Comput Econ (2009) 34:243–271
DOI 10.1007/s10614-009-9169-3

Auctions and Differential Pricing: Optimal Seller and Bidder Strategies in Second-chance Offers

Yu-An Sun · Poorvi Vora

Accepted: 27 January 2009 / Published online: 17 February 2009

Abstract The *second chance offer* is a common seller practice on eBay. It consists of price discrimination against the losing bidder, who is offered an identical item at the value of his or her highest bid. Prior work has shown that, if the goods are private-value goods and price discrimination is certain, rational bidders can anticipate it and accordingly modify their bidding strategies; this results in revenue loss for the seller. This paper hence examines the impact of randomized price discrimination. It examines a similar, more general problem: a seller has k items of the private-value good. They are sold to N bidders in a two-stage game. The first stage is a sealed-bid first-price auction with N bidders. The second stage is a take-it-or-leave-it offer to each of $k-1$ losing bidders; randomized between a fixed-price offer and a second-chance offer. Showing that analytic techniques do not provide complete solutions because bidding strategies are not always monotonic increasing, this paper uses genetic algorithm simulations to determine the Bayesian (near-Nash) equilibrium strategies for bidders and sellers, for $N=8$ and different values of k. It finds that price discrimination is beneficial to the seller when k is small, as item scarcity increases competition among bidders. The paper's use of randomized seller strategies and genetic algorithm simulations is unique in the study of the second-chance offer.

Keywords Auctions · Second-chance offer · Genetic algorithms · Privacy · Price discrimination

Y.-A. Sun (✉) · P. Vora
Department of Computer Science, The George Washington University, 800 22nd St, NW, Washington, DC 20052, USA
e-mail: ysun.hosp@gmail.com

1 Introduction

Consider an open-cry ascending-price auction such as on eBay. If the auction is a stand-alone game not connected to any other, and timing effects are ignored, each losing bidder's highest bid reveals some information on his valuation. On eBay, this fact may be exploited by the seller through the *second-chance offer*: a second stage of price discrimination where a losing bidder is offered another item of the same good at his highest failed bid. A question that immediately arises is the following: what is the Nash equilibrium for the second-chance offer? That is, what are symmetric bidding strategies? This paper takes an analytical approach to a problem similar to the second-chance offer, making several simplifying assumptions in order to obtain analytical results. While the results may not directly be applied to eBay's second-chance offer, the analysis provides insight into the impact of price discrimination on bidding strategies in auctions.

Analytical studies of Nash equilibria for the second-chance offer (Joshi et al. 2005; Salmon and Wilson 2008) have generalized the second stage and have made the simplifying assumptions of risk-neutral bidders, and private-value goods. In particular, Joshi et al find that, when the number of bidders is greater than twice the number of available items, equilibria exist such that bidders lower their bids in anticipation of the second-chance offer. Salmon and Wilson study a second stage price offer that inverts the known symmetric bidder strategy (including one that anticipates the second-chance offer) and thus charges the bidder his valuation. They are unable to find equilibrium bidding strategies for the game, but find that the only existing Nash equilibria for two bidders are mixed-strategy equilibria. In an experimental study, they also find that real bidders do not modify their bids even if they are aware that a second-chance offer is to follow. Joshi et al suggest that the existence of automated bidding tools based on symmetric equilibrium strategies may enable bidders to bid so as to decrease the incentive for the seller to price discriminate.

Continuing the line of reasoning of the existing analytical approaches to the problem, this paper addresses a game, *Game Price Discrimination*, similar to the second-chance offer, where the seller randomizes the second-chance. A two-stage game is played between N bidders and one seller. In the first stage, all bidders enter a sealed-bid, first-price auction, and one item is sold. After the first stage is over, all non-winning bidders enter the second stage for $k-1$ items identical to the one sold in the first stage. In this stage, the seller makes a second chance offer with probability α, or a fixed-priced offer, P, with probability $1-\alpha$. Bidders accept any offer that is not greater than valuation. The game is defined formally in Sect. 2.

Two major differences exist between *Game Price Discrimination* and the motivating example of eBay's second-chance offer. First, the first stage of *Game Price Discrimination* consists of a first-price sealed-bid auction. It is hence not influenced by the timing effects of open-cry auctions such as those on eBay, where bidders may wait till the last few minutes to make their bids. Further, while ascending-price auctions such as those on eBay are similar to second-price sealed-bid auctions (Ockenfels and Roth 2002; Duffy and Utku 2008), this paper studies a first-price sealed-bid auction. We make this choice to simplify the analysis. Second, in order to obtain analytical expressions for Nash equilibria, this paper makes some simplifying assumptions.

In particular, it assumes that bidders are risk-neutral, and that the goods are private-value; that is, a bidder knows his own valuation, but not that of another. In general, however, this is not true for eBay or for most real auctions. The valuations for most eBay goods possess a common-value component (Bajari and Hortaçsu 2003). Further, typically, valuations in real auctions are influenced by both: private-values as well as a common-value signal (Goeree and Offerman 2002). Finally, human bidders in auctions do not behave in a risk-neutral fashion, and Nash equilibria have been shown to be of limited use in predicting the behavior of bidders in real-world auctions. Thus, the assumptions made in the paper limit the extent to which the results may be directly applied to the problem of the second-chance offer on eBay. However, these assumptions are typically made in the analytical study of auctions in order to simplify the problem so that it is mathematically tractable (see, for example, Krishna 2002). Additionally, Nash equilibrium approaches may continue to be valuable for the study of auctions where bidders are not human, but are automated agents acting on behalf of the bidders.

This paper demonstrates that analytical techniques, which assume monotonic increasing bidding strategies, do not provide complete solutions to the game. Genetic algorithms are hence used to solve the two-stage game and determine the optimal Bayesian strategies for both bidders and sellers. Two-population GA experiments are conducted where bidders and sellers form the two populations. Bidders' chromosomes are represented by a 20-entry lookup table, and form their bidding strategies. Sellers' chromosomes are represented by the probability to make a second chance offer, α, and the fixed-price offer, P, and represent the seller's probabilistic strategy. When the population converges, near Nash equilibria[1] are obtained. The results of the GA experiments demonstrate that when items are scarce—that is, less than or equal to half of the number of bidders ($k \leq 4$ when $N = 8$)—optimal bidding strategies are monotonic increasing. Otherwise, the strategies are not monotonic increasing, and, for $k > 5$, preliminary results imply the existence of mixed strategies. Further, item scarcity motivates bidders to bid high, and price discrimination ($\alpha = 1$) is an optimal seller strategy when items are scarce.

The contributions of this paper are as follows. First, it obtains near Nash equilibrium results for a general game similar to the second-chance offer, where the seller's strategy is randomized. It is unique in considering randomization of the second-chance offer, and thus enables a more complete understanding of when it is optimal for a seller to follow an auction stage with a differential pricing stage. Second, it is unique in its use of GAs to address the second-chance offer and price-discrimination following an auction, presenting an approach for obtaining solutions when equilibrium strategies are not monotonic increasing.

This paper is organized as follows; Sect. 2 presents the general approach, and Sect. 3 related work. The analysis is presented in Sect. 4 and GA simulation results in Sect. 5. Section 6 contains conclusions and future research directions.

[1] GA experiments result in a solution state that is close to a Nash equilibrium, but is not completely perfect (Riechmann 2001). This is because, in a GA experiment, due to mutation, the entire population does not adopt the Nash symmetric strategy, but almost all of it does.

2 The Model

We model a two stage price discrimination game—an auction stage followed by a seller-offering stage—as follows.

Game Price Discrimination

1. Stage I: N bidders join a first-price sealed-bid auction. All bidders simultaneously and independently make bids that each bidder is committed to—i.e., a bidder has to buy the item if it is offered to him at the value of his bid. The bidder with the highest bid wins the auction and pays his bid. On the occurrence of a tie, the winner is chosen at random. The remaining $N-1$ bidders enter Stage II.
2. Stage II: The seller offers an identical item to some of the remaining bidders:
   (a) The *failed bid option:* With probability α chosen by the seller, the $k-1$ objects are offered to the highest $k-1$ ($k>1$) of the losing bidders, each at the respective bidder's highest bid in Stage I.
   (b) The *uniform price option:* With probability $1-\alpha$, the $k-1$ objects are offered at a uniform price $P \neq 0$ chosen by the seller.

Bidders have to accept an offer in Stage I. In Stage II, bidders accept those offers that are not greater than valuation.

2.1 Notation

We follow the notation of Krishna (2002) and denote the bidder's valuation by x, the bid by b, the payoff by Π, and the expectation operator by $E[.]$. $\beta(x)$ is the bidding function. $G(x)$ denotes the probability that a given valuation x is the highest among N bidders; $g(x)$ denotes its derivative. $H(x)$ denotes the probability that a given valuation x is among the highest k ones, but is not the highest; $h(x)$ denotes its derivative. P denotes the uniform-price offer made in Stage II.

2.2 Assumptions

We make the following assumptions:

1. Goods are private-value. x is independent and identically distributed with a uniform distribution over interval $[0, 1]$, with cumulative distribution function $F(x) = x$, and the corresponding probability distribution function f. Hence,

$$G(x) = [F(x)]^{N-1}$$

and

$$H(x) = \sum_{i=1^{k-1}}^{N-1} C_i [F(x)]^{N-i-1} [1 - F(x)]^i$$

where ${}^jC_i$ is the number of ways of choosing i items from j identical items.

2. $\beta(x)$ is identical for all bidders; i.e. bidders are symmetric.
3. Bidders are risk-neutral.
4. Bidders reject any offer exceeding valuation, and accept any other.

Each bidder's payoff is calculated as the difference between the price paid for the item and the bidder's valuation. In this two-stage game, bidders are seeking to maximize their expected payoff and sellers are seeking to maximize their expected revenue over both stages. We seek equilibrium strategies for both bidders (β) and sellers (α and P).

As mentioned in the introduction, bidders in real auctions are not typically risk-neutral, and do not use symmetric strategies. As is common in the literature on the analysis of auctions and bidding strategies (see, for example, Krishna 2002), we make these assumptions in order to obtain analytical results.

2.3 The General Approach

We first approach the problem analytically with a simpler version of the game, and then, because the analytical approach does not provide a complete solution, we use GAs to examine a more general version of the game. In this section we sketch both approaches.

2.3.1 The Analytical Approach: Case 1

In Case 1, the bidding strategy is assumed pure and monotonic increasing (this is also standard in analytical approaches to bidding strategies—see, for example, Krishna 2002). Further, the $k-1$ objects in Stage II are offered to the $k-1$ highest bidders in Stage II (that is, to the $k-1$ of the k highest bidders who did not win the object in Stage I). To analytically determine the equilibrium strategies, we would first fix the equilibrium strategy of the seller—α and P, assume it is known to the bidder, and determine the corresponding equilibrium strategy of the bidder. Finally, we would determine the values of α and P that maximize seller revenue.

When all other bidders use monotonic increasing pure strategy β, the probability of winning when bidding b is $G(\beta^{-1}(b))$, and the probability of being among the top k bidders but not the top bidder is $H(\beta^{-1}(b))$. Hence the expected payoff to the bidder with bid b and valuation x is:

$$E[\Pi(b,x)] = \begin{cases} G(\beta^{-1}(b))(x-b) + H(\beta^{-1}(b))[\alpha(x-b) + (1-\alpha)(x-P)] & x > P \\ G(\beta^{-1}(b))(x-b) + H(\beta^{-1}(b))[\alpha(x-b)] & x \le P \end{cases} \tag{1}$$

The symmetric equilibrium bidding strategy is the value of β that is a mutual best response for all bidders, and is typically obtained by maximizing (1) with respect to b, assuming $b = \beta(x)$. We find, however, (see Sect. 4) that the bidding strategy solutions thus obtained for a given value of α, P and k are not always monotonic increasing—that is, they do not always satisfy the assumptions under which they were derived, and are hence not always valid solutions. Hence, we are not able to derive closed-form expressions for revenue as a function of α and P, which are needed to determine the seller's equilibrium strategy (the values of α and P that maximize revenue). Thus the analytical approach does not provide a complete solution. In Sect. 4 we describe this approach and its limitations in detail.

2.3.2 The Genetic Algorithm Approach: Case 2

We next examine, using GA simulations, arbitrary (non-monotonic-increasing) strategies in a more general game. In Case 2, the bidding strategy β is not assumed to be monotonic increasing. It is assumed that $0 \leq \beta(x) \leq x$. Because β is not necessarily monotonic increasing, it is possible that a bidder with a lower bid might be able to pay the fixed-price for the object while a bidder with a higher bid might not. Hence, in Case 2, the fixed-price offer in Stage II, made with probability $1-\alpha$, is made not only to the $k-1$ highest bidders, but to all $N-1$ losing bidders. A random $k-1$ are chosen from all bidders who accept the fixed price offer. Thus the game we can study using the GA approach is more general than the game theoretically analyzed.

When bidding strategies are not monotonic increasing, the probability of winning Stage I with bid b is not $G(\beta^{-1}(b))$ as in (1). The following formula provides the expected payoff:

$$E[\Pi(\beta(x),x)] = \begin{cases} (x-\beta(x))Pr[\beta(x)=b_1] \\ +\alpha(x-\beta(x))Pr[b=b_2\ OR\ b_3\ OR\ ...\ OR\ b_k] \\ +(1-\alpha)(x-P)Pr[bidder\ chosen\ fOR\ fixed\ price\ offer] & x>P \\ \\ (x-\beta(x))Pr[\beta(x)=b_1] \\ +\alpha(x-\beta(x))Pr[b=b_2\ OR\ b_3\ OR\ ...\ OR\ b_k] & x \leq P \end{cases}$$

Again, the equilibrium strategy is the mutual best response strategy.

The solution space of player strategies is expected to be large, as the bidding strategy is not constrained to a particular form. We use GA experiments to perform a search over the large solution space. In the GA experiments, the bidder's fitness function is payoff, and the seller's fitness function is revenue. The bidder's chromosomes define his unconstrained strategy β. Using the GA experiments, equilibrium strategies for bidder and seller are obtained when the chromosomes converge. Our approach and the experimental results are described in detail in Sect. 5.

3 Related Work

There have been several applications of genetic algorithms to the solution of economic problems. One application is mechanism design and evaluation. The genetic algorithm approach has been adopted to investigate an optimal mechanism for an online auction trading environment (Cliff 2003, 2007). Cliff noted that software agents can be used in online auctions and current human-developed auctions are not necessarily optimal; in particular, Cliff discovered a hybrid auction market evolving in his experiments, which is much more market efficient and results in more overall market profit than any human-designed auction mechanisms. Similarly, genetic algorithms have been adopted to evaluate various auctions including first-price and second-price sealed bid auctions, to observe the evolution of a hybrid mechanism in simulations (Byde 2003). It is important to note that Byde established that the GA-based solution is optimal regardless of whether it is a human-trading market or an agent-based trading market. Byde also noted that the advantage of the GA-based approach is that not-theoretically-analyzable factors can be taken into consideration during the simulation process by evolving. Similar to our approach, Byde used a 1-to-1 mapping table to represent the bidding strategy where the entry of the table is the bidders' valuation, and the outcome of the lookup table is the bid.

Genetic algorithms have also been applied to well-established problems in economics, such as the prisoner's dilemma, auctions, the cobweb model and other microeconomic problems (Riechmann 2001; Dawid 1996, 1999; Arifovic 1994). In prisoner's dilemma, defection is the dominant strategy if the game is only played once. Optimal strategies in the iterated prisoner's dilemma (IPD) game evolved from genetic algorithm experiments have similar properties as TIT FOR TAT, a strategy submitted by Anatol Rapoport in a previous IPD strategy contest (Axelrod 1984, 1987). Axelrod concluded that the genetic algorithm is an effective optimization technique in a large problem space.

Genetic algorithm experiments have also been conducted to explain the anomaly of human auctions (Andreoni and Miller 1995). The bidders gradually learn the optimal strategy by evolving. Andreoni and Miller examined the evolved strategies in auctions with common-value goods, affiliated private-value goods and independent private-value goods over a period of 1,000 generations. Each generation consists of twenty rounds of auctions and the fitness function is defined as the total profit over twenty rounds. Andreoni and Miller also examined experiments with settings of eight and four bidders. They concluded that bidders in the experiments did reach Nash equilibrium in the common-value, first price auctions, and noted it is relatively difficult to converge to equilibrium in auctions because of poor feedback in the auction environment.

The genetic algorithm approach has been found to be useful to analyze economic problems because it always initiates a heterogeneous population to begin further simulations; this is a major advantage over theoretical analysis (Dawid and Kopel 1998).

Bidding strategies in a multi-stage game similar to the second-chance offer have also been studied (Salmon and Wilson 2008; Joshi et al. 2005). Joshi et al demonstrate that, when bidder and seller strategies are deterministic, bidders lower their bids considerably, penalizing the seller for privacy infringement. Salmon and Wilson studied a similar game, the English-Ultimatum game, that has a first price auction as first stage and a take-it-or-leave-it offer at the second stage. Salmon and Wilson showed that there is only a mixed strategy equilibrium in the theoretical 2-item-2-bidder, English-Ultimatum game.

4 The Analytical Approach: Case 1

This section deals with the analysis of Case 1. The genetic algorithm approach to solving Case 2 is described in detail in Sect. 5.

We first observe that a dominant pure strategy does not exist.[2] We hence attempt to analytically obtain the Bayesian–Nash equilibrium bidding strategy.

Let β^* denote the (symmetric) equilibrium bidding strategy for a fixed, known value of α and P (in equilibrium, the values of α and P are known to the bidder). For the purposes of this analysis, we use the standard assumption (see, for example, Krishna 2002) that β^* is monotonic increasing. For a more general strategy in a more general game, see Sect. 5. We obtain the following result.

Theorem 1 *If α and P are fixed and known to the bidder, and*

$$\beta^*(x) = \begin{cases} \dfrac{\int_0^x yg(y)dy + \alpha \int_0^x yh(y)dy + (1-\alpha)\int_P^x (y-P)h(y)dy}{G(x)+\alpha H(x)} & x > P \\[2ex] \dfrac{\int_0^x yg(y)dy + \alpha \int_0^x yh(y)dy}{G(x)+\alpha H(x)} & x \le P \end{cases}$$

is non-negative, monotonic increasing, and not greater than x, $\beta^*(x)$ is a Bayesian–Nash equilibrium bidding strategy for Game Price Discrimination.

The proof may be found in the appendix.

Through numerical evaluations, we find that there are several values of the triplet α, P, k such that $\beta^*(x)$ is not monotonic increasing (non-monotonic strategies $\beta^*(x)$ for $k = 5, 6$ are plotted in Figs. 5 and 6). Hence $\beta^*(x)$ does not always provide the equilibrium strategy. We cannot use it to compute the revenue for each value of α and P in order to determine the optimal α and P for the seller. We hence need an alternate approach to determine the equilibrium seller strategy, and the corresponding bidder strategy. In particular, we need an approach that does not require the bidding strategy to be monotonic increasing. The next section describes the use of genetic algorithm simulations to derive equilibrium strategies that are not constrained to be monotonic increasing in a more general game.

5 Genetic Algorithms: Case 2

In this section, we present our approach to Case 2. Notice that, as expected, by arguments similar to those for Case 1, there is no dominant deterministic strategy for

[2] Suppose a dominant strategy did exist. Consider a valuation $x < P$. Denote the corresponding bid b. Then, either $b = x$ or $b < x$. If $b < x$, there would be a distribution of other bids such that b does not win either stage, but a b', $x > b' > b$, would win a stage and provide a non-zero payoff. If $b = x$, the payoff of all stages except the fixed-price stage is zero. For a given distribution of other bids, a bid b', such that $x = b > b'$, can be found such that the rank of b' is identical to that of b, but its payoff is greater because it possesses a non-zero payoff in Stage I and in the price discriminatory Stage II.

the Case 2 game.[3] Genetic algorithm simulations are used to find Bayesian–Nash equilibria among sellers and bidders in a multi-stage game of incomplete information. From an evolutionary game theory point of view, the survived strategies result in mutual best responses for all the players because the strategies with lower utility are eliminated during the evolving process; the near-Nash equilibria obtained from our experiments are hence also evolutionarily-stable.

5.1 Experimental Settings

We performed 7 GA experiments, with $N = 8$ and $k = 2, 3, ...8$. The fitness functions are expected payoff and expected revenue for bidder and seller respectively. The seller's chromosome consists of the values of α and P. The bidder's chromosome is coded as a 20-entry lookup table. It represents the bidder's bidding function, i.e. the bidder's valuation is the index of the lookup table and the bid is the content of the corresponding entry. Because Stage I requires a bid that the bidder can commit to, the bid is not greater than valuation. All chromosome entries lie between 0 and 1.

The stop condition for our simulation is number of generations = $4,000$ for $k = 2, 3, 4$ and $10,000$ for all other values of k. This is because of larger variation in the bidding population for larger values of k. The population is set at 500 sellers for every experiment, and $N = 8$ distinct bidders for each seller. Bidder chromosomes are randomly generated at the beginning of the simulation. Bidder valuations are randomly generated *each generation*. 500 auctions are conducted during each generation for experiments 1, 2 and 3; 1,000 auctions for the other experiments, where the variation among bidding strategies is larger even at convergence.

The optimal mutation rate is set to $\frac{1}{22}$ because there are 22 different variables in the sellers and bidders chromosomes (see Muhlenbein and Schlierkamp-Voosen 1993). We also implement the mutation operator suggested by Muhlenbein and Schlierkamp-Voosen:

$$Var_i^{Mut} = Var_i + s_i \cdot a_i \cdot r_i$$

where $i \in 1, 2, ..., 4{,}000$. A variable Var_i is selected with probability $\frac{1}{22}$ to mutate. For every variable, we flip a biased coin; with probability $\frac{1}{22}$, the variable mutates. If the variable is not selected to mutate, its value does not change. Otherwise, the new variable Var_i^{Mut} is computed by adding or subtracting a small value $a_i \cdot r_i$. Addition or subtraction is randomly chosen with probability 0.5, and is represented by random variable s_i.

$$s_i \in \{-1, \ +1\}$$
$$r_i : mutation\ range = 0.1 * (1 - 0)$$

[3] The difference in the two games is that, in Case 2, the fixed-price offer is independent of the bid. However, all other pay-offs are identical to those of Case 1.

r_i is set to be $0.1 * (1 - 0)$ because the range of bidders' chromosomes and sellers' chromosomes are all between 0 and 1, and the mutation range is fixed to be an optimal 10% of this range.

$$a_i = \sum_{l=1}^{L} u_l 2^{-l}, u_l \in \{0, 1\} \textit{ chosen at random}, \ L : \textit{mutation precision}, \quad L = 16$$

u_l is initially 0 and is randomly chosen to be 1 with probability $\frac{1}{L}$. On average, there will only be one value of u_l with value 1, which makes $a_i = 2^{-l}$ for some $l \in \{1, 2, ..., L\}$.

We also performed some experiments with different settings to investigate the existence of mixed strategies when our GA experiments do not converge for $k = 6, 7, 8$. These are described in more detail in Sect. 5.3.5.

5.2 Detailed Steps

After the population is initialized, each bidder submits a bid according to his bidding function (the lookup table). The seller's fitness score is the total revenue over both stages, while that of the bidder is the individual payoff after both stages.

The bidder's payoff is calculated as the difference between his purchasing price—either his submitted bid at stage I or a fixed price offer—and his valuation. If the seller offers a fixed price higher than the bidder's valuation, the offer is rejected, and does not contribute to the seller's revenue, and the bidder's payoff is zero. Because the fitness function takes inputs from the chromosomes of both sellers and bidders, sellers and bidders can be viewed as two species that affect each other while evolving over generations. Figure 1 illustrates the detailed steps in one generation. Seven different experiments are conducted.

5.3 Simulation Results

5.3.1 Experiments One, Two and Three—500 Sellers, 8 Bidders, $k = 2, 3$, and 4 Items

In all three experiments, the seller's α chromosome converges to $\alpha = 0.9999$ and the P chromosome does not converge. This is because when α is nearly 1, a second chance offer always occurs in Stage II; hence the fixed price, P, is never tested; and therefore does not converge.

Note that the analytically-obtained strategy β^* of Case 1, for $\alpha = 1$, is independent of P, and is monotonic increasing, non-negative and $\beta^*(x) \leq x \ \forall x \in [0, 1]$. Hence, we expect that the bidder strategy for Case 2 obtained through the GA should be close to β^*, up to some error due to mutation. (The difference between the two games—a random choice of winning bidders in the fixed-price offer for Case 2, and the choice of highest bidders in Case 1—does not play a role here. For $\alpha = 0.9999$, the two games are virtually identical because a fixed-price offer is almost never made.) Figures 2, 3 and 4 show both the analytically-obtained strategy β^* as well as the bidding strategy

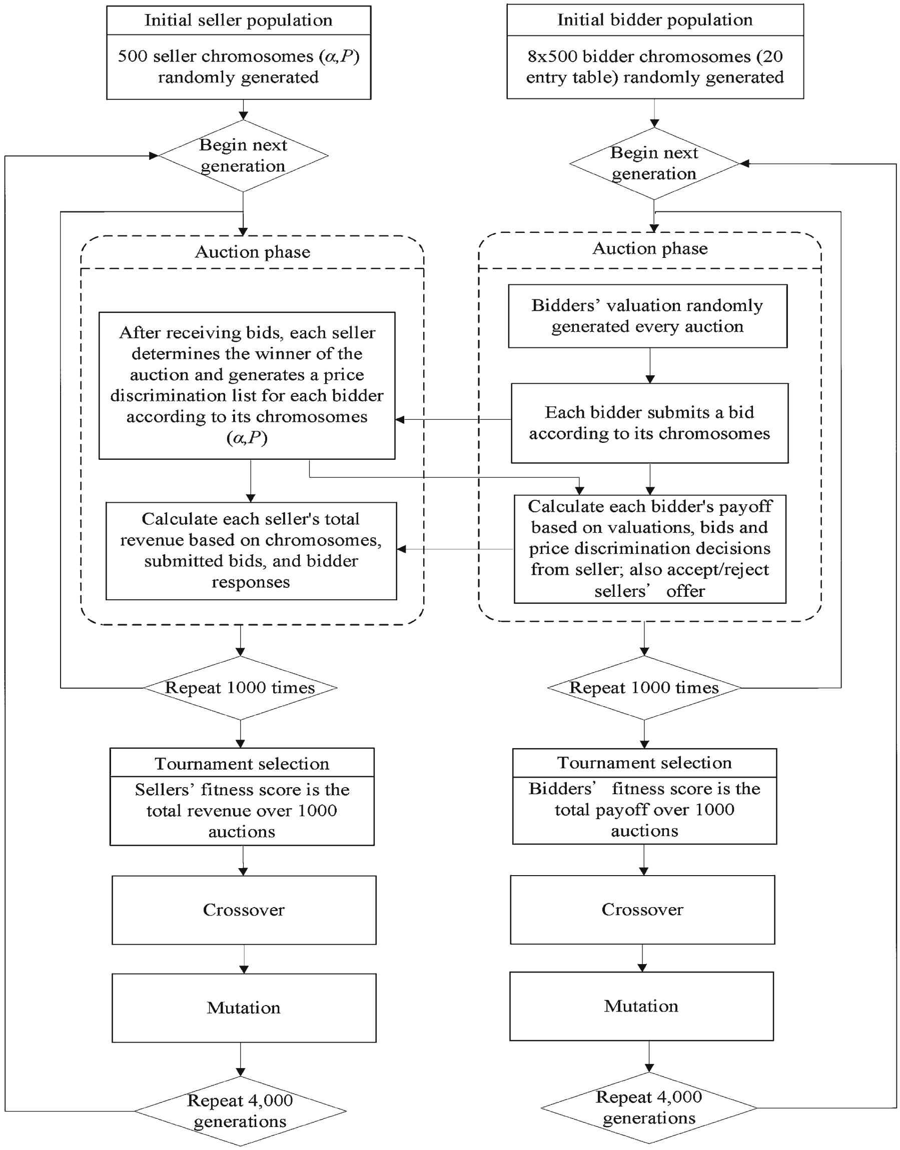

Fig. 1 Detailed process of evolutionary programming experiments

obtained from GA simulations for all three experiments. The results are as expected, the strategy obtained through the GA is indeed close to β^* in all three cases.

Note that the fact that β^* is close to the strategy obtained through GA simulations does not imply that the study of Case 2 was unnecessary. Through numerical calculations, we find that, for each value of $k = 2, 3, 4$, there exist values of α and P for which β^* is not monotonic increasing. Hence it is not possible to analytically compare seller revenues for all values of α and P in Case 1 to determine those that provide the

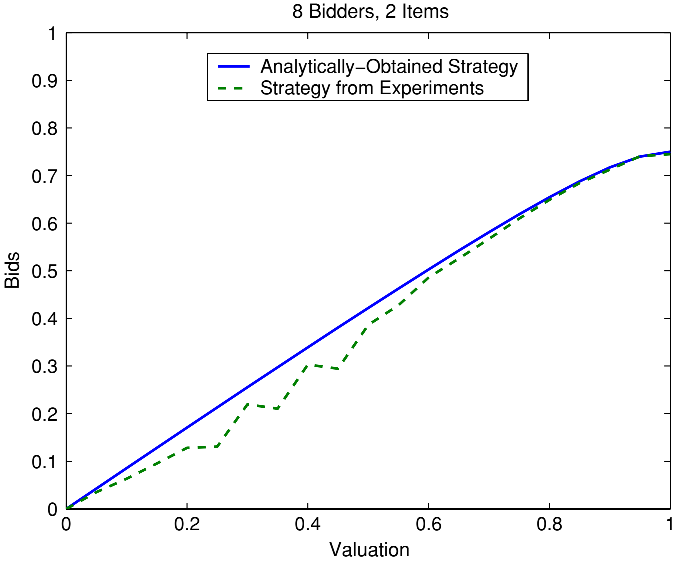

Fig. 2 Strategy comparison of experiment one

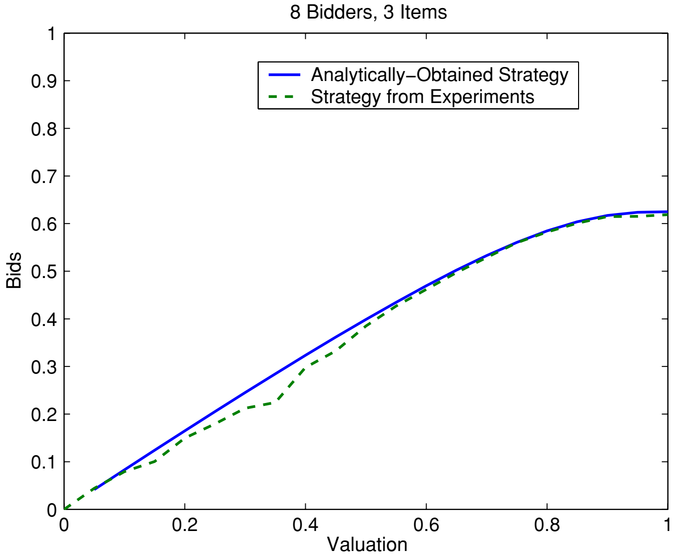

Fig. 3 Strategy comparison of experiment two

best revenue. On the other hand, by using the revenue as the seller's fitness function, and by not requiring bidder strategies to be monotonic increasing in Case 2, the GA simulations do enable the implicit comparison of all revenues. It so happens that the final equilibrium strategy obtained is one that is monotonic increasing.

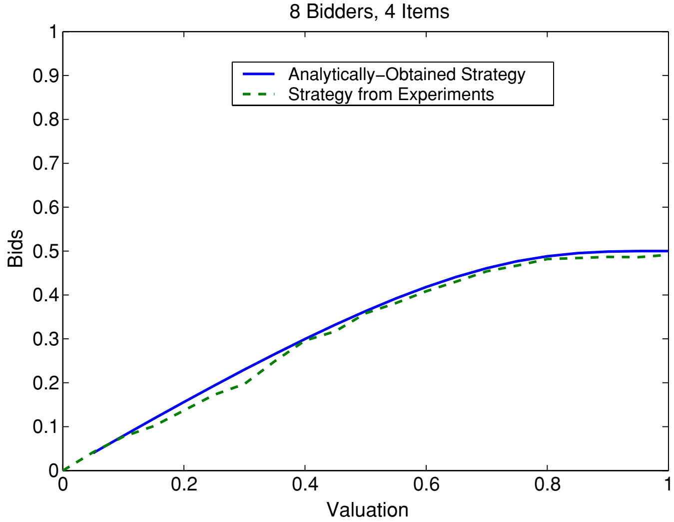

Fig. 4 Strategy comparison of experiment three

*5.3.2 Experiments Four, Five and Six—*500 *Sellers,* 8 *Bidders,* $k = 5, 6$ *and* 7 *Items*

In some of the following experiments, the high variation among initial results indicated the possibility of the existence of mixed strategies. For the previous experiments, we had encoded bidders' chromosomes as a 20-entry look-up table and implemented linear interpolation among the 20 entries. We seek a more accurate approach in the presence of mixed strategies—because a single bidder's lookup table could have a value from one strategy in one position and that of another strategy in another; linear interpolation would not give a correct value of either of the two mixed strategies. Therefore, in the following experiments, bidder valuations are rounded up to multiples of 0.05, and no linear interpolation is performed. We also increase the number of auctions to 1,000 per generation to reduce the variance.

In the experiment with $k = 5$ items, the seller's α chromosome converged to 0.0015. With $k = 6$ items, the seller's α chromosome does not converge.[4] For $k = 7$ items, the seller's α chromosome converged to 0.09. In all three experiments, the seller's P chromosome converged to 0.5. Figures 5, 6 and 7 provide graphs of the bidding function, averaged over the population of 4,000 bidders, for Case 2, obtained from the experiment with $k = 5, 6, 7$ items respectively, and the corresponding value of β^* for Case 1. While β^* is not expected to be similar to the results obtained with genetic simulations, as the games in Case 1 and Case 2 are distinct for $\alpha \neq 1$, the plots of β^* demonstrate that β^* is not always monotonic increasing. Observe that the average bidding functions for Experiments Five and Six ($k = 6, 7$, Figures 6 and 7) are considerably less smooth than those of the other experiments. We show in

[4] We computed seller revenue for values of α from 0 to 0.5, and confirmed that all values of revenue were close.

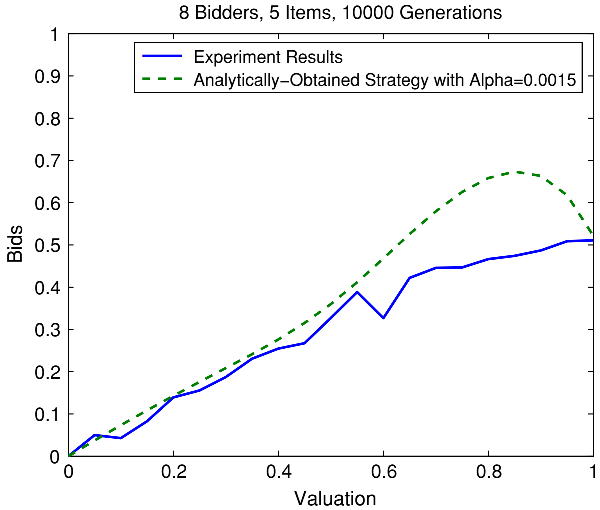

Fig. 5 Strategy comparison of experiment four

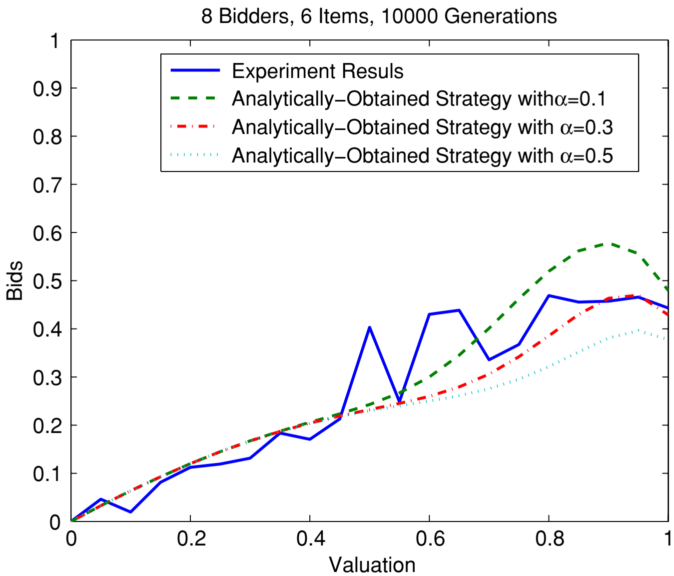

Fig. 6 Experiment five: comparison of β^* with average (unconverged) population strategy

Sect. 5.3.4 that the population variance for these cases is very high, and that it does not change with an increase in number of generations. This implies that the strategies of Figs. 6 and 7 are not the equilibrium bidding strategies for $k = 6$ and $k = 7$ respectively, and that bidding strategies are likely to be mixed strategies. Hence we need a different approach to parameterizing the strategies. Details on the study of mixed strategies may be found in Sect. 5.3.5.

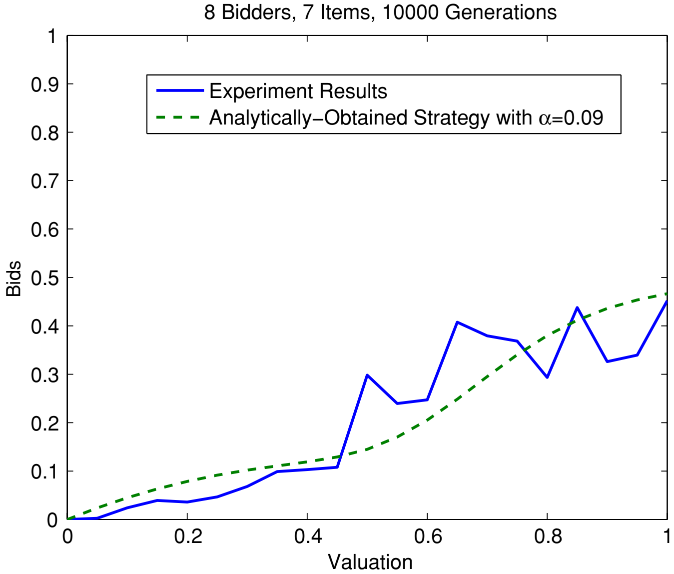

Fig. 7 Experiment six: comparison of β^* with average (unconverged) population strategy

*5.3.3 Experiment Seven—*500 *Sellers,* 8 *Bidders,* $k = 8$ *Items*

In the experiment with $k = 8$ items, the seller's α chromosome converges to $4.9e-07$, and the P chromosome to 0.501. Figure 8 provides a graph of the bidding function, averaged over the population of 4,000 bidders, for Case 2, obtained from the experiment with $k = 8$ items, and the corresponding value of β^* for Case 1. As with Experiments Five and Six, the average bidding function is considerably less smooth than those of Experiments One-Four. We examine this further in Sect. 5.3.4.

5.3.4 Variation Among Strategies

In this section we study the variation in the bidding strategy. The strategies plotted in Figs. 2, 3, 4, 5, 6, 7 and 8 are average strategies over the bidder population at the end of the experimental simulation. We observe that the variance about the average strategy, across the population of 4000 bidders, is larger for $k = 6, 7, 8$ than it is for other values of k. See Fig. 9 for some examples; the value of the average variance over the bidder's chromosome is 0.001353 for $k = 5$ and 0.001892, 0.003665 and 0.004679 for $k = 6, 7, 8$ respectively. We also observe that, for these values of k, the variance decreases sharply from 5,000 generations to 6,000 generations, but that it does not decrease appreciably as the number of generations increases thereafter; Figure 10 provides an example plot of the variance for $k = 7$, where the average variance over the bidder's chromosome is 0.01469, 0.004145, 0.003951 and 0.003665, for 5,000, 6,000, 8,000 and 10,000 respectively. To ensure that this particular result is not anomalous, we carry out the experiment a few times, and obtain similar differences in population variance each time.

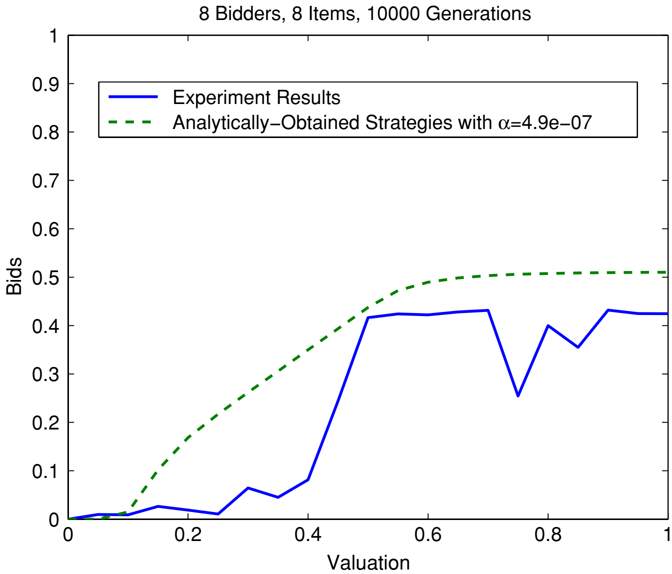

Fig. 8 Experiment seven: comparison of β^* with average(unconverged) population strategy

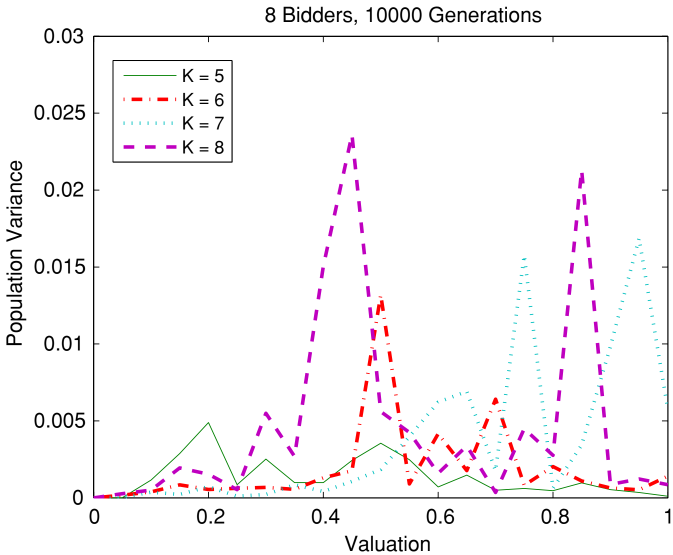

Fig. 9 Variance in strategies across the bidding population; $k = 5, 6, 7, 8$

As the variance is larger than that of other values of k, as it does not decrease appreciably with an increase in number of generations, and as these results are obtained for different instances of the experiment, we conclude that the algorithm does not converge for $k = 6, 7, 8$. This is likely due to the fact that the equilibrium strategy for

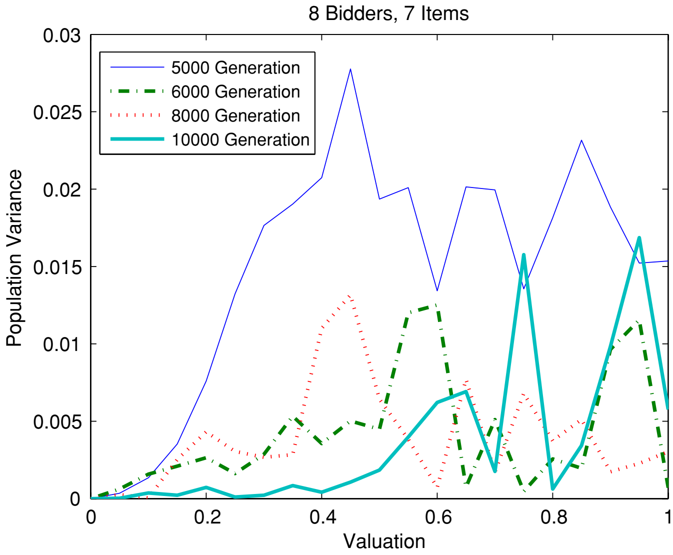

Fig. 10 Variance in strategies across the bidding population; $k = 7$; 5,000–10,000 generations

these values of k is not captured accurately by our chromosome encoding scheme. In the next section, we describe the encoding of a mixed strategy.

5.3.5 Mixed Strategies

In this section we describe the preliminary results we obtain by changing the design of the bidder chromosome in order to allow mixed strategies; the detailed study of mixed bidder strategies is outside the scope of this paper. We attempt to capture a (mixed) bidding strategy of the form:

$$Pr[\beta(x) = b] = p(x, b)$$

where, for a fixed value of x, $p(x, b)$ is non-zero for a finite number of values of b. That is, a bidder with valuation x_0 bids one of a finite number of values b with probability $p(x_0, b)$. Because $p(x, b)$ represents a probability,

$$\sum_{b|p(x_0,b)\neq 0} p(x_0, b) = 1 \quad \forall\, x_0$$

That is, for any given valuation x_0, the total of the probabilities over all possible bids is unity.

We do not restrict the possible strategies by parameterizing them; instead, we sample the bid and valuation spaces. Thus, we limit ourselves to the case where the bidder has one of 11 possible valuations: $0.1i$ for $i = 0, 1, ..., 11$. The bids are also one of 11 possible values: $0.1j$ for $j = 0, 1, ..., 11$. The bidder chromosome is encoded as an

Table 1 Experiment result: strategy matrix for $k = 6$

	$b=0$	$b=0.1$	$b=0.2$	$b=0.3$	$b=0.4$	$b=0.5$	$b=0.6$	$b=0.7$	$b=0.8$	$b=0.9$	$b=1.0$
$v=0$	1.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
$v=0.1$	0.1000	0.9000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
$v=0.2$	0.0658	0.0504	0.8838	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
$v=0.3$	0.0661	0.0555	0.0940	0.7845	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
$v=0.4$	0.0941	0.1264	0.2190	0.0926	0.4679	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
$v=0.5$	0.0794	0.0634	0.2091	0.2075	0.0840	0.3567	0.0000	0.0000	0.0000	0.0000	0.0000
$v=0.6$	0.0539	0.0645	0.1071	0.1239	0.0605	0.3975	0.1926	0.0000	0.0000	0.0000	0.0000
$v=0.7$	0.0663	0.0639	0.1390	0.0507	0.0864	0.4959	0.0456	0.0523	0.0000	0.0000	0.0000
$v=0.8$	0.0604	0.0625	0.0416	0.0664	0.0830	0.6037	0.0552	0.0087	0.0186	0.0000	0.0000
$v=0.9$	0.0748	0.0633	0.0491	0.0570	0.1116	0.5660	0.0450	0.0131	0.0106	0.0096	0.0000
$v=1.0$	0.0356	0.0582	0.0548	0.0581	0.0905	0.6465	0.0317	0.0088	0.0080	0.0038	0.0039

Table 2 Experiment result: strategy matrix for $k = 7$

	$b=0$	$b=0.1$	$b=0.2$	$b=0.3$	$b=0.4$	$b=0.5$	$b=0.6$	$b=0.7$	$b=0.8$	$b=0.9$	$b=1.0$
$v=0$	1.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
$v=0.1$	0.0676	0.9324	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
$v=0.2$	0.0724	0.1394	0.7881	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
$v=0.3$	0.0934	0.1855	0.0579	0.6632	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
$v=0.4$	0.0491	0.0493	0.1066	0.0939	0.7011	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
$v=0.5$	0.1078	0.1379	0.2506	0.1308	0.0968	0.2760	0.0000	0.0000	0.0000	0.0000	0.0000
$v=0.6$	0.0530	0.0623	0.1492	0.1384	0.2167	0.2245	0.1559	0.0000	0.0000	0.0000	0.0000
$v=0.7$	0.0237	0.0698	0.0466	0.0600	0.0420	0.6764	0.0226	0.0589	0.0000	0.0000	0.0000
$v=0.8$	0.0264	0.0486	0.0591	0.0247	0.1079	0.6580	0.0226	0.0144	0.0382	0.0000	0.0000
$v=0.9$	0.0726	0.0650	0.0499	0.0841	0.1085	0.5163	0.0648	0.0098	0.0157	0.0134	0.0000
$v=1.0$	0.0659	0.0786	0.1148	0.0730	0.0869	0.5025	0.0408	0.0123	0.0110	0.0034	0.0107

11×11 matrix. The $(i, j)^{th}$ entry in the matrix is the probability that the bid is $0.1j$ when the valuation is $0.1i$. That is, if $A(i, j)$ represents the strategy,

$$A(i, j) = Pr[\beta(0.1i) = 0.1j] = p(0.1i, 0.1j)$$

and

$$\sum_{j=0}^{j=11} A(i, j) = 1 \quad \forall i = 0, 1, \ldots, 11$$

The results we obtain for 5,000 generations are listed in Tables 1, 2 and 3. For the seller, $\alpha = 0.005$, $P = 0.59$ for $k = 6$; and $P = 0.59$ and $\alpha = 0.0001$ for $k = 8$. When $k = 7$, $P = 0.55$, and α converges to 0 for a small number of generations, but

Table 3 Experiment result: strategy matrix for $k = 8$

	$b=0$	$b=0.1$	$b=0.2$	$b=0.3$	$b=0.4$	$b=0.5$	$b=0.6$	$b=0.7$	$b=0.8$	$b=0.9$	$b=1.0$
$v=0.0$	1.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
$v=0.1$	0.0595	0.9405	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
$v=0.2$	0.0774	0.1316	0.7910	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
$v=0.3$	0.0732	0.1566	0.0778	0.6924	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
$v=0.4$	0.0715	0.0611	0.1863	0.1011	0.5800	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
$v=0.5$	0.1304	0.1559	0.1298	0.1176	0.0793	0.3870	0.0000	0.0000	0.0000	0.0000	0.0000
$v=0.6$	0.0561	0.0707	0.0992	0.0877	0.0608	0.5166	0.1090	0.0000	0.0000	0.0000	0.0000
$v=0.7$	0.0560	0.0918	0.0413	0.0491	0.0783	0.6168	0.0232	0.0435	0.0000	0.0000	0.0000
$v=0.8$	0.0694	0.0696	0.0791	0.0681	0.1066	0.5479	0.0226	0.0127	0.0240	0.0000	0.0000
$v=0.9$	0.1084	0.0847	0.1198	0.0686	0.1091	0.4408	0.0213	0.0227	0.0065	0.0180	0.0000
$v=1.0$	0.0451	0.1036	0.0634	0.0783	0.0886	0.5731	0.0199	0.0149	0.0064	0.0022	0.0046

then appears to diverge again, to $\alpha = 0$ and to $\alpha = 0.5$. A representative result for the bidding strategy is plotted in Fig. 11, for $k = 8$ and valuations 0.2, 0.5, 0.7 and 1.0. We observe that strategies for low valuations are closer to deterministic; bidders bid valuation with high probability, and other values with small probability. On the other hand, strategies for higher valuations are more mixed, and bidders bid 0.5 with higher probability and other values below valuation with equal probability; this is more pronounced for larger values of k. This may be explained as follows. Because the number of items is large, and $\alpha = 0$ or α small, and $P \approx 0.6$, bidders with high valuations are more likely than not to obtain the item in the second stage at a fixed price of $P \approx 0.6$. Hence the first stage consists of attempts by the high valuation bidders to obtain the item for a price smaller than P; this effect is enhanced as the value of k increases. In addition to attempting to obtain the item, bidders are also willing to confuse the seller for large values of k (and hence to use mixed strategies); the decreased competition enhances the value of privacy, and the willingness of high valuation bidders to bid lower. Bidders with valuation smaller than P do not, however, hope to get the item in the second stage, and are hence not willing to bid low, nor to attempt to protect their privacy by using mixed strategies.

The population variance at 5,000 generations is also listed in Tables 4, 5 and 6. The fractional standard deviation is smaller than that of the corresponding pure strategies for the same number of generations.

5.3.6 Modified Fitness Function for Experiments One, Two and Three

Recall that the bidding strategies obtained in Experiments One, Two and Three ($k = 2, 3, 4$—Sect. 5.3.1) were very similar to β^*. These results were, however, closer to β^* for high valuations than for low valuations (see Figs. 2, 3 and 4). A reasonable explanation is as follows. The bidders' fitness function is the sum of the payoff over several successful auctions; hence it is equivalent to the sum of payoffs weighted by

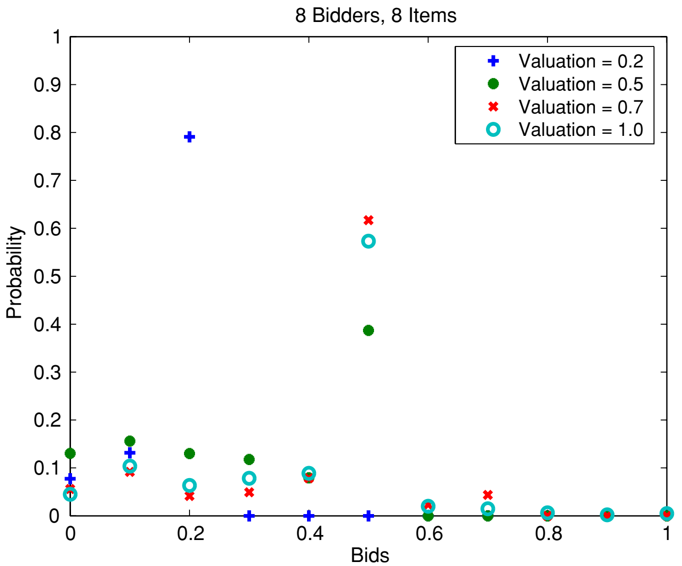

Fig. 11 Mixed strategy $k = 8$, $x = 0.2, 0.5, 0.7, 1.0$

Table 4 Population variance for $k = 6$

	$b = 0$	$b = 0.1$	$b = 0.2$	$b = 0.3$	$b = 0.4$	$b = 0.5$	$b = 0.6$	$b = 0.7$	$b = 0.8$	$b = 0.9$	$b = 1.0$
$v = 0.0$	0.0123	–	–	–	–	–	–	–	–	–	–
$v = 0.0$	0.0081	0.0093	–	–	–	–	–	–	–	–	–
$v = 0.0$	0.0052	0.0062	0.0127	–	–	–	–	–	–	–	–
$v = 0.0$	0.0052	0.0040	0.0085	0.0123	–	–	–	–	–	–	–
$v = 0.0$	0.0113	0.0152	0.0169	0.0113	0.0077	–	–	–	–	–	–
$v = 0.0$	0.0090	0.0116	0.0226	0.0205	0.0108	0.0065	–	–	–	–	–
$v = 0.0$	0.0054	0.0058	0.0105	0.0140	0.0101	0.032	0.0048	–	–	–	–
$v = 0.0$	0.0058	0.0080	0.0119	0.0066	0.0123	0.0325	0.0085	0.0018	–	–	–
$v = 0.0$	0.0035	0.0036	0.0067	0.0110	0.0108	0.0306	0.0061	0.0012	0.00070	–	–
$v = 0.0$	0.0064	0.0058	0.0076	0.0064	0.0118	0.0242	0.0070	0.0021	0.0010	0.00040	–
$v = 0.0$	0.0032	0.0053	0.0051	0.0077	0.0101	0.0245	0.0063	0.0012	0.0011	0.00010	0.00030

the probability of a win. Because α is close to unity, the bidders' fitness function may be denoted as $(x - b)[G(x) + H(x)]$. The higher a valuation, the higher the probability of a win $(G(x) + H(x))$, the larger the contribution of $x - b$ to the payoff, the greater the sensitivity of the payoff to the correct value of b for that value of x, and the more accurate the value of b in the converged bidding strategies. One may undo the effect of the weightage by adding not the payoff over several auctions, but the payoff weighted by the inverse of $G(x) + H(x)$. This would result in a fitness function which is $\sum_{b=b_1}(x - b)$. Figure 12 shows the preliminary results from an experiment that only conducts 200 auctions per generation after 500 generations.

Table 5 Population variance for $k = 7$

	$b=0$	$b=0.1$	$b=0.2$	$b=0.3$	$b=0.4$	$b=0.5$	$b=0.6$	$b=0.7$	$b=0.8$	$b=0.9$	$b=1.0$
$v=0.0$	0.0108	–	–	–	–	–	–	–	–	–	–
$v=0.0$	0.0058	0.0149	–	–	–	–	–	–	–	–	–
$v=0.0$	0.0043	0.0062	0.0090	–	–	–	–	–	–	–	–
$v=0.0$	0.0112	0.0174	0.0067	0.010	–	–	–	–	–	–	–
$v=0.0$	0.0051	0.0070	0.0118	0.0133	0.0131	–	–	–	–	–	–
$v=0.0$	0.0146	0.0156	0.0344	0.0166	0.0119	0.0051	–	–	–	–	–
$v=0.0$	0.0054	0.0067	0.0185	0.0204	0.0246	0.0247	0.0036	–	–	–	–
$v=0.0$	0.0028	0.0062	0.0044	0.0073	0.0072	0.0315	0.0030	0.0022	–	–	–
$v=0.0$	0.0036	0.0036	0.0049	0.0034	0.0150	0.0310	0.0030	0.0023	0.0019	–	–
$v=0.0$	0.0064	0.0074	0.0078	0.0129	0.0138	0.0333	0.0093	0.0008	0.0021	0.0007	–
$v=0.0$	0.0065	0.0103	0.0114	0.0051	0.0077	0.0203	0.0053	0.0013	0.0014	0.0002	0.0017

Table 6 Population variance for $k = 8$

	$b=0$	$b=0.1$	$b=0.2$	$b=0.3$	$b=0.4$	$b=0.5$	$b=0.6$	$b=0.7$	$b=0.8$	$b=0.9$	$b=1.0$
$v=0.0$	0.0154	–	–	–	–	–	–	–	–	–	–
$v=0.0$	0.0053	0.0131	–	–	–	–	–	–	–	–	–
$v=0.0$	0.0074	0.0135	0.0126	–	–	–	–	–	–	–	–
$v=0.0$	0.0089	0.0120	0.0126	0.0103	–	–	–	–	–	–	–
$v=0.0$	0.0070	0.0089	0.0223	0.016	0.0086	–	–	–	–	–	–
$v=0.0$	0.0173	0.0219	0.0154	0.0168	0.0113	0.0080	–	–	–	–	–
$v=0.0$	0.0082	0.0076	0.0059	0.0103	0.0104	0.0292	0.0026	–	–	–	–
$v=0.0$	0.0049	0.0062	0.0042	0.0069	0.0110	0.0231	0.0050	0.0014	–	–	–
$v=0.0$	0.0065	0.0090	0.0089	0.0124	0.0271	0.0375	0.0020	0.0023	0.0013	–	–
$v=0.0$	0.0100	0.0078	0.0098	0.0093	0.0108	0.0227	0.0025	0.0016	0.0004	0.0015	–
$v=0.0$	0.0034	0.0112	0.0068	0.0084	0.0149	0.0329	0.0034	0.0017	0.0008	0.0002	0.0009

We conduct a numerical analysis with the bidding function obtained from experiments with the modified fitness function. We first show that any different combination of P and α does not increase total revenue (see Fig. 13, which provides average revenue over 10,000 auctions).

In Figs. 14 and 15, we show that neither random over-bidding nor random under-bidding with any percentage range increases bidder payoff. Bidder No. 1 over-bids and under-bids while bidder Nos. 2 to 8 bid according to the bidding function. We do not show that no other strategy is better than the obtained strategies; however, our results lead us to believe more strongly that we have obtained Nash equilibria using fewer auctions per generation and fewer generations, than with the expected payoff fitness function. As the modified fitness function provides results similar to the original fitness function in fewer generations, it would be very useful in experiments with a larger number of bidders, if one believed that $\alpha = 1$.

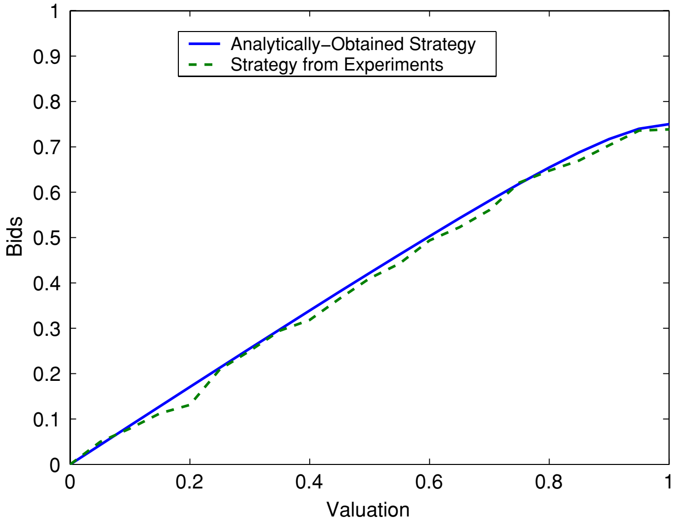

Fig. 12 Bidding function comparison of modified fitness function when $k = 2$

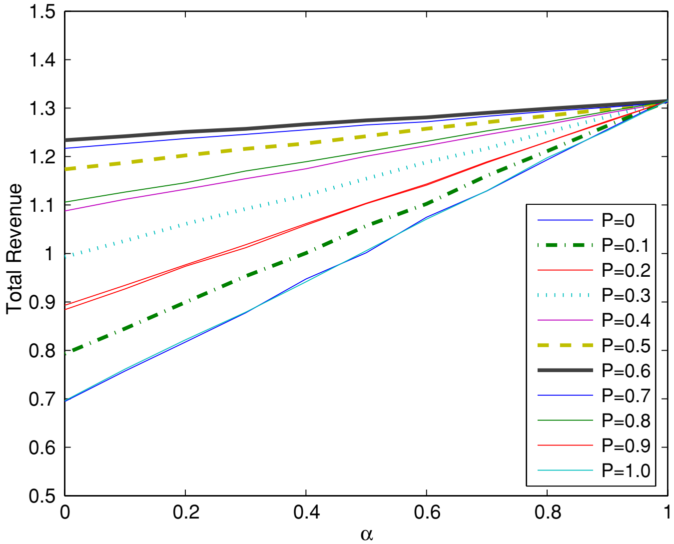

Fig. 13 Total revenue comparison for different P and α

5.4 Experimental Results: Summary

We summarize our experimental results as follows, in the 8-bidder-2-stage game of Case 2:

1. $k = 2, 3, 4$: The seller's best response is to price discriminate. Because items are scarce, bidders will not lower bids significantly in response to price discrimination. The bidder's best response is pure and monotonic increasing and consistent

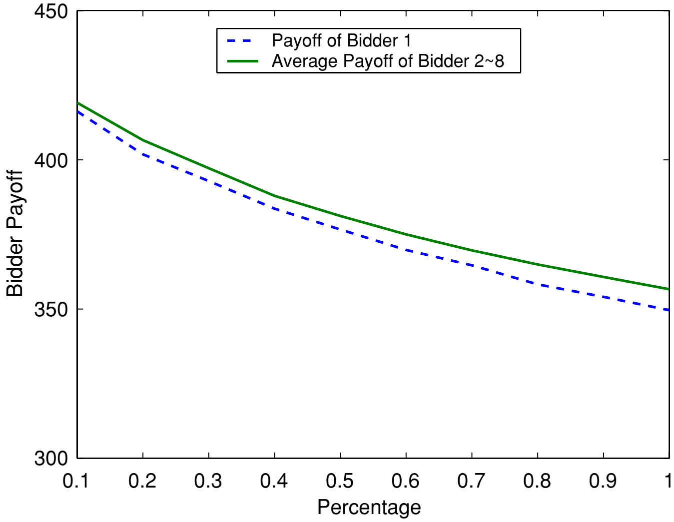

Fig. 14 Bidder payoff with over-bidding

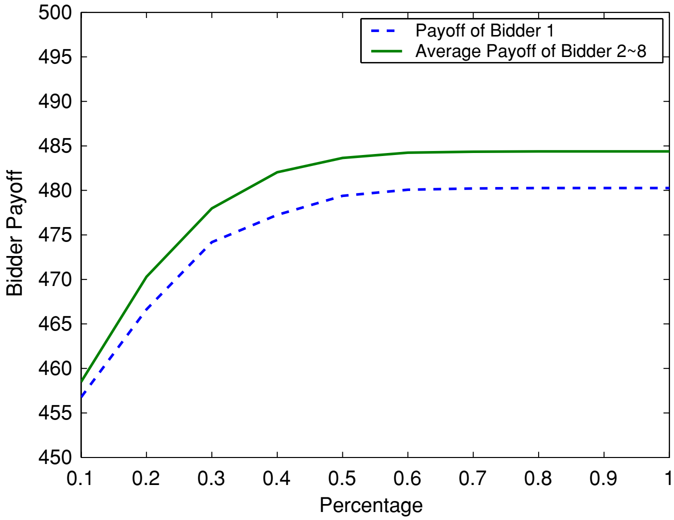

Fig. 15 Bidder payoff with under-bidding

with the analytically-derived strategy of Case 1 (the two games of Case 1 and Case 2 are identical for $\alpha = 1$).

2. $k = 5$: The seller's optimal strategy is to always make uniform-price offers. The availability of items makes the seller reluctant to price discriminate, as bidders would be more willing to bid low. The (pure) bidding strategy is not monotonic increasing.

3. $k = 7$: Preliminary results indicate that the seller's optimal strategy is to make a second chance offer with probability $\alpha = 0$ or $\alpha = 0.5$, and that bidder strategies are mixed.
4. $k = 6, 8$: Preliminary results indicate that the seller's best response is to always make uniform-price offers, as any indication of price discrimination reduces bids. Further, bidder strategies are mixed.

It is clear that different degrees of item scarcity make a difference for the seller's choice between the second chance offer and the uniform price. Further, a greater number of available items leads to the presence of mixed strategies, and average equilibrium strategies are not always monotonic increasing.

6 Conclusions and Future Work

From the GA experiments for the two-stage price discrimination game, we have the following observations:

1. When $k > \frac{n}{2}$, and bidders are aware that there is no scarcity, the seller's optimal strategy is to always offer a fixed price. Bidding strategies obtained are not always pure or monotonic increasing.
2. When $k \leq \frac{n}{2}$ it is the seller's best response to always choose a second chance offer in Stage II. That is, when items are scarce, the seller need not worry about bidders bidding low, as competition among bidders will allow the seller to be successful with price discrimination. Because of the scarcity, bidders do not lower their bids sufficiently. Further, equilibrium bidding strategies that are pure and monotonic increasing exist.

Interesting questions include the further characterization of mixed strategies, and whether similar behavior would be observed in other auctions (second-price, for example) and in auctions with goods where bidder valuations are not completely private.

Acknowledgments The authors would like to thank Sumit Joshi for several valuable conversations, and an anonymous reviewer for several valuable suggestions.

Appendix

Proof: Theorem 1

Proof Theorem 1
The expected payoff due to bid b, when all others are bidding according to strategy β, is, from (1):

$$E[\Pi(b,x)] = \begin{cases} G(\beta^{-1}(b))(x-b)+H(\beta^{-1}(b))[\alpha(x-b)+(1-\alpha)(x-P)] & x > P \\ G(\beta^{-1}(b))(x-b)+H(\beta^{-1}(b))[\alpha(x-b)] & x \leq P \end{cases} \tag{2}$$

To find the bidding strategy β^* that maximizes the expected payoff, we differentiate (2) wrt b, equate to zero, and, assuming a symmetric strategy among bidders, replace b with $\beta^*(x)$. This gives us:

$$\alpha h(x)\beta^*(x) + \alpha\beta^{*\prime}(x)H(x) + \beta^*(x)g(x) + G(x)\beta^{*\prime}(x)$$
$$= \begin{cases} xg(x) + \alpha x h(x) + (1-\alpha)(x-P)h(x) & x > P \\ xg(x) + \alpha x h(x) & x \le P \end{cases}$$

Integrating both sides wrt x gives:

$$\alpha\beta^*(x)H(x) + \beta^*(x)G(x)$$
$$= \begin{cases} \int_0^x yg(y)dy + \alpha \int_0^x yh(y)dy + (1-\alpha)\int_P^x h(y)(y-P)dy & x > P \\ \int_0^x xg(x)dx + \alpha\int_0^x xh(x)dx & x \le P \end{cases}$$

Solving for β^* gives:

$$\beta^*(x) = \begin{cases} \frac{\int_0^x yg(y)dy + \alpha\int_0^x yh(y)dy + (1-\alpha)\int_P^x h(y)(y-P)dy}{G(x)+\alpha H(x)} & x > P \\ \frac{\int_0^x yg(y)dy + \alpha\int_0^x yh(y)dy}{G(x)+\alpha H(x)} & x \le P \end{cases} \tag{3}$$

When x is uniformly distributed, $G(x) = x^{N-1}$; and

$$\frac{\int xG'(x)dx}{G(x)} = \frac{N-1}{N}x$$

may be substituted in (3).

We now show that unilateral deviation by a single bidder does not provide benefit. Clearly, there is no value in submitting a bid greater than the highest value of $\beta^*(x)$ because the bidder can certainly win with a bid that is equal to this highest value. Similarly, there is no value in submitting a bid lower than the lowest value of $\beta(x)$, as this bid will certainly not be a winning bid. Hence a deviating bidder will only provide a bid from the range of β^*.

Suppose a bidder with valuation x bids $\beta^*(z)$, $z \neq x$. We consider cases when both x and z are smaller than and greater than P, and, also when $x < P < z$ and $z < P < x$.

Consider the case when $x < P$ and $z < P$. The bidder's payoff is:

$$\Pi(\beta^*(z), x) = G(z)(x - \beta^*(z)) + H(z)[\alpha(x - \beta^*(z))] \tag{4}$$

Further,

$$\Pi(\beta^*(x), x) = G(x)(x - \beta^*(x)) + H(x)[\alpha(x - \beta^*(x))] \tag{5}$$

To show that bidders do not have an incentive to unilaterally deviate, we need to show that $\Pi(\beta^*(x), x) - \Pi(\beta^*(z), x) \geq 0$ for $x \leq z$, and for $x > z$.

Substituting (3) in (5) and (4) gives:

$$E[\Pi(\beta^*(x), x)] = \int_0^x G(y)dy + \alpha \int_0^x H(y)dy \tag{6}$$

and

$$E[\Pi(\beta^*(z), x)] = G(z)(x - z) + \alpha H(z)(x - z) + \int_0^z G(y)dy + \alpha \int_0^z H(y)dy \tag{7}$$

Subtracting (7) from (6) gives:

$$\begin{aligned}
&\Pi(\beta^*(x), x) - \Pi(\beta^*(z), x) \\
&\quad = (z - x)\,(G(z) + \alpha H(z)) - \int_x^z (G(y) + \alpha H(y))\,dy \\
&\quad = (z - x)\,[(1 - \alpha)G(z) + \alpha K(z)] - \int_x^z [(1 - \alpha)G(y) + \alpha K(y)]\,dy
\end{aligned}$$

where $K(x)$ is the probability that x is among the k highest bidders, $K(x) = G(x) + H(x)$. The expression above is non-negative, by an argument similar to that provided for the derivation of the equilibrium bidding strategy in a classical first-price auction (Krishna 2002), for both $x \geq z$ and $x \leq z$, because both $G(x)$ and $K(x)$ are monotonic increasing.

We now examine the case for $x \geq P$ and $z \geq P$. The expected payoff for bidding $\beta^*(z)$, when $x \geq P$ is:

$$E[\Pi(\beta^*(z), x)] = G(z)(x - \beta^*(z)) + H(z)[\alpha(x - \beta^*(z)) + (1 - \alpha)(x - P)] \tag{8}$$

Further,

$$E[\Pi(\beta^*(x), x)] = G(x)(x - \beta^*(x)) + H(x)[\alpha(x - \beta^*(x)) + (1 - \alpha)(x - P)]$$

Again, we need to show that $E[\Pi(\beta^*(x), x)] - E[\Pi(\beta^*(z), x)] \geq 0$, whether $x > z$ or $z > x$. Again, by substituting the expression for $\beta^*(x)$, we obtain:

$$\begin{aligned} E[\Pi(\beta^*(x), x)] &= \int_0^x G(y)dy + \alpha \int_0^x H(y)dy + (1-\alpha)[(x-P)H(x) \\ &\quad - \int_P^x (y-P)h(y)dy] \\ &= \int_0^x G(y)dy + \alpha \int_0^x H(y)dy + (1-\alpha)\int_P^x H(y)dy \end{aligned} \tag{9}$$

Similarly,

$$\begin{aligned} E[\Pi(\beta^*(z), x)] &= G(z)(x-z) + \alpha H(z)(x-z) + \int_0^z G(y)dy + \alpha \int_0^z H(y)dy \\ &\quad + (1-\alpha)\left[H(z)(x-P) - \int_P^z (y-P)h(y)dy\right] \end{aligned}$$

or

$$\begin{aligned} E[\Pi(\beta^*(z), x)] &= K(z)(x-z) + \int_0^z G(y)dy + \alpha \int_0^z H(y)dy \\ &\quad +(1-\alpha)\int_P^z H(y)dy \end{aligned} \tag{10}$$

Hence

$$E[\Pi(\beta^*(x), x)] - E[\Pi(\beta^*(z), x)] = (z-x)K(z) - \int_x^z K(y)dy \tag{11}$$

Again, the above expression is non-negative because $K(x)$ is monotonic increasing.

Now we examine the case when $x < P < z$.

Using (10) and (6) we obtain:

$$E[\Pi(\beta^*(x), x)] - E[\Pi(\beta^*(z), x)] = (z - x)K(z) - \int_x^z K(y)dy + (1 - \alpha)\int_0^P H(y)dy \geq 0$$

Finally, we examine the case when $z < P < x$.

Using (7) and (9) we obtain:

$$E[\Pi(\beta^*(x), x)] - E[\Pi(\beta^*(z), x)] = (x - z)[G(z) + \alpha H(z)] + \int_0^z [G(y) + \alpha H(y)]dy + (1 - \alpha)\int_P^x H(y)dy \geq 0$$

by arguments as with the case $z \leq P$ and $x \leq P$.

Thus we have shown that bidders do not have an incentive to deviate from β^* when it is the strategy used by all other bidders. Hence β^* is an equilibrium strategy if $0 \leq \beta^* \leq x$ and β^* is monotonic increasing. □

Mixed Bidding Strategies for $k = 6, 7, 8$

We tabulate the mixed bidding strategies obtained for $k = 6, 7, 8$, and the corresponding population variance.

References

Andreoni, J., & Miller, J. H. (1995). Auctions with artificial adaptive agents. *Games and Economic Behavior, 10*(1), 39–64.

Arifovic, J. (1994). Genetic algorithm and the Cobweb model. *Journal of Economic Dynamics and Control, 18*(1), 3–28.

Axelrod, R. (1984). *Evolution of cooperation*. New York: Basic Books.

Axelrod, R. (1987). The evolution of strategies in the iterated prisoner's dilemma. In L. Davis (Ed.), Genetic algorithms and simulated annealing (Chapter 3, pp. 32–41). Los Altos, CA: Morgan Kaufman.

Bajari, P., & Hortaçsu, A. (2003). Winner's curse, reserve prices and endogenous entry: Empirical insights from eBay auctions. *RAND Journal of Economics, 34*(2), 329–355.

Byde, A. (2003). *Applying evolutionary game theory to auction mechanism design*. In Proc. of IEEE International Conference on Electronic Commerce (CEC03).

Cliff, D. (2003). Explorations in evolutionary design of online auction market mechanisms. *Journal of Electronic Commerce Research and Applications, 2*(2), 162–175.

Cliff, D. (2007). ZIP60: Further explorations in the evolutionary design of trader agents and online auction-market mechanisms. http://eprints.ecs.soton.ac.uk/14078/. Last modified Apr 13, 2008.

Dawid, H. (1996). Genetic algorithms as a model of adaptive learning in economic systems. *Central European Journal for Operations Research and Economics, 4*(1), 7–23.

Dawid, H. (1999). *Adaptive learning by genetic algorithms: Analytical results and applications to economic models*. Berlin: Springer.

Dawid, H., & Kopel, M. (1998). On economic applications of the genetic algorithm: A model of the Cobweb type. *Journal of Evolutionary Economics, 8*(3), 297–315.

Duffy, J., & Utku, M. U. (2008). Internet auctions With artificial adaptive agents: A study on market design. *Journal of Economic Behavior & Organization, 67*(2), 394–417.

Goeree, J. K., & Offerman, T. (2002). Efficiency in auctions with private and common values: An experimental study. *American Economic Review, 92*(3), 625–643.

Joshi, S., Sun, Y. A., & Vora, P. L. (2005). *The privacy cost of the second chance offer*. In Proc. of the ACM workshop on privacy in the electronic society (WPES 05). Alexandria, VA: ACM.

Krishna, V. (2002). *Auction theory*. New York: Academic Press.

Muhlenbein, H., & Schlierkamp-Voosen, D. (1993). Predictive models for the Breeder genetic algorithm: I. continuous parameter optimization. *Evolutionary Computation, 1*(1), 25–49.

Ockenfels, A., & Roth, A. E. (2002). The timing of bids in Internet auctions: Market design, bidder behavior, and artificial agents. *AI Magazine, 23*, 79–88.

Riechmann, T. (2001). *Learning in economics: Analysis and application of genetic algorithms*. Heidelberg: Physica-Verlag.

Salmon, T. C., & Wilson, B. J. (2008). Second chance offers versus sequential auctions: Theory and behavior. *Economic Theory, 34*(1), 47–67.