entitled to trade one item per day.

The heuristic payoff table is approximated with a linear program as described in Section 3.2 which leads to the normal form game representation given in Figure 3. This normal form representation is transformed




Figure 2: This figure shows the original replicator dynamics from the heuristic payoff table (left) and those from the normal form game approximation (right) in the clearing house auction.

	<i>ZIP</i>	<i>MRE</i>	<i>GD</i>
<i>ZIP</i>	93.8	102.7	52.9
<i>MRE</i>	94.9	100.0	38.3
<i>GD</i>	66.2	60.5	81.8

Figure 3: The symmetric two-player normal form game approximation of the heuristic payoff table for a clearing house auction with the three strategies ZIP, MRE and GD.

back into a heuristic payoff table as described in Section 3.1 and compared to the original table, leading to the differences indicated by ΔU_s in Table 1. Furthermore, the replicator dynamics are derived from both tables and compared in Figure 2. A more intuitive interpretation is given in Figure 4 which shows the basins of attraction that arise from the replicator dynamics.

Figure 2 shows that the differences in the replicator dynamics are very small and can hardly be observed by inspection of the vector field plots. Therefore, Figure 4 visualizes the basins of attraction which show a clear qualitative correspondence of the dynamics in the observed heuristic payoff table and the reconstruction from the normal form approximation. Only one mixed attractor moves slightly but this hardly impacts its basin of attraction. In the context of evolutionary game theory, evolutionary stable strategies provide a concept to find stable solutions in normal form games. The attractor $(0, 0, 1)$ which corresponds to a pure population of GD and the attractor $(0.7, 0.3, 0)$ are evolutionary stable in the normal form game and predict the attractors that are observed in the auction game dynamics.

5 Discussion

The results show that heuristic payoff tables in the domain of auctions may be approximated by normal form games with a reasonably small error. However, this case study is rather a proof of concept and does not necessarily generalize. Therefore, this section starts with a discussion of the limitations of this approach. Eventually, leveraging the newly gained insights for strategic choice is illustrated.




Figure 4: This figure shows the basins of attraction of the heuristic payoff table (left) and of the normal form game approximation (right) in the clearing house auction.

N_{ZIP}	5	4	4	3	3	3	2	2	2	2	1	1	1	1	1
N_{MRE}	0	1	0	2	1	0	3	2	1	0	4	3	2	1	0
N_{GD}	1	1	2	1	2	3	1	2	3	4	1	2	3	4	5
$U_{average}$	74.9	72.3	74.1	70.7	76.4	74.0	73.9	72.5	76.8	74.5	71.3	74.1	74.5	76.1	72.5

Table 2: The average payoffs of the optimal strategy $\pi^* = (0.3, 0, 0.7)$ against all discrete profiles for which the payoffs for the strategies ZIP and GD are known. The lower bound is given by 70.7 against the profile $N_{min} = (3, 2, 1)$ and the average is 73.9.

5.1 Limitations of the approximation

The proposed approach can be applied to any number of strategies. However, the approximation of heuristic payoff tables by normal form games imposes a less complex model on the data, which may be an oversimplification for the actual dynamics. Consequently, the precision of the approximation is likely to deteriorate when the number of trading strategies to choose from is increased.

5.2 Strategic choice

Consider the normal form representation of the auction game as given in Figure 3. It is possible to derive an *optimal strategy* that gives a lower bound on the profit that can be guaranteed even if nothing is known about the opponents. This profit is also known as the matrix game value and can be determined with standard algorithms from game theory. It equals 73.1 for this example and can be guaranteed by the optimal strategy $\pi^* = (0.3, 0, 0.7)$. This means that a rational trader who is playing ZIP with probability 0.3 and GD with probability 0.7 will get an expected payoff of at least 73.1 against any opponent that mixes between ZIP, MRE and GD. For any other probability distribution than π^* , he may encounter an opponent that gives him a lower expected payoff¹.

In order to validate these results, the optimal strategy may be applied to compute the average payoff against the distributions given in the heuristic payoff table. However, we need to restrict the consideration to those profiles for which all trading strategies that are used in the optimal strategy are present. Only then, the payoffs are known and the average can be computed accurately from the heuristic payoff table. These average profits are given in Table 2 for all these profiles, resulting in a minimum of 70.7 which is close to the approximated 73.1. Furthermore, the average overall profit against a uniformly distributed population is as high as 73.9. Calculation of the matrix game value and the optimal strategy can be achieved by linear programming since the agent simply wants to choose his probabilities in such a way that the minimal payoff over all columns is maximized, i.e. he is maximizing a linear function with respect to some linear inequalities.

If an agent knew the current actual mix of trading strategies in the population he faces, he could even make more than with the optimal strategy because the optimal strategy is based upon a worst case analysis. Under the assumption that all agents have the same goal and the payoff matrix is known to all other agents as well, one can further assume that all other agents apply a similar reasoning and arrive at the same probability distribution. Therefore, a symmetric evolutionary equilibrium is more interesting in this context, especially if we want to understand how the distribution of traders will look like in the long run.

6 Conclusions

This article has modeled trading in auctions by considering a population of traders that repeatedly participate in an auction. A set of trading strategies is made available to the agents who make their choice according to the relative profit of these strategies.

The contributions can be summarized as follows: A methodology to approximate heuristic payoff tables by normal form games has been introduced. This smaller game representation is easier to analyze and fills in a gap of missing payoffs in the blind spots of the heuristic payoff table. Rather than merely participating myopically, a rational agent can now inspect the game strategically and reasoning from game theory can be applied.

¹Note that the optimal strategy is not symmetric and therefore does not appear in the replicator dynamics.

The progress that has been made gives room for interesting opportunities. This approach needs to be tested on other auctions and domains, possibly applying it to higher dimensions as it is general in the number of strategies. This would allow inspecting the game in auctions that yield more than 3 strategies which cannot be visualized easily. Another direction extends the game theoretic analysis, e.g. investigate symmetric equilibria analytically in the normal form game. Furthermore, we aim to argue for the described approach on a more theoretical level and look for structure in the deviation from the linear model, in particular where and why qualitative changes occur.

7 Acknowledgments

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

References

- [1] Dave Cliff and Janet Bruten. Minimal-intelligence agents for bargaining behaviours in market-based environments. Technical report, Hewlett-Packard Research Laboratories, 1997.
- [2] Ido Erev and Alvin E. Roth. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *The American Economic Review*, 88(4):848–881, 1998.
- [3] New York Stock Exchange. Stock market activity. Technical report, New York Stock Exchange, 2000. available at http://www.nyse.com/...pdfs/02_STOCKMARKETACTIVITY.pdf.
- [4] H. Gintis. *Game Theory Evolving*. Princeton University Press, 2000.
- [5] Steven Gjerstad and John Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22(1):1–29, January 1998.
- [6] Tomas B. Klos and Gerrit Jan van Ahee. Evolutionary dynamics for designing multi-period auctions (short paper). In *Proceedings 7th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2008)*. IFAAMAS, 2008.
- [7] John McMillan. Selling spectrum rights. *Journal of Economic Perspectives*, 8(3):145–162, Summer 1994.
- [8] 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.
- [9] Simon Parsons, Juan Rodriguez-Aguilar, and Mark Klein. A bluffer’s guide to auctions. Technical report, Center for Coordination Science, MIT, 2004.
- [10] S. Phelps, M. Marcinkiewicz, and S. Parsons. A novel method for automatic strategy acquisition in n-player non-zero-sum games. In *AAMAS ’06: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pages 705–712, Hakodate, Japan, 2006. ACM.
- [11] Steve Phelps. Java auction simulator api. <http://www.csc.liv.ac.uk/sphelps/jasa/>, 2005.
- [12] Tim Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.
- [13] J. Rust, J. Miller, and R. Palmer. Behavior of trading automata in a computerized double auction market. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories, and Evidence*. Addison-Wesley, 1993.
- [14] William E. Walsh, Rajarshi Das, Gerald Tesauro, and Jeffrey O. Kephart. Analyzing complex strategic interactions in multi-agent systems. In Piotr Gmytrasiewicz and Simon Parsons, editors, *Proceedings of the Workshop on Game Theoretic and Decision Theoretic Agents*, pages 109–118, 2002.