Approximating Mixed Nash Equilibria using Smooth Fictitious Play in Simultaneous Auctions

(Short Paper)

Enrico H. Gerding[*] Zinovi Rabinovich[*] Andrew Byde[*,†] Edith Elkind[*]
Nicholas R. Jennings[*]
{eg,zr,ab06v,ee,nrj}@ecs.soton.ac.uk
[*]School of Electronics and Computer Science, University of Southampton, UK
[†]Hewlett-Packard Laboratories, Bristol, UK.

ABSTRACT

We investigate equilibrium strategies for bidding agents that participate in multiple, simultaneous second-price auctions with perfect substitutes. For this setting, previous research has shown that it is a best response for a bidder to participate in as many such auctions as there are available, provided that other bidders only participate in a single auction. In contrast, in this paper we consider equilibrium behaviour where all bidders participate in multiple auctions. For this new setting we consider mixed-strategy Nash equilibria where bidders can bid high in one auction and low in all others. By discretising the bid space, we are able to use smooth fictitious play to compute approximate solutions. Specifically, we find that the results do indeed converge to ϵ-Nash mixed equilibria and, therefore, we are able to locate equilibrium strategies in such complex games where no known solutions previously existed.

1. INTRODUCTION

The rapid increase of online auctions such as eBay, QXL, and Yahoo! has spawned considerable research in the field of auctions and automated bidding agents. In such auctions we increasingly observe different sellers offering similar or even identical goods and services at the same time. In eBay alone, for example, the Nintendo Wii game console has nearly 2000 listings at the time of writing, of which over 1500 are proper auctions. In addition to the web, such auctions are also considered a key approach to achieve effective allocation of tasks and resources within a number of research areas of multi-agent systems, including Grid computing and multi-robot coordination. Against this background, it is important to develop intelligent agents that are able to bid effectively in such auctions. In particular, this paper considers bidding strategies when multiple auctions selling substitutable goods are held simultaneously. Whereas most of the previous research in this domain focuses on best-response or heuristic strategies, here we extend this research by considering equilibrium outcomes when several agents optimise their utility by participating in multiple auctions. To this end, we compute the equilibrium using *fictitious play*, a game-theoretic learning algorithm that optimises behaviour based on the opponents' history of play. This algorithm has the property that, *if* the strategies converge, in the limit these strategies are a Nash equilibrium solution [4, Ch 2, Prop 2.1]. Specifically, here we apply an approach called smoothed or cautious fictitious play which is able to converge to approximate, also called ϵ-Nash mixed, equilibria. Formally, an equilibrium is ϵ-Nash if any single agent cannot gain more than ϵ by deviating from it.

To date, much of the existing research dealing with multiple simultaneous auctions typically assumes that bidders choose one of them and then bid optimally in that auction. Previous research has shown, however, that if all opponents follow this strategy and bid in a single auction, and given that the auctions do not have reserve prices, it is a best response to bid in *all* available auctions [5]. Here we extend this work by investigating equilibrium strategies where all the agents participate in multiple auctions. Finding an equilibrium outcome in this setting is a challenging problem, however, since no closed-form solution exists even for the best response case, and finding the equilibrium by brute-force search is computationally intractable (especially when considering mixed strategies). As a result, this setting has received very little attention in the literature and the work that exists operates in very limited cases.

In this case, the seminal paper by Engelbrecht-Wiggans and Weber [2] provides one of the starting points for the game-theoretic analysis of markets where buyers have substitutable goods. They derive a mixed Nash equilibrium for the special case where the number of buyers is large. Moreover, they assume that bidders have the same valuations and not all bidders can bid in all auctions. Our analysis, on the other hand, does not make these assumptions. Following this, [6] studied the case of simultaneous auctions with complementary goods. The setting provided in [6] is further extended to the case of common values in [8]. However, neither of these works extend easily to the case of substitutable goods which we consider. This case is studied in [9], but the scenario considered is restricted to three sellers and two bidders and with each bidder having the same value (and thereby knowing the value of other bidders). The space of symmetric mixed equilibrium strategies is derived for this special case, but these results do not generalise to settings with more bidders and sellers, and, most importantly, to settings where bidders have different valuations.

In more detail, this work advances the state-of-the-art in the following ways. First, we derive equations for the bidder's expected utility in the case when all bidders use mixed strategies. In this way we can compute the equilibrium using smooth fictitious play without actually simulating the auctions. These equations cannot be easily computed for very large inputs, however, and therefore we limit the strategy space by assuming that bidders bid at most

Cite as: Approximating Mixed Nash Equilibria using Smooth Fictitious Play in Simultaneous Auctions (Short Paper), Gerding, Rabinovich, Byde, Elkind, and Jennings, *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. 1577-1580.

two different values: high in one auction, and low or equal in the remaining ones. This assumption is based on empirical work described in [5] showing this to be a best response in many settings. We then show empirically that the learning algorithm converges to ϵ-Nash equilibria.

2. THE MODEL

The model consists of m sellers, each of whom acts as an auctioneer. Each seller auctions one item; these items are complete substitutes (i.e., they are equal in terms of value and a bidder obtains no additional benefit from winning more than one of them). The m auctions are executed simultaneously; in particular, no information about the outcome of any of the auctions becomes available until the bids are placed (in real-life settings, when some of the auctions close at almost the same time, there is insufficient time to obtain the results of one auction before proceeding to bid in the next one). We assume that all the auctions are identical (i.e., a bidder is indifferent between them). Furthermore, bidders are committed to buy the items they win and thus cannot withdraw their bids. However, we also assume free disposal, meaning bidders do not incur *additional* costs for disposing of unwanted items. Finally, we assume that bidders maximise their expected profit.

2.1 The Auctions

This paper focuses on second-price sealed bid auctions, where the highest bidder wins but pays the second-highest price, although we briefly address the first-price variant when we compare the auctioneer's expected revenue in Section 4. The second-price format has several advantages for agent-based settings. First, it is communication-efficient. Second, for the single-auction case (i.e., where a bidder places a bid in at most one auction), the optimal strategy is to bid the true value and thus it requires no computation (once the valuation of the item is known). This strategy is also weakly dominant (i.e., it is independent of the other bidders' decisions), and therefore it requires no information about the preferences of other agents (such as the distribution of their valuations). Also, the auction is strategically equivalent to online auctions such as eBay by using proxy bidding.

2.2 Bidder Strategies and Expected Utility

In this section we formalise the bidding strategies and derive a bidder's expected utility when bidding in one or more auctions. We note that the equations are based on the continuous bids and valuations, whereas the numerical results are based on a discrete setting. However, the equations can be easily converted into discrete ones.

In what follows, the number of sellers (auctions) is $m \geq 2$ and the number of bidding agents (or simply bidders) is $n \geq 2$. Let $M = \{1, 2, \ldots, m\}$ and $N = \{1, 2, \ldots, n\}$ denote the set of auctions and bidders respectively. Each agent i's private valuation v_i is independently and randomly drawn from a probability distribution with support $V = [0, v_{max}]$. Let $F(x) = \int_0^x f(x)dx$ denote the cumulative distribution function. F is assumed to be common knowledge.

In general, a mixed strategy is defined as one which randomly selects between pure, deterministic strategies with a certain probability. In this problem domain, we define a mixed strategy as a mapping from a bidder's valuation to a distribution over bid vectors, where the bid vectors describe the bids for each auction. Let $\mathbf{b}$ denote the *joint bids* of all bidders for all auctions. That is, $\mathbf{b}$ is a matrix in which each element b_j^i is bidder i's bid in auction j. Furthermore, let $\mathbf{b}^i$ denote bidder $i's$ vector of bids, $\mathbf{b}^{-i}$ the bids of all agents and auctions except those of bidder i, and $\mathbf{b}_j^{-i}$ all the bids in auction j except i's. Now, the utility of agent k, provided that all the bids are known, is given by:

$$u_k(v, \mathbf{b}^k, \mathbf{b}^{-k}) = v\left[1 - \prod_{j\in M}\left(1 - P^w(\mathbf{b}_j^k, \mathbf{b}_j^{-k})\right)\right] - \sum_{j\in M} C(\mathbf{b}_j^k, \mathbf{b}_j^{-k}), \quad (1)$$

where P^w is the probability of winning a particular auction given the bids of all players placed in that auction (the left part of the equation thus denotes the probability of winning *at least* one item), and C denotes the unconditional (expected) costs for that auction given the bids. Note that this formulation of the utility is fairly general and captures a wide range of simultaneous auctions for complete substitutes, including first-price and second-price. Furthermore, although P^w generally reduces to a deterministic function (e.g., highest bidder always wins), tie breaking rules can be included, for example in the case of discrete bids, in which case the probability of winning, and thus also the utility, are probabilistic.

The above equation assumes that all bids are known. We can extend the equation to express uncertainty about the bids of other bidders as follows. Let $\mathbb{B}^i$ denote the space of bids for bidder i, and $\mathbb{B} = \mathbb{B}^1 \times \mathbb{B}^2 \times \ldots \times \mathbb{B}^n$ denote the space of all possible joint bids. Given that optimal bids are no more than a bidder's true valuation in first and second-price auctions (clearly this also applies in the case of multiple auctions), without loss of generality we assume that $\mathbb{B}^i = V^m$. Now, the expected utility U_k becomes:

$$U_k(v) = \int_{\mathbb{B}} P(\mathbf{b}^k|v)P(\mathbf{b}^{-k})u_k(v, \mathbf{b}^k, \mathbf{b}^{-k})d\mathbf{b} \quad (2)$$

where $P^k(b^k|v)$ denotes the probability of bid vector b^k occurring given the valuation v. Note that this is essentially bidder k's *mixed bidding policy* since it specifies the probabilistic strategy as a function of his valuation. Furthermore, $P(b^{-k})$ is the probability of the joint bids of the opponents. We can express this probability in terms of the bidders' valuation density function f as follows:

$$P(\mathbf{b}^{-k}) = \prod_{i\in N\setminus\{k\}}\left(\int_V P^i(\mathbf{b}^i|v')f(v')dv'\right). \quad (3)$$

Again, $P^i(\mathbf{b}^i|v')$ refers to bidder i's policy. Note that the above equations can be easily applied to the discrete case by replacing integrals with summations.

3. SMOOTH FICTITIOUS PLAY

The fictitious play concept was introduced into game theory by George W. Brown in 1951 [1]. He describes fictitious play as an iterative process formed by "two statisticians ... playing many plays" of the same game. Each statistician assumes that the opponent maintains a constant, though potentially probabilistic (mixed) strategy of game, and estimates the adversaries' strategies by the frequency of actions in the history of play, also interpreted as a player's *beliefs* about the other players' actions. The statistician then selects a best response given these beliefs.

Brown set this intuition into a formal *fictitious play* algorithm, and successfully applied it to recover a pure Nash equilibrium of several simple game instances. It has been later shown that if all the agents in the system adopt the algorithm and the game is played repeatedly, then in some types of games fictitious play converges to a pure Nash equilibrium (see e.g. [7]).

Unfortunately, a pure Nash equilibrium does not always exist in a game, and in such games fictitious play is not guaranteed to converge. Furthermore, in our problem domain the actual strategies are mappings from valuations to (distributions of) bids, but only the actual bids are observed. Therefore, even if the *beliefs* converge, it is not possible to reproduce the actual strategies from the beliefs since many different mappings result in the same set of beliefs. However, there exists a class of fictitious play modifications, called *smooth fictitious play*, which can resolve the aforementioned complications, and explicitly concentrates on mixed-strategy profiles. This algorithm class is discussed next.

Assume that the interaction between the agents in the system is described by a multi-dimensional utility function, $u : \prod_{i=1}^{n} B_i \rightarrow \mathcal{R}^n$, mapping actions independently selected by the agents into a vector of payoffs. Given that each agent selects action $b_i \in B_i$, the utility becomes $(u_1, ..., u_n) = \mathbf{u} = u(b_1, ..., b_n) \in \mathcal{R}^n$, and agent i receives utility u_i[1]. Furthermore, let $p_i(b_i)$ denote the *relative observed frequency* of action b_i by agent i, where $\sum_{b_i \in B_i} p(b_i) = 1$. The values $p_i(\cdot)$ are the beliefs about agent i's *mixed strategy*. Now, given these beliefs, the *expected utility* becomes:

$$(E(u_1), ..., E(u_n)) = E(u) = \sum_{\mathbf{b} \in \prod B_i} u(\mathbf{b}) \prod p_i(b_i)$$

Now, standard fictitious play selects the pure best response action that maximises expected utility. In this case, however, infinitesimal variation in the beliefs about adversary strategies may cause a radical change in the best response action. Allowing best response strategies to vary *smoothly* forms the core of the *smooth fictitious play* algorithm class [3]. The main idea is that, instead of taking a pure best response, the agents select their action with a probability proportional to the expected utility of that action. Although there are many variations, in this paper we apply *k-exponential cautious fictitious play* [3]. This method is commonly used since it has a natural interpretation; it is equivalent to adding an entropy component to the utility function, and this component thus promotes the usage of mixed strategies in a principled way.

In more detail, let $\sigma_i(b)$ denote the probability that (pure) action b is played. Furthermore, let $E(u(b))$ denote the expected utility from playing this action. Then the k-exponential fictitious play response is given by specifying a set of fixed weights $0 < w_{b \in B_i}$ and setting the probability of applying action b to:

$$\sigma_i(b) = \frac{w_b e^{\frac{1}{\tau} E(u(b))}}{\sum\limits_{b} w_b e^{\frac{1}{\tau} E(u(b))}}, \tag{4}$$

where τ is termed the *temperature of exploration*. Note that as $\tau \rightarrow 0$ the probability of playing the exact best response action approaches 1, while for higher temperature values alternative actions may be selected with higher probability.

When $w_b = 1$ for all $b \in B_i$, k-exponential fictitious play corresponds to computing the best (mixed) response with a modified utility function: $U(\sigma(\cdot)) = \mathbb{E}[u(\mathbf{a})] + \tau H(\sigma)$ where $H(\sigma) = -\sum_{b \in B_i} \sigma(b) \log \sigma(b)$ is the entropy of the mixed strategy. In addition to this, Fudenberg and Levine [3] have shown that k-exponential fictitious play converges to a ϵ-Nash mixed equilibrium[2], and we build on this convergence property to obtain the experimental results of this paper.

We now apply the above approach to our domain of bidding in simultaneous auctions. The standard description of smoothed fictitious play relies on sampling the actions from the mixed strategies. Furthermore, since the policies depend on an agent's type, i.e., the value it assigns to the auctioned item, we would need to sample from that as well. However, using the equations from Section 2.2, we can immediately compute the (mixed) best response strategies for each type without the need to actually play the auction game, making the update based exclusively on theoretical computations. This has the advantage of considerably reducing the computational costs. Furthermore, we can take advantage of the fact that all opponents are assumed to be symmetric (since we consider symmetric equilibria), which means that the beliefs and best responses need to be calculated only once for a single agent in each iteration (and computational complexity is independent on the number of bidders). In this case the smooth fictitious play is given by Figure 1.

Initialise:
Set iteration count to $t = 0$
for $v \in V, \mathbf{b} \in \mathbb{B}$ **do**
  Initialise beliefs $P(\mathbf{b}|v)$
end for{ Players symmetry implies $\mathbb{B} = V^m$.}
Main:
loop
  for all $v \in V$ **do**
    for all $\mathbf{b} \in \mathbb{B}$ **do**
      Compute $\sigma^*(\mathbf{b}|v)$ w.r.t. expected utility, $U(v, \mathbf{b})$, given that the opponents use mixed-strategy $P(\cdot)$, and by using Equation 4.
    end for
  end for
  for all $v \in V$ **do**
    for all $\mathbf{b} \in \mathbb{B}$ **do**
      Update $P(\mathbf{b}|v) = \frac{1}{n+1}(n * P(\mathbf{b}|v) + \sigma(\mathbf{b}|v))$
    end for
  end for
end loop

Figure 1: The Smooth Fictitious Play Algorithm.

4. EMPIRICAL RESULTS

The results in this section are based on the following settings. The bidder valuations (types) and bid values are discrete, and we use integer values ranging from 1 to 500. Each valuation occurs with equal probability, equivalent to a uniform valuation distribution in the continuous case. To allow for tractable results when $m > 2$, for each type the pure-strategy bid space is reduced to two bid values: a high bid in one of the auctions and a low bid in all others. We use a tie-breaking rule such that if more than one bidder bids the same, none of them win the item. However, given the number of possible bid values, the effect of this rule is negligibly small. For the temperature parameter we use a schedule which starts with $\tau = 1$ (i.e., almost completely random), and decreases logarithmically until it ends with $\tau = 2^{-4}$ (a logarithmic schedule is standard in smooth fictitious play). The algorithm runs for 1000 iterations.

We now evaluate the learning algorithm by showing that the smoothed fictitious play converges to an ϵ-Nash mixed equilibrium. To this end, Figure 2 depicts a typical result comparing regular and smooth fictitious play for second-price auctions. This shows that the latter has considerably better convergence properties.

A typical example of the equilibrium strategies that the agents converge to is visualised in Figure 3. Whereas the mixed strategy is in fact a 3-dimensional matrix (where the dimensions are valuation, low bid, and high bid), this figure shows a projected view of the strategies onto a 2-dimensional surface. The converged strategies are surprisingly more complex than was initially expected. Previous research showed that, by iteratively calculating the best re-

[1]Notice that the utility is based on a single instance of the game, and does not imply, nor depend on, the fact that the game will be repeated.

[2]More specifically, they show that for any $\epsilon > 0$, there exists a temperature τ such that the results converge. Furthermore, the closer ϵ comes to zero, the longer convergence may take.

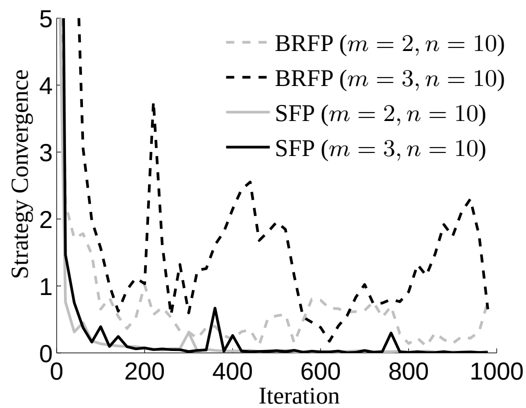

Figure 2: Convergence of the regular or best response fictitious play (BRFP) and the smoothed variant (SFP) for various settings. This is measured by first averaging the bidding strategies over a period of 20 iterations, and then taking the standard deviation. A low standard deviation means that little change has occurred.

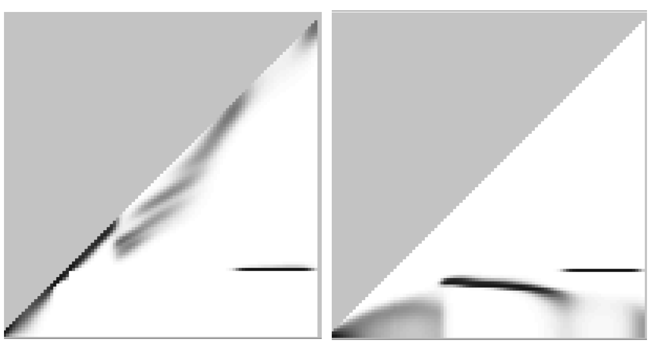

Figure 3: Projected view of a mixed strategy after 100 iterations for $m = 4, n = 6$, depicting the probabilities of each of the high bids (left) and low bids (right) given the bidder's valuation, where darker corresponds to a higher probability. Note that the grey planes indicate those parts of the strategy which are not used, since bidders have no incentives to not bid above their valuations.

sponse strategy (this is similar to best response fictitious play but without any history), the strategies cycled between two states: one in which bidders bid uniformly, and another where a single bifurcation occurs: bidders with relatively low valuations bid uniformly, whereas bidders with high valuations bid close to the true value in one auction, and low in all others [5]. Although a similar behaviour is observed in the initial iterations of the smooth fictitious play algorithm, the converged strategies are more complex.

This is shown in the example of Figure 3 where the strategy consists of 4 distinct parts. For low valuations the strategy is to close to the true value in one auction, and for the remaining auctions the bidders are indifferent between a large number of strategies (as indicated by the grey area). Here, the low bids differ significantly from the best response strategies observed in [5], where it is optimal to bid uniformly and close to the true value in all auctions when valuations are low. The discrepancy occurs because the probability of winning is very low in this case, and even a small disturbance in the utility function (e.g., by adding the entropy) can cause a large deviation in the best response strategy. For bidders with slightly higher valuations, however, a type of bifurcation is observed similar to that of the best response strategy, and bidders bid high in one auction, and low in all others. However, a bidder randomises between various high-bid strategies. A closer examination of this part of the strategy shows that a higher value for the high bid is coupled with a slightly lower value for the other bids and visa versa. That is, the bidders are indifferent between a number of pure strategies where the values for the high and low bids are negatively correlated. In the third part of the strategy, bidders bid mostly uniformly, but also randomise with a bifurcated strategy. This is consistent with the iterative best response process, which shows an alternation between uniform and bifurcated bidding. Finally, bidders with very high valuations have a strategy where they bid truthfully in one auction, and a low value in all others, which is again consistent with the best response strategy. A similar pattern is observed in other settings. To conclude, although parts of the strategy are consistent with best response dynamics, the emerging mixed strategies show some interesting and unexpected patterns which suggest that the equilibrium strategy is in fact more complex than a mixture of the pure strategies found by performing iterative best response.

5. CONCLUSIONS

In this paper we analyse equilibrium bidding strategies when intelligent bidding agents participate in multiple, simultaneous second-price auctions. We show empirically that, using best-response fictitious play, the strategies do not converge to a pure Nash equilibrium. As a result, we turn our attention to mixed Nash equilibria. Since finding such equilibria is computationally intractable, we use a learning approach called smooth fictitious play to numerically approximate the equilibrium. By combining the learning algorithm with equations about the expected utility, we can relatively quickly compute the mixed strategies without the need to simulate the auctions. The empirical results show that the strategies converge to ϵ-Nash mixed strategies.

In future work we intend to formally analyse the (non)-existence of symmetric pure Nash equilibria. Also, we would like to verify our conjecture that the equilibrium strategies consist of at most two different bid values: a low bid and a high bid, and in so doing show that the incentive to deviate from such a reduced strategy, if it exists, is very small.

6. REFERENCES

[1] G. W. Brown. Iterative solutions of games by fictitious play. In T. C. Koopmans, editor, *Activity Analysis of Production and Allocation*, pages 374–376. Wiley, 1951.

[2] R. Engelbrecht-Wiggans and R. Weber. An example of a multiobject auction game. *Management Science*, 25:1272–1277, 1979.

[3] D. Fudenberg and D. K. Levine. Consistency and cautious fictitious play. *Journal of Economic Dynamics and Control*, 19:1065–1089, 1995.

[4] D. Fudenberg and D. K. Levine. *The Theory of Learning in Games*. MIT Press, 1999.

[5] E. H. Gerding, R. K. Dash, D. C. K. Yuen, and N. R. Jennings. Bidding optimally in concurrent second-price auctions of perfectly substitutable goods. In *Proc. 6th Int. J. Conference on Autonomous Agents and Multi-Agent Systems, Hawaii, USA*, pages 267–274, 2007.

[6] V. Krishna and R. Rosenthal. Simultaneous auctions with synergies. *Games and Economic Behaviour*, 17:1–31, 1996.

[7] D. Monderer and L. S. Shapley. Fictitious play property for games with identical interests. *Journal of Economic Theory*, 68(1):258–265, 1996.

[8] R. Rosenthal and R. Wang. Simultaneous auctions with synergies and common values. *Games and Economic Behaviour*, 17:32–55, 1996.

[9] B. Szentes and R. Rosenthal. Three-object two-bidder simultaneous auctions: Chopsticks and tetrahedra. *Games and Economic Behaviour*, 44:114–133, 2003.