Using Simultaneous Best Response to Find Symmetric Bayes-Nash Equilibria in Auctions

Victor Naroditskiy Amy Greenwald

August 5, 2008

Abstract

Finding Nash and Bayes-Nash equilibria in games is a hard problem both analytically and computationally. We restrict out attention to symmetric Bayes-Nash equilibria in auctions and propose a computational method that takes advantage of the symmetry of equilibria and structure of auction games. The method is iterated best-response where all players move simultaneously. We present experimental results for single unit first- and second-price auctions with discrete values and bids.

1 Introduction

Bayes-Nash equilibria (BNE) have been derived analytically only for the simplest auction settings [3]. Such settings include single-item first- and second-price auctions with continuous distributions of bidders' values[1]. However, very little research has been devoted to auctions with discrete bids and values. There is a general game solver Gambit (http://gambit.sourceforge.net) which is the state-of-the-art solver for finite normal and extensive form games. [7] and [8] report that GAMBIT is only capable of solving relatively small games (e.g., finding BNE in the first-price sealed bid auction with 9 types and 9 values take an hour).

We propose a fast iterated best-response algorithm that takes advantage of

- structure of payoffs in first and second-price auctions
- the symmetry of equilibria

Experimental results in this paper are for single auction settings and symmetrically distributed bidders' values. In the future, we plan to search for equilibria in the settings with multiple one-shot auctions where bidders have combinatorial valuations.

1.1 Symmetric Pure Strategy BNE

BNE can be interpreted as NE of the game where player's actions are strategies [6]. For example, consider an auction with 2 symmetric bidders, the set of values $V = \{0, 2\}$, and the set of bids $B = \{1, 2\}$. We can represent the Bayesian game in the normal form where strategies represent actions, i.e., an action is a vector of bids for each value:

	1,1	2,1	1,2	2,2
1,1	*			
2,1		*		
1,2			*	
2,2				*

Nash equilibria of the normal form representation of the game are the BNE of the corresponding Bayesian game.

Symmetric pure strategies correpond to cells along the diagonal marked with '*'. Thus, an exhaustive search for a pure strategy symmetric BNE would check each cells along the diagonal. The number of

---

[1] Some literature refers to values as types or signals.

symmetric pure strategy profiles is the same as the number of actions of a player in the normal form game - the number of bids raised to the number of values $|B|^{|V|}$. Note that the number of symmetric strategies does not depend on the number of players.

We can search the space of symmetric strategies using *simultaneous* best response: starting from any cell in the payoff matrix, we calculate the best response of a player and let all players choose the best-response strategy. Because all players select the same best-response strategy, the payoff is on the diagonal.

2 Existence

The seminal paper by Nash [5] proves that finite *symmetric* games have a *symmetric* mixed strategy equilibria. However, pure strategy symmetric (or non-symmetric) equilibrium does not have to exist. At this stage, we focus on the games where we can find a pure-strategy symmetric equilibria or show that one does not exist. When it exists, symmetric pure strategy equilibrium is an appealing solution concept from the implementation point of view.

3 Model

We focus on a single object independent private-value model. There are n risk neutral bidders with quasi-linear utility functions. Each bidder's value of the object is a random variable distributed according to the same discrete probability distribution function f with support V. The distribution f is common knowledge. All bidders have the same finite set of bids B resulting in the total number of distinct pure strategies equal to $|B|^{|V|}$. Utility $u_i(s; v_i)$ of bidder i is

$$\begin{array}{ll} v_i - q_i & \text{if} \quad s_i(v_i) > \max\limits_{j \neq i} s_j(v_j) \\ (v_i - q_i)\dfrac{1}{\#\text{ties}} & \text{if} \quad s_i(v_i) = \max\limits_{j \neq i} s_j(v_j) \end{array}$$

where q_i is the bid in the first-price auction and the second highest bid in the second-price auction. In words, the utility of a bidder is her value minus the price when the bidder submits the highest bid. When the bidder ties with #ties other bidders, the winner is chosen uniformly at random from the highest bidders.

4 Utility and BNE in Auctions

The strategy profile $s^* = (s_1^*, \ldots, s_n^*)$ is a BNE if for each bidder i and for each of i's values v_i, $b_i = s_i^*(v_i)$ solves

$$\max_{b_i \in B_i} \sum_{v_{-i} \in V_{-i}} u_i(s_1^*(v_1), \ldots, b_i, \ldots, s_n^*(v_n); v_i) f(v_{-i}) \tag{1}$$

One key observation is that, in the absence of ties[2], the only information relevant to a bidder's utility in first and second-price auctions is the maximum bid of the other bidders.

$$u_i(s_i(v_i), s_{-i}(v_{-i}); v_i) = u_i(s_i(v_i), \max_{j \neq i} s_j(v_{-i}); v_i)$$

We will refer to the maximum bid of the other bidders as the price. The price distribution g can be derived from the distribution of the other bidders' values.

$$\forall p \in B_i \quad g(p) = \sum_{v_{-i} \in V_{-i} |\, \max_{j \neq i} s_j^*(v_j) = p} f_i(v_{-i}) \tag{2}$$

---

[2] We address the issue of ties in Section 7

The probability $g(p)$ that the price is p is the sum of the probabilities of all combinations of values of the other bidders that result in the maximum bid equal to p. Rewriting Equation 1 with the price distribution we get

$$\max_{b_i \in B_i} \sum_{p \in B_i} u_i(b_i, p; v_i) g(p) \tag{3}$$

The strategies $s^* = (s_1^*, \ldots, s_n^*)$ are a BNE if for each bidder i and for each of i's values v_i, $b_i = s_i^*(v_i)$ solves Equation 3. Note that $g(p)$ in Equation 3 is determined by the strategies of the other bidders.

Given g, Equation 3 is more compact than Equation 1. The summation in Equation 1 has V_i^{n-1} terms. The summation in Equation 3 has at most B_i terms. This difference becomes important when there are more than 2 bidders. We will show that in the symmetric case g can be calculated much faster than in Equation 2.

5 Iterated Best Response

Strategy profile $s = (s_1, ..., s_n)$ is a symmetric BNE if for all i, s_i is the best response to the price distribution g resulting from $n-1$ bidders playing the same strategy s_i. We search for such profile via an iterated simultaneous best-response procedure starting from some initial distribution of prices.

```
• initialize $g$
• repeat until $s$ are a symmetric equilibrium
    – for every $v_i \in V_i$
        * bidder $i$ finds the bid $s_i(v_i)$ that is the best response to $g$.[3]
    – $h$ is the frequency distribution of bids $s_i$
    – $H$ is the corresponding cumulative distribution
    – $g(p) = H^{n-1}(p) - H^{n-1}(p-1)$
```

The strategies $s = (s_1, \ldots, s_n)$ are a symmetric BNE if s_i is a best response to the price distribution resulting from $n-1$ bidders playing s_i. We check this condition at the start of the loop.

In each iteration we calculate bidder i's best response s_i to the price distribution g. The best response is calculated for each value $v_i \in V_i$. At the end of an iteration, g is set to the price distribution resulting from $n-1$ bidders playing s_i. This price distribution is calculated using the cumulative distribution H of bids submitted by bidder i. The cumulative distribution of the maximum of $n-1$ bids distributed according to H is H^{n-1}. For integer prices p the probability function $g(p)$ can be calculated from H: $g(p) = H^{n-1}(p) - H^{n-1}(p-1)$.

6 Experiments

We test the procedure in first- and second-price auctions with risk-neutral bidders. The number of bidders ranges between 2 and 11. Bidders have integer values uniformly distributed between 0 and k. The parameter k ranges from 2 to 50. Given a value v and a price distribution, an optimal bid is calculated by comparing profits from bidding each of the $v+1$ bids $0, 1, \ldots, v$. If the bid is equal to the price, we use $\frac{1}{n}$ as an approximation of the probability that the bidder wins the auction[4]. The bidder pays his bid in the first-price auction and the price in the second-price auction. We tried two ways of initializing the price distribution: to zero and to the value distribution.

---

[3] If multiple bids are best responses, we set $s_i(v_i)$ to the mixed strategy of submitting any of the best-response bids with equal probability.

[4] The probability is exact only when there are two bidder. Section 7 describes how the probabilities can be calculated exactly. However, at the time of running the experiments, we used $\frac{1}{n}$ as an approximation. We re-ran some of the experiments with correct probabilities and confirmed that the results reported here still hold.

6.1 Second-Price Auctions

The second-price auction has a known dominant strategy of bidding the true value. The iterative procedure converges to the dominant strategy after 2 iterations when the price distribution is initialized to zero and after 1 iteration when the price distribution is initialized to the value distribution.

6.2 First-Price Auctions

The procedure converges to an equilibrium in most of our experiments with varying number of bidders and initial price distributions for values of k below 9. It takes under 8 iterations (usually 1 iteration) to reach convergence when the price distribution is initialized to the value distribution. It takes under 27 iterations (5 iterations on average) to reach convergence when the price distribution is initialized to zero. Table 1 illustrates the steps of the iterative procedure in the auction with 2 bidders, $k = 4$, and initial belief that the price is zero. Convergence is reached after 3 iterations.

price/value	0	1	2	3	4
price distr	1.0	0	0	0	0
BR strategy	0,1,2,3,4	1,2,3,4	1,2,3,4	1,2,3,4	1,2,3,4
price distr	0.04	0.24	0.24	0.24	0.24
BR strategy	0	1	2	3	4
price distr	0.2	0.2	0.2	0.2	0.2
BR strategy	0	1	2	3	4
price distr	0.2	0.2	0.2	0.2	0.2
BR strategy	0	1	2	3	4

Table 1: Sample Run of Iterated Myopic Best Response. "price distr" is the probability distribution of prices resulting from the opponent playing the best-response (BR) strategy to the price distribution from the previous iteration.

Bidding $\frac{n-1}{n}$ of the value is a BNE of a first-price auction when bidders' values are identically distributed according to a continuous uniform distribution [3]. Not surprisingly, the equilibrium strategies we find are similar to the equilibrium strategies for the continuous case. Examples of equilibria are in Table 2. The continuous counterparts for these discrete strategies are to bid $\frac{10}{11}$ and $\frac{1}{2}$ of the value respectively.

k	# bidders	equilibrium strategy
8	11	0:0,1:0,2:1,3:2,4:3,5:4,6:5,7:6,8:7
7	2	0:0,1:0,2:1,3:1,4:2,5:2,6:3,7:3

Table 2: Equilibria in First-Price Auctions. 3:1, 4:2,... denotes the equilibrium strategy of bidding 1 when the value is 3 and bidding 2 when the value is 4.

In the cases when myopic best response does not converge to an equilibrium, it converges to a cycle. Surprisingly, when the price distribution was initialized to value distribution, the procedure converged to an equilibrium for any odd value of k.

7 Ties

Here we address the issue of ties, which was ignored when we said that only the bid and the price affect one's utility.

$$\Pr(\text{tie and win}) = \sum_{m=1}^{n-1} \underbrace{\frac{1}{m+1}}_{} \underbrace{\binom{n-1}{m}}_{} \underbrace{h^m(b_i) H^{n-m-1}(b_i - 1)}_{} \tag{4}$$

The first term is the probability of winning with tied with m bidders. The second term is the number of ways to choose m bidders to tie with. The last term is the probability of a tie with exactly m bidders. This computation takes $n-1$ steps assuming that all the terms inside the summation are computed once per iteration of best response. The computation of ties has been previously done in [1].

8 Related Work

The related work pertains to symmetric risk-neutral bidders and the independent private value model.

8.1 Analytical Results

Existence results on equilibria are summarized in the table. References and comments appear below.

	DB DV	DB CV	CB DV	CB CV
FP	M	P	M	P
SP	M	P	D	D

Table 3: Equilibria in Auctions: M - symmetric mixed, P - symmetric pure, D - dominant; DB - discrete bids, DV - discrete values, CB - continuous bids, CV - continuous values

Dominant Strategy Solvable Second Price (SP) with continuous bids and continuous values [10] and SP with continuous bids and discrete values.

Symmetric Pure Strategy Equilibrium First Price (FP) with continuous bids and continuous values [3].

Evenly Spaced Discrete Bids FP ([1]) and SP with continuous values [11].

Note that there is no dominant strategy equilibrium in SP with continuous values. We can illustrate this with an example showing that for some value there is no bid that is dominating all the other bids. Suppose there are 2 bidders, the possible bids are 0 and 1, and the values are uniformly distributed between 0 and 1. A bidder with the value close to one (say, value is .9) prefers bidding 1 if the expected bid of the other bidder is close to zero, but he prefers bidding 0 when the expected bid of the other bidder is above .9.

Arbitrary Discrete Bids SP with continuous values [4].

Symmetric Mixed Strategy Equilibrium FP with continuous bids and discrete values [2]. FP and SP with discrete bids and discrete values (normal form game).

8.2 Computational Results

[7] describes a procedure for computing best-response strategies and finding Bayes-Nash equilibrium for two-player infinite games with types drawn from a piecewise-uniform distribution. [9] uses iterated tabu best response to search for NE in normal form games. Their work is not restricted to symmetric equilibria and considers smaller normal form games.

9 Future Work

We envision three directions for future work.

- design a procedure with better convergence properties

- extend to settings with multiple one-shot simultaneous auctions and combinatorial valuations
- search for a mixed strategy BNE

References

[1] Michael Suk-Young Chwe. The discrete bid first auction. *Economics Letters*, 31(4):303–306, December 1989.

[2] D. Fudenberg and J. Tirole. *Game Theory.* MIT Press, Cambridge, 1991.

[3] Vijay Krishna. *Auction Theory.* Academic Press, March 2002.

[4] Timothy Mathews and Abhijit Sengupta. Sealed bid second price auctions with discrete bidding. *Applied Economics Research Bulletin*, Special Issue I (Auctions):31–52, 2008.

[5] J. Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, 1951.

[6] M. Osborne and A. Rubinstein. *A Course in Game Theory.* MIT Press, Cambridge, 1994.

[7] Daniel M. Reeves and Michael P. Wellman. Computing best-response strategies in infinite games of incomplete information. In *UAI*, 2004.

[8] Daniel M. Reeves, Michael P. Wellman, Jeffrey K. Mackie-Mason, and Anna Osepayshvili. Exploring bidding strategies for market-based scheduling. *Decis. Support Syst.*, 39(1):67–85, March 2005.

[9] Ashish Sureka and Peter R. Wurman. Using tabu best-response search to find pure strategy nash equilibria in normal form games. In *AAMAS '05: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 1023–1029, New York, NY, USA, 2005. ACM.

[10] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

[11] Jin Yu. *Discrete approximation of continuous allocation mechanism.* PhD thesis, California Institute of Technology, 1999.