

Developing Adaptive Auction Mechanisms

David Pardoe and Peter Stone

The University of Texas at Austin

{dpardoe, pstone}@cs.utexas.edu

Mechanism design has traditionally been a largely analytic process, relying on assumptions such as fully rational bidders. In practice, however, these assumptions may not hold, making bidder behavior difficult to model and complicating the design process. To address this issue, we propose a different approach to mechanism design. Instead of relying on analytic methods that require specific assumptions about bidders, our approach is to create a self-adapting mechanism that adjusts auction parameters in response to past auction results. In this paper, we describe our approach and then present an example of its implementation to illustrate its efficacy.

Categories and Subject Descriptors: K.4.4 [**Computers and Society**]: Electronic Commerce

General Terms: Economics, Algorithms

Additional Key Words and Phrases: auctions, machine learning

1. INTRODUCTION

Mechanism design has traditionally been a largely analytic process. Assumptions such as full rationality are made about bidders, and the resulting properties of the mechanism are analyzed in this context [Parkes 2001]. Even in large-scale real-world auction settings such as the FCC Spectrum auctions, game theorists have convened prior to the auction to determine the best mechanism to satisfy a set of objectives. Historically, this process has been incremental, requiring several live iterations to iron out wrinkles, and the results produced have been mixed [Cramton 1997; Weber 1997].

A large component of this incremental design process involves reevaluating the assumptions made about bidders in light of auction results. These assumptions pertain to things such as:

- Bidders' motivating factors such as valuation distributions and risk aversion;
- Information that is available to the bidders; and
- Bidder rationality.

Even when the assumptions about bidders can be successfully modified to explain past results, the process requires human input and is time consuming, delaying the speed with which changes can be made to the mechanism.

Perhaps the largest challenge involves the case where the assumption of fully

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2005 ACM /2005/-0001 \$5.00

rational bidders is violated. In practice, bidders are not able to attain full rationality in complex, real-world settings. Rather, they employ heuristic strategies that are in general opaque to the seller, certainly a priori, and often even after the auction. Modeling such bidder behavior analytically is challenging at best.

To address these difficulties, we propose a substantially different approach to mechanism design - developing self-adaptive mechanisms that require minimal assumptions about bidders. We describe this approach at a high level in the next section, and then present an illustrative example of its implementation in Section 3. We discuss how this approach compares to related work in Section 4, and Section 5 concludes.

2. AN ADAPTIVE APPROACH

The strategies employed in an auction by bidders are often unknown to the seller. Nonetheless, the effectiveness of the mechanism can vary drastically as a function of the bidding strategies used. As a result, we view adaptive mechanism design as an online empirical process whereby the mechanism adapts to maximize a given objective function based on observed outcomes. Because we do not assume any insight into the bidders' strategies, their behaviors cannot be simulated. Consequently, this process is necessarily online.

Our view of adaptive mechanisms is illustrated in Figure 1. A parameterized mechanism is defined such that an *evaluator* module can revise parameters in response to observed results of previous auctions. Upon execution, the parameterized mechanism clears one or more auctions involving a population of bidders with various, potentially unknown, bidding strategies. The results of the auction are then taken as input to the evaluator as it revises the mechanism parameters in an effort to maximize an objective function such as seller revenue. 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. (For an extensive parameterization of the auction design space, see [Wurman et al. 2001].)

We view the mechanism selection module as the key active element in this picture. It is essentially an online machine learning module aiming to characterize the function from mechanism parameters to expected revenue (or any other objective function). Because the learner can select its training examples and the target output is in general a continuous value, the problem is an active learning [Saar-Tsechansky and Provost 2004] regression problem. A key characteristic is that the learning is all done online, so that excessive exploration can be costly.

The bidders in Figure 1 use a variety of different bidding strategies, including heuristic, analytic, and learning-based approaches. For the latter to make sense, the same bidders must interact repeatedly with the mechanism, leading to a potential co-evolutionary scenario in which the bidders and mechanism continue to adapt in response to each other [Phelps et al. 2002]. However, our approach does not depend on repeated interactions with the same bidders. The only assumption about the bidders is that their behavior either remains consistent or changes slowly enough that it is possible to learn to predict auction results as a function of the mechanism, at least in expectation.

In the following section we instantiate an adaptive mechanism in the most straight-

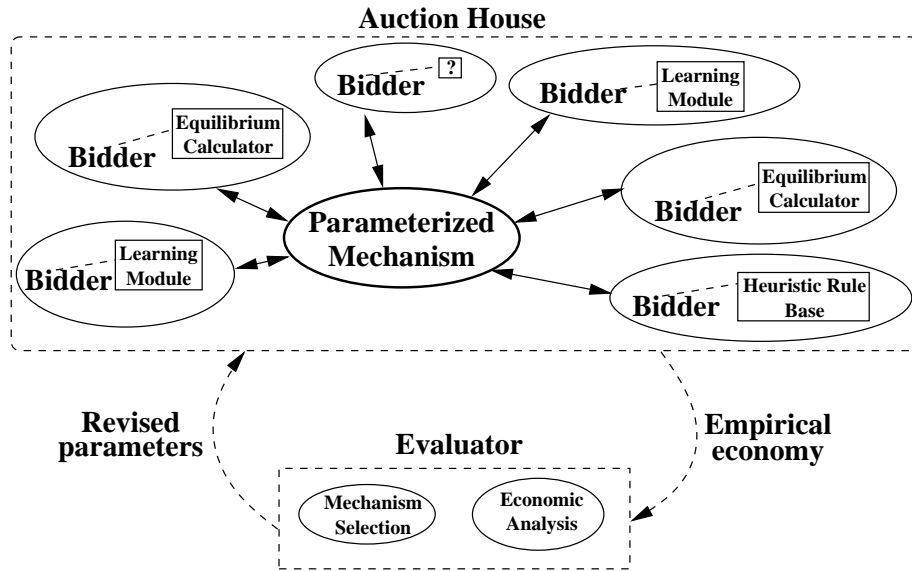


Fig. 1. A high-level illustration of the concept of adaptive mechanisms. From the point of view of the evaluator, the various bidder behaviors are unknown aspects of the environment.

forward way possible, considering a case in which there is a small discrete set of possible mechanism parameters.

3. AN EXAMPLE IMPLEMENTATION

In this section, we describe how a seller could use an adaptive mechanism to maximize its revenue in a series of sequential auctions. After giving some background on bidder behavior in sequential auctions, we describe the seller’s situation in detail and present experimental results.

3.1 Sequential auctions

Sequential auctions are a common method of selling groups of similar or identical items. In sequential auctions, items are sold one at a time, with separate bidding on each item. Milgrom and Weber [1982] were the first to derive equilibrium solutions for sequential auctions. For the case of identical items, independent private values, and unit demand, they determined that prices follow a martingale¹, remaining constant on average throughout the auction. With affiliated values, prices tend to drift upwards, due to the information revealed by winning prices.

In practice, however, prices in sequential auctions often decrease. This phenomenon is known as the “price decline anomaly”. Ashenfelter [1989] first observed

¹<http://mathworld.wolfram.com/Martingale.html>

the anomaly in wine auctions, noting that prices were twice as likely to decrease as they were to increase when pairs of identical items were considered. Additional studies since then have identified this phenomenon in other auction settings and attempted to offer explanations. For example, McAfee and Vincent [1993] consider the possibility of risk aversion, while von der Fehr [1994] explores the effects of participation costs. Jones et al. [2004] study Australian wool auctions and detect a statistically significant price change over the course of an auction approximately thirty percent of the time. Surprisingly, when there is a change, it is just as likely to be positive as negative, and no properties of the auctions appear sufficient to predict this effect.

This potential for unpredictable bidder behavior in sequential auctions motivates the following scenario.

3.2 An auction design choice

We consider a somewhat artificial but illustrative scenario in which a seller is repeatedly required to auction off large sets of items. Within each set, items are identical. For the sake of simplicity, we will assume that each set contains the same number of items, in this case 60. The seller would like to conduct uniform-price sealed-bid auctions, where each winning bidder must pay the amount of the highest losing bid (known as an $n+1$ st-price auction), and each bidder may only purchase a single item. However, the seller is unsure whether it would be more profitable to sell entire sets at once, or to take a more sequential approach and sell each set over a number of periods, dividing the set equally among the periods and announcing the winning price after each period. For example, the seller could sell all 60 items in a single auction, or sell 30 items, announce the winning bidders and the price they will pay, and then auction off the remaining items to the remaining bidders. Thus, the only free auction parameter to be experimented upon by the seller is the number of periods. The seller considers between one and four periods.

We assume that the seller has almost no knowledge concerning the bidders that will participate in the auctions. In particular, the seller knows nothing about the potential bidder types, where a bidder's type consists of its signal and the strategy it employs. This lack of knowledge rules out an analytical approach to choosing the number of periods. The only assumption made by the seller is that for each auction, the number of bidders and the bidders' types are drawn from the same distributions.

The seller's goal is to maximize its revenue over a large number of auctions. In order to accomplish this goal, the seller must identify the optimal value of a discrete parameter (the number of periods) through online experimentation producing noisy results. This situation is an instance of the k -armed bandit problem (here k equals four), a classic reinforcement learning problem. One approach to k -armed bandit problems, and the one we use here, is sample averaging with ϵ -greedy action selection [Sutton and Barto 1998].² This method records the average revenue re-

²More sophisticated approaches, such as softmax action-selection and interval estimation [Kaelbling and Nilsson 1990], also exist. Our purpose in this paper is not to compare approaches to the k -armed bandit problem, but rather to demonstrate that adaptive mechanism design is feasible in our setting.

ceived for each choice over the seller’s entire history. At each step, the choice with the highest average is taken with probability $1-\epsilon$, and a uniformly random choice is made otherwise. We now demonstrate that by using this approach, the seller can come close to obtaining the optimal revenue for any given situation, and can on average obtain a higher total reward than would be expected from sticking with one parameter choice.

3.3 Experiment

In order to test the value of adaptivity in this scenario, we need to identify the properties of the bidders. We assume that 120 bidders will participate in each auction, and that the bidders’ types will be drawn from the same distribution in each auction. A type consists of a signal, chosen uniformly randomly from the interval $[0,1]$, and one of five heuristic strategies to use, chosen randomly according to a fixed distribution. The five strategies are designed to represent a variety of plausible, but not optimal, bidding approaches that might be employed by a bidder. Recall that a key motivation for this work is the fact that bidders are not able to bid optimally in complex real-world scenarios. The five bidder strategies are as follows.

- (1) A bidder using the equilibrium bidding strategy for the case of independent private values drawn from a uniform distribution with unit demand [Milgrom and Weber 1982]. The strategy is to bid one’s value in the final period, and to bid

$$\frac{\text{total bidders} - \text{total items}}{\text{total bidders} - \text{items sold}} * \text{value}$$

in previous periods, where *items sold* includes those sold up to and including the current period. Note that if all bidders use this strategy, the auction will be revenue-equivalent regardless of the number of periods chosen.

- (2) A bidder that tries to play the same equilibrium strategy, but overestimates the number of bidders participating (which is not announced) by 40. This strategy results in higher bids in all but the last period.
- (3) A bidder that similarly underestimates the number of bidders by 40.
- (4) A bidder that is either impatient or faces some type of participation cost, and therefore has a tendency to drop out in later periods. This is modeled by simply setting a fixed probability of .15 for the bidder to drop out after each period. When bidding, the aforementioned equilibrium strategy is used.
- (5) A simple approximation of an affiliated values case. A dealer buys items to sell later at a profit. Although the dealer has little knowledge of the value of each item, the price at which he can resell items can be predicted from the winning prices. The dealer’s heuristic is to bid conservatively in the first period (one fourth his signal), and afterwards to bid according to the aforementioned equilibrium strategy using the previous period’s winning price as his value.

Note that these strategies are defined only for the purpose of permitting simulation, and are completely unknown to the seller we are modeling.

With these strategies defined, we can now measure the average revenue obtained for each choice of number of periods for any given distribution of bidder types,

and compare the results to those obtained using our adaptive method. Figures 2 and 3 show the average results for 1000 runs of 1000 auctions for two different bidder distributions P , where P_i represents the probability that a bidder will have the i th strategy from the list above. A value of 0.1 is used for ϵ . The average revenue per auction is graphed as a constant for each fixed period choice, but is shown varying over time for the adaptive method. The optimal choice of periods for each distribution is different, but in each case our adaptive method is able to approach the revenue of the optimal choice. The average revenue does not actually converge to that of the optimal choice because of the ϵ probability of exploring at each step. Decreasing ϵ over time or using a more sophisticated learning method would remedy this limitation.

The true measure of the success of our adaptive method is whether the seller would expect a higher total revenue from using it than from choosing a fixed number of periods for an unknown bidder distribution. To determine this, we found the average revenue per auction for each choice of number of periods over a large number of distributions. We randomly chose 1000 bidder distributions, and ran each period choice for 1000 auctions on each distribution. The average revenue per thousand-auction trial for each period choice is shown in Table I. We can see that our adaptive method does in fact result in a higher expected revenue than any fixed choice. The difference between the results of the adaptive method and each fixed choice is statistically significant at the 95% confidence level according to a paired t-test.

We also consider a more challenging scenario in which we relax the assumption that the distribution of bidder types is stationary. It is reasonable to assume that over a period of time, the number of bidders might increase or decrease, and bidders might refine their strategies in response to past auction results. A robust auction mechanism should be able to handle such variation. With a changing bidder distribution, the expected revenue for each period choice changes from one auction to another. As a result, we are now faced with a nonstationary k-armed bandit problem. To solve this problem, we again use the ϵ -greedy method, but with recency-weighted averaging of the results for each parameter choice, as described in Table III.

We evaluated our method as before, by running 1000 auctions for 1000 random starting bidder distributions. This time, however, we allowed the distribution to change over time by having the probability of each bidder strategy follow a random walk, as shown in Table IV. A value of 0.01 is used for *step_size*, and ϵ is set to 0.1 again. The results are shown in Table II. Again, our adaptive method is able to outperform any fixed choice, and the differences are statistically significant at the 95% confidence level according to paired t-tests.

With this illustrative example we have demonstrated that by applying a form of adaptivity to the design of an auction mechanism, a seller may be able to increase its revenue.

4. 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. In this section, we




 Fig. 2. Distribution $P = (.2, .2, .2, .2, .2)$

 Fig. 3. Distribution $P = (.2, .1, .3, .3, .1)$

Table I. Avg. revenue, fixed bidder dist.

# Periods	Revenue	Std. dev.
1	24,707	$\pm 2,531$
2	24,933	$\pm 2,447$
3	24,867	$\pm 2,345$
4	24,757	$\pm 2,331$
adaptive	25,334	$\pm 2,341$

Table II. Avg. revenue, bidders change

# Periods	Revenue	Std. dev.
1	24,281	± 966
2	24,294	± 887
3	24,368	± 881
4	24,268	± 878
adaptive	24,500	± 877

briefly survey that work and relate it to our own.

Cliff [2001] explores a continuous space of auction mechanisms defined by a parameterized version of the continuous double auction, where the parameter represents the probability that a seller will make an offer during any time slice. The mechanism parameter and the parameters of the simulated bidding agents used are evolved simultaneously using a genetic algorithm. For different underlying supply and demand schedules, the system converges to different values of the auction parameter. Phelps et al. [2002] also address continuous double auctions, using genetic programming to co-evolve buyer and seller strategies and auction rules from scratch.

Byde [2003] takes a similar approach in studying the space of auction mechanisms between the first and second-price sealed-bid auction. The winner's payment is determined as a weighted average of the two highest bids, with the weighting determined by the auction parameter. For a given population of bidders, the revenue-maximizing parameter is approximated by considering a number of parameter choices over the allowed range, using a genetic algorithm to learn the parameters of the bidders' strategies for each choice, and observing the resulting average revenues. For different bidder populations (factors considered include variable bidder counts, risk sensitivity, and correlation of signals), different auction parameter values are found to maximize revenue.

The primary difference between these previous approaches and the method advocated in this paper is that these approaches adapt mechanisms in simulation, while our aim is to develop mechanisms that are self-adapting in an online setting. Although the auction mechanisms developed by these approaches may work well under the assumed conditions, when they are used in real-life settings the same

-
- Let $step_size$ represent the weight given to new samples.
 - For each period choice i , initialize the variables sum_i and $weight_i$ optimistically to encourage initial exploration. We used the values $weight_i = 5 * step_size$ and $sum_i = 35 * weight_i$.
 - After each auction, perform the following updates for each period choice i :
 - $sum_i = sum_i * (1 - step_size)$
 - $weight_i = weight_i * (1 - step_size)$
 - Then, if the auction used period choice j and yielded revenue r , perform these updates:
 - $sum_j = sum_j + step_size * r$
 - $weight_j = weight_j + step_size$
 - To find the expected revenue for period choice i , compute $sum_i/weight_i$.

Table III. Implementation of recency-weighted averaging

-
- Let P be the distribution over bidder strategies, and let R control the rate at which P changes.
 - Initialize each P_i randomly in $[0,1]$, and normalize P to obtain a valid distribution.
 - Set R to a random integer between 0 and 500.
 - After each auction, repeat R times:
 - Choose i and j as random integers between 1 and 5
 - If $P_i > .001$, set $P_i = P_i - .001$ and $P_j = P_j + .001$

Table IV. Implementation of the random walk

problem may arise as with analytical mechanism design: bidders' goals, beliefs, and strategies may be different from those assumed, leading to unexpected results. Although the adaptive measures used in these approaches could be applied in an online setting, they would likely be found unsuitable. For example, evolutionary methods frequently explore highly suboptimal solutions that could be disastrous if actually tried. Our goal is to design adaptive mechanisms that are both safe to use and capable of quickly finding the parameters best suited to the participating bidders, all while making as few assumptions as necessary about the behavior of these bidders.

The lack of assumptions about bidders also differentiates this work from Conitzer and Sandholm's *automated mechanism design* [2004], in which mechanisms are computationally developed for specific situations based on complete knowledge of the distribution over bidders' types.

The particular case of a seller responding to customer behavior by using a k-armed bandit algorithm has been well studied (e.g. [Rothschild 1974]), but generally in the context of price selection. The aim of this paper is not to analyze a particular

learning algorithm, but to propose the use of such algorithms to adaptively tune parameters for any chosen auction parameterization.

5. CONCLUSIONS AND FUTURE WORK

In this paper, we have proposed 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 demonstrated the efficacy of this approach in a situation where a seller must choose from a space of sequential auctions in order to maximize its revenue. This illustrative example represents a first step toward applying adaptive mechanisms to more challenging problems. The fact that we were able to have success using a straightforward adaptive method with no parameter-tuning is encouraging; it suggests that the success or failure of our method is not dependent on such choices, while offering hope of even better results from the use of more sophisticated methods.

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 continuous parameters and multidimensional parameterizations. In those cases, more sophisticated online learning methods will be required. 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.

REFERENCES

- ASHENFELTER, O. 1989. How auctions work for wine and art. *Journal of Economic Perspectives* 3, 3 (Summer), 23–36.
- BYDE, A. 2003. Applying evolutionary game theory to auction mechanism design. In *Proceedings of the 4th ACM conference on Electronic commerce*. ACM Press, 192–193.
- CLIFF, D. 2001. Evolution of market mechanism through a continuous space of auction types. Tech. Rep. HPL-2001-326, HP Labs.
- CONITZER, V. AND SANDHOLM, T. 2004. Self-interested automated mechanism design and implications for optimal combinatorial auctions. In *Proceedings of the 5th ACM Conference on Electronic Commerce*. ACM Press, 132–141.
- CRAMTON, P. C. 1997. The FCC spectrum auctions: An early assessment. *Journal of Economics and Management Strategy* 6, 3, 431–495.
- JONES, C., MENEZES, F., AND VELLA, F. 2004. Auction price anomalies: Evidence from wool auctions in australia. *The Economic Record* 80, 127 (09), 271–288.
- KAELBLING, L. P. AND NILSSON, N. J. 1990. Learning in embedded systems. Ph.D. thesis.
- MCAFEE, R. P. AND VINCENT, D. 1993. The declining price anomaly. *Journal of Economic Theory* 60, 1 (6), 191–212.
- MILGROM, P. AND WEBER, R. 1982. A theory of auctions and competitive bidding, ii, unpublished manuscript.
- PARKES, D. C. 2001. Iterative combinatorial auctions: Achieving economic and computational efficiency. Ph.D. thesis, Department of Computer and Information Science, University of Pennsylvania.

- PHELPS, S., MC BURNLEY, P., PARSONS, S., AND SKLAR, E. 2002. Co-evolutionary auction mechanism design. In *Agent Mediated Electronic Commerce IV: Designing Mechanisms and Systems*. Lecture Notes in Artificial Intelligence, vol. 2531. Springer Verlag.
- ROTHSCHILD, M. 1974. A two armed bandit theory of market pricing. *Journal of Economic Theory* 9, 185–202.
- SAAR-TSECHANSKY, M. AND PROVOST, F. 2004. Active learning for class probability estimation and ranking. *Machine Learning*.
- SUTTON, R. S. AND BARTO, A. G. 1998. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA.
- VON DER FEHR, N.-H. M. 1994. Predatory bidding in sequential auctions. *Oxford Economic Papers* 46, 3 (July), 345–56.
- WEBER, R. J. 1997. Making more from less: Strategic demand reduction in the FCC spectrum auctions. *Journal of Economics and Management Strategy* 6, 3, 529–548.
- WURMAN, P. R., WELLMAN, M. P., AND WALSH, W. E. 2001. A parameterization of the auction design space. *Journal of Games of Economic Behavior* 35, 304–338.