

Adaptive Auctions: Learning to Adjust to Bidders

David Pardoe¹, Peter Stone¹, Maytal Saar-Tsechansky², and Kerem Tomak²

The University of Texas at Austin, Austin TX 78712, USA

¹Department of Computer Sciences, ²McCombs School of Business

{dpardoe, pstone}@cs.utexas.edu, {Maytal.Saar-Tsechansky, Kerem.Tomak}@mcombs.utexas.edu

Abstract: Auction mechanism design has traditionally been a largely analytic process, relying on assumptions such as fully rational bidders. In practice, however, bidders behave unpredictably, making them difficult to model and complicating the design process. To address this challenge, we present an adaptive auction mechanism: one that *learns* to adjust its parameters in response to past empirical bidder behavior so as to maximize an objective function such as auctioneer revenue. In this paper, we give an overview of our general approach and then present an instantiation in a specific auction scenario. The algorithm is fully implemented and tested. Results indicate that the adaptive mechanism is able to outperform any single fixed mechanism.

1 Introduction

Recent years have seen the emergence of numerous auction platforms that cater to a variety of markets such as business to business procurement and consumer to consumer transactions. Depending on factors such as bidder strategies and product types, varying the parameters of the auction mechanism, such as auctioneer fees, minimum bid increments, and reserve prices, can lead to widely differing results. This paper considers *learning* auction parameters to maximize auctioneer revenue as a function of empirical bidder behavior.

Mechanism design has traditionally been largely an analytic process. Assumptions such as full rationality are made about bidders, and the resulting properties of the mechanism are analyzed in this context [1]. Typically, the design process is incremental, involving reevaluating the assumptions made about bidders in light of auction outcomes. In particular, these assumptions pertain to bidders' intrinsic properties and to the manner by which these properties are manifested in bidding strategies.

Even when the assumptions about bidders can be successfully modified to explain past results, the process requires human input and is time consuming, undermining the efficiency with which changes can be made to the mechanism. In e-commerce settings in which a large number of auctions for similar goods may be held within a short time frame, such as auctions on e-Bay or Google keyword auctions, this is a serious drawback.

To address these challenges, we propose a substantially different approach to mechanism design: self-adaptive auction mechanisms that change in response to observed bidder behavior. In this paper, we consider an auction with a single continuous parameter, and present a *metalearning* process by which the method of parameter optimization is itself parameterized and optimized based on simulated experiences with different *populations* of bidders. The main contribution of this paper is the specification, implementation, and empirical testing of an adaptive mechanism designed to maximize auctioneer revenue in the face of an unknown population of bidders.

2 An Adaptive Approach

The strategies employed in an auction by bidders are often unknown to the seller. Nonetheless, the effectiveness of the auction mechanism can vary drastically as a function of the bidding strategies used. In settings in which a large number of similar auctions are held, it may be reasonable to assume that the behavior of bidders remains somewhat consistent, suggesting the possibility of learning about bidder behavior through experience. For example, the bidders on a particular Google keyword may remain the same for some time, and identical items on eBay will likely attract similar buyers. For such settings, we propose *adaptive mechanism design*, an online empirical process whereby the mechanism adapts to maximize a given objective function based on observed outcomes. Because

we allow for situations in which bidder behavior cannot be predicted beforehand, this process must be performed online during interactions with real bidders. (In this paper, the term “online” refers to the fact that adaptation takes place during the course of actual auctions, and not the fact that auctions take place electronically – although that may also be the case.)

In our view of adaptive mechanisms, a parameterized mechanism is defined such that an adaptive method can be used to revise parameters in response to observed results of previous auctions, choosing the most promising parameters to be used in future auctions. Any number of continuous or discrete auction parameters may be considered, such as reserve prices, auctioneer fees, minimum bid increments, and whether the close is hard or soft.

The adaptive method is an online machine learning algorithm aiming to characterize the function from mechanism parameters to expected revenue (or any other objective function). Because the learner can choose different auction parameters at each step (thus effectively selecting its own training examples), and the target output is continuous, the problem is an active learning [2] regression problem. A key characteristic is that the learning is all done online during actual auctions, so that excessive exploration of various parameter settings can be costly.

The only assumption about bidders is that their behavior is consistent in some way (e.g. bidders associated with a particular industry tend to bid similarly) so that it is possible to learn to predict auction results as a function of the mechanism, at least in expectation.

The use of an adaptive mechanism provides the possibility of identifying optimal auction parameters even without explicitly modeling the bidders. However, when predictions can be made about the types of behavior to be expected, this knowledge can usefully influence the method of adaptation. Specifically, one can use a method of adaptation that is itself parameterized, and then choose the parameters that result in the best performance under expected bidder behavior.

The steps in the “metalearning” process of choosing an adaptive auction mechanism to maximize a particular objective function are thus as follows: 1) Choose the parameterization of the auction. 2) Make predictions about possible bidder behavior that allow for simulation. Sources for these predictions may include analytically derived equilibrium strategies, empirical data from past auctions in a similar setting, and learned behaviors. 3) Choose the method of adaptation and its parameters. 4) Search the space of parameters of the adaptive method to find those that best achieve the objective in simulation.

We now present an illustrative application of this approach to a particular auction scenario.

3 An Auction Scenario

We consider an English auction in which bidders submit ascending bids, and assume that the seller may set a *reserve price* indicating the minimum acceptable bid. For the sake of simplicity, we assume that two bidders participate in each auction. We base the behavior of these bidders on the model of *loss averse* bidders described by Dodonova [3]. A loss averse bidder considers the utility from a gain to be lower than the disutility from a loss (“losing” an item for which he previously had the highest bid) of the same magnitude. Specifically, if the marginal utility from winning an auction is x , then the marginal disutility from losing the same object is αx , where $\alpha > 1$.

Under the equilibrium derived by Dodonova, the first bidder will submit a bid in the beginning of the auction if his valuation is higher than the reserve price, while the second bidder enters the auction only if by doing so he can guarantee a positive expected utility. This equilibrium can cause the seller’s optimal reserve price to be 0 under certain conditions, and can also result in a non-convex revenue as a function of reserve price, with one maximum close to zero and another at a much higher reserve price, as will be illustrated in Figure 1. Thus the auctioneer has potential

incentives to set both a low reserve price and a high reserve price, a conflict that must be taken into account when choosing a method of searching for the optimal reserve price.

We consider a scenario in which a seller interacts repeatedly with bidders drawn from a fixed population. In particular, the seller has 1000 identical items that will be sold one at a time through a series of English auctions. The seller sets a (potentially different) reserve price for each auction, thus indirectly affecting the auction's outcome. The seller's goal is to set the reserve price for each auction so that the total revenue obtained from all the auctions is maximized. If a complete model of the behavior of the population of bidders were available, the seller could determine the optimal reserve price analytically. However, as this information is not available, the seller must identify the optimal reserve price through online experimentation guided by an adaptive mechanism.

A bidder is characterized by i) an independent, private value v for the sold item, and ii) a degree of loss-aversion α . The seller knows that bidders have independent, private values, and are likely loss averse, but does not know the actual distributions from which α and v are drawn, or the strategies bidders will employ. However, the seller assumes that the *population* of bidders (characterized in this case by distributions over valuations and α) does not change over time. Thus, the behavior exhibited by bidders will be the same for each auction *in expectation*, allowing the seller to draw inferences from past auction results.

Although the seller cannot completely characterize the bidder population, we assume that the seller can predict and simulate a possible *distribution over populations*. As an example of how a such a distribution might be generated, a seller introducing a new product to the market might identify similar items that have been sold in the past and observe the behavior of bidders on each item, treating each group as a distinct population. In our experiments, the seller simulates a bidder population as having Gaussian distributions over valuations and α values, chosen according to the distribution over populations. Further details are omitted due to space limitations; however, from the standpoint of the adaptive mechanism described in Section 4, it is only important that the seller is able to simulate a population drawn from this distribution and wants to find the adaptive parameters that give the best performance under this distribution.

To illustrate the task faced by the seller, we generated 10,000 bidder populations according to the seller's distribution, and found the average revenue for each reserve price between 0 and 1 at intervals of 0.01. The average revenue for each reserve is shown by the solid line in Figure 1. A reserve price of 0.54 yields the highest average revenue, 0.367. If we were required to select a single reserve price for the seller to use, we would chose this price. However, for each individual bidder population there is a distinct choice of reserve that yields the highest average revenue. In particular, the dotted line in Figure 1 shows the number of times that each reserve was optimal. Two important observations can be made: i) despite the variety in bidder populations, the optimal reserve price is frequently in one of two small regions (including near zero, as is expected with loss averse bidders); ii) nevertheless, most choices of reserve are optimal for some population. The second observation motivates our use of an adaptive mechanism – we want to identify the auction parameters that are best for the particular population that the seller actually encounters. Our goal in learning the parameters of the mechanism is to take advantage of the first observation – we can focus the mechanism's exploration on those auction parameter settings that appear most promising given the seller's beliefs about possible bidder populations.

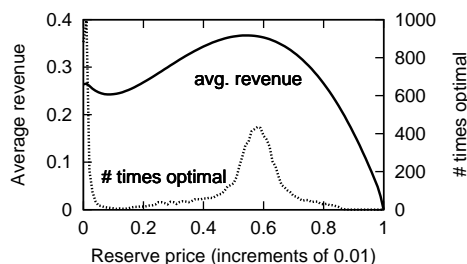


Fig. 1.

4 Implementation and Results

As specified at the end of Section 2, for the auction scenario with the goal of maximizing revenue over 1000 auctions, we have 1) chosen the auction parameterization (the reserve price represents a single, continuous parameter), and 2) the seller has provided a means of generating bidder behavior. In this section, we complete the remaining tasks of 3) specifying our adaptive method and its parameters, and 4) presenting a means of identifying the parameters that result in optimal performance. We then present the results of applying the approach described to the auction scenario.

4.1 Method of adaptation

We now describe an adaptive method that discretizes the problem by restricting the seller to choosing one of k choices for the reserve price at each step, where the i th choice is a price of $(i - 1)/(k - 1)$. (An extension of this method that does not require discretization is straightforward and appears promising, but we leave its discussion to future work.) The resulting problem can be viewed as an instance of the k -armed bandit problem, a classic reinforcement learning problem [4]. In such problems, the expected value of each choice is assumed to be independent, and the goal of maximizing the reward obtained presents a tradeoff between exploring the choices, in order to increase the knowledge of each one's result, and exploiting the choice currently believed to be best.

The approach to solving k -armed bandit problems that we use is sample averaging with softmax action selection using the Boltzmann distribution. In this approach, the average revenue for each choice, avg_i , is recorded, and at each step the probability of choosing i is $(e^{avg_i/\tau})/(\sum_{j=1}^k e^{avg_j/\tau})$, where τ represents a *temperature* determining the extent to which exploitation trumps exploration. The temperature is often lowered over time to favor increasing exploitation due to the fact that estimates of the result of each choice improve in accuracy with experience.

Softmax action selection has parameters controlling the temperature and controlling the initial estimates of each choice's reward. We vary the temperature throughout an episode by choosing starting and ending temperatures, τ_{start} and τ_{end} , and interpolating linearly. To calculate the average revenue for each choice, we require for each choice a record of both the average revenue, avg_i , and the number of times that choice has been tried, $count_i$. Although the straightforward approach would be to initialize the averages and counts to zero, one common technique, known as *optimistic initialization* [4] is to set all initial averages to a value higher than the predicted value of the largest possible revenue. Each choice is therefore likely to be explored at least once near the beginning of the episode. We employ a variation on this technique in which we choose values for the averages and counts that encourage heavy initial exploration of those choices believed most likely to be optimal given the predictions of bidder behavior. For instance, if the revenue from a particular choice is expected to be high on average but have a high variance, assigning a high initial count and average to that choice would ensure that it is explored sufficiently: several trials resulting in low revenue would be needed to significantly lower the computed average. This approach amounts to starting out with what we will call *initial experience*. The choice of initial experience and temperatures are made by the search procedure we will describe shortly. Thus for a given choice of k , this will be a search over $2k + 2$ parameters (including τ_{start} and τ_{end}).

4.2 Parameter search

Now that we have chosen a method of adaptation and have a means of generating bidder behavior, we are ready to search for the set of parameters that results in the best expected performance. For any given set of parameters, we can obtain an estimate of the expected revenue from an episode by generating a population of bidders and running an episode using those parameters. This estimate

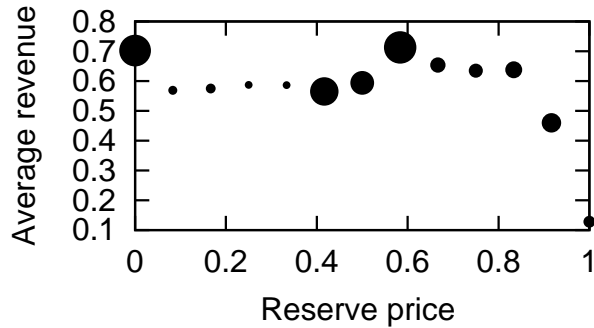


Fig. 2. Learned parameters - $\tau_{start} = .0423$, $\tau_{end} = .0077$

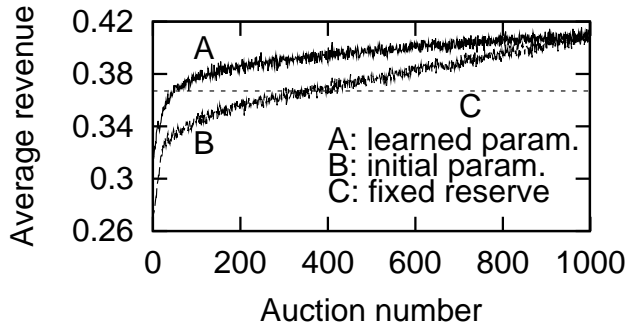


Fig. 3. Average revenue per auction over the course of an episode for each method.

will be highly noisy, due to the large number of random factors involved in the process, and so we are faced with a stochastic optimization task.

To solve this task, we use Simultaneous Perturbation Stochastic Approximation (SPSA) [5], a popular method of stochastic optimization based on gradient approximation. For initial parameters, we use a somewhat optimistic value of 0.6 for each avg_i and a value of 1 for each $count_i$. τ_{start} and τ_{end} are set to 0.1 and 0.01, respectively. Ideally, the parameter k would be part of the search process as well, but as our search method requires a fixed number of parameters, we have chosen what appears to be the best value after running searches with several values of k .

It should be noted that although this process of searching for the optimal parameters can be time consuming (in our experiments, a few hours were required), the process takes place in offline simulation before the actual auctions begin. When the adaptive method is applied during the actual auctions using the resulting parameters, each choice of a new reserve price takes only a small fraction of a second.

4.3 Results

To evaluate our adaptive method, we first searched for the best possible set of parameters, including k , as described above. We found that a value of 13 was optimal for k . The learned parameters are presented in Figure 2. Initial experience is displayed visually by plotting a circle for each avg_i with area proportional to $count_i$. The initial experience appears reasonable given Figure 1. The values of avg are mostly similar and fairly high, but the values of $count$ are much higher for the choices in the more promising regions. As a result, it will take longer for the computed average revenue of these choices to fall, and so these choices will be explored more heavily early in an episode.

We next generated a set of 10,000 bidder populations, and found the average revenue per episode using both the initial and the learned parameters. The average revenues per auction are shown in Figure 4, while a plot of the average revenue for each auction over an entire episode is shown in Figure 3. The average total revenue in each case is higher than the revenue resulting from using the best fixed reserve price, 0.54, indicating that the use of an adaptive mechanism is indeed worthwhile in this scenario. The difference observed between each pair of methods is statistically significant at the 99% confidence level according to paired t-tests comparing results for the same bidder population. From Figure 3 we can see that while both adaptive methods approach the same revenue by the last auction

Adaptive method	Total Revenue
best fixed reserve price (0.54)	0.367
adaptive, initial parameters	0.374
adaptive, learned parameters	0.394

Fig. 4. Average revenue per auction for each adaptive method.

in an episode, using learned parameters leads to much higher revenues during the early part of an episode. Thus, the learned parameters are effective at focusing initial exploration; providing sufficient initial experience to permit a higher initial degree of exploitation; or both.

5 Related Work

To our knowledge, only a few recent articles have begun to explore the subject of adapting auction mechanisms in response to bidder behavior. Cliff [6] and Phelps et al. [7] consider continuous double auctions, using genetic algorithms and genetic programming, respectively, to evolve both bidder strategies and auction rules. Byde [8] studies the space of auction mechanisms between the first and second-price sealed-bid auction, using a genetic algorithm to learn the bidders' strategies in response to different mechanisms. The primary difference between these previous approaches and the method advocated in this paper is that these approaches use simulation to produce fixed mechanisms, while our aim is to develop mechanisms that are self-adapting in an online setting.

The process of identifying the parameters of the adaptive mechanism can be viewed as an instance of *metalearning* [9]. In *metalearning*, the goal is to improve the performance of a learning system for a particular task through experience with a family of related tasks. In our case, the learning system is the adaptive mechanism, and the family of related tasks is the set of different bidder populations generated during simulation.

6 Conclusions and Future Work

In this paper, we have presented a novel approach to mechanism design. Instead of relying on analytical methods that depend on specific assumptions about bidders, our approach is to create a self-adapting mechanism that adjusts auction parameters in response to past auction results. We have analyzed and experimented with a specific auction scenario involving loss averse bidders and varying seller reserve prices. We have shown how information about potential bidder behavior can guide the selection of the method of adaptation and significantly improve auctioneer revenue.

There are several directions in which this work could be extended. Many auction parameters are available for tuning, ranging from bidding rules to clearing policies. The problem becomes more challenging in the face of multidimensional parameterizations.

Our on-going research agenda also includes examining the effects of including some adaptive bidders in the economies that are treated by adaptive mechanisms.

Acknowledgments

This research was supported in part by NSF CAREER award IIS-0237699 and NSF grant EIA-0303609.

References

1. Parkes, D.C.: Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency. PhD thesis, Department of Computer and Information Science, University of Pennsylvania (2001)
2. Saar-Tsechansky, M., Provost, F.: Active learning for class probability estimation and ranking. *Machine Learning* (2004)
3. Dodonova, A., Khoroshilov, Y.: Optimal auction design when bidders are loss averse. (Working Paper. University of Ottawa.)
4. Sutton, R.S., Barto, A.G.: *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA (1998)
5. Spall, J.C.: An overview of the simultaneous perturbation method for efficient optimization. *Johns Hopkins APL Technical Digest* **19** (1998) 482–492
6. Cliff, D.: Evolution of market mechanism through a continuous space of auction types. Technical Report HPL-2001-326, HP Labs (2001)
7. Phelps, S., Mc Burnley, P., Parsons, S., Sklar, E.: Co-evolutionary auction mechanism design. In: *Agent Mediated Electronic Commerce IV*. Volume 2531 of *Lecture Notes in Artificial Intelligence*. Springer Verlag (2002)
8. Byde, A.: Applying evolutionary game theory to auction mechanism design. In: *Proceedings of the 4th ACM conference on Electronic commerce*, ACM Press (2003) 192–193
9. Vilalta, R., Drissi, Y.: A perspective view and survey of meta-learning. *Artificial Intelligence Review* **18** (2002) 77–95