Adaptive Auctions: Learning to Adjust to Bidders

David Pardoe[1], Peter Stone[1], Maytal Saar-Tsechansky[2], and Kerem Tomak[2]

The University of Texas at Austin, Austin TX 78712, USA
[1]Department of Computer Sciences, [2]McCombs School of Business
{dpardoe, pstone}@cs.utexas.edu, {Maytal.Saar-Tsechansky, Kerem.Tomak}@mccombs.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 parameters of the auction mechanism can lead to widely differing results. Common parameters include auctioneer fees, minimum bid increments, rules governing when the auction closes, and whether sellers can set a reserve price. 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]. 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 have been mixed [2, 3]. An important component of this incremental design process involves 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. For example, assumptions are often made about

- 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, undermining the efficiency with which changes can be made to the mechanism. Perhaps the biggest challenge results from the fact that, in practice, bidders are not able to attain full rationality in complex, real-world settings [4]. Rather, they employ heuristic strategies that are in general opaque to the seller, certainly a priori, and often even after the auction.

To address these challenges, we propose a substantially different approach to mechanism design: self-adaptive mechanisms that change in response to observed bidder behavior. In previous work [5], we introduced this approach in a context with discrete auction mechanisms, thus, from a learning perspective, reducing to a k-armed bandit. In this paper, we extend our approach to an auction with a continuous parameter, thus enabling a parameter optimization approach. In this context, we propose a *metalearning* process by which the method of parameter optimization is itself parameterized and optimized based on 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 with varying degrees of loss-aversion. We describe our approach to designing adaptive mechanisms at a high level in the next section. Section 3 describes an auction scenario involving loss averse bidders, and we present an illustrative application of our adaptive approach to this scenario in Section 4. We discuss how our approach compares to related work in Section 5, and Section 6 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 allow for the possibility of unexpected bidder behavior, this process must be performed online during interactions with real bidders.

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, generally 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 [6].)

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 because the target output is, in general, continuous, the problem is an active learning [7]

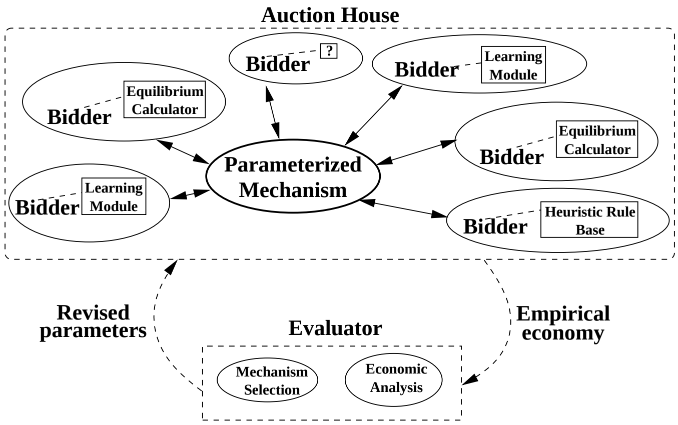

Fig. 1. A high-level illustration of the concept of adaptive mechanisms. From the point of view of the evaluator, the bidder behaviors are unknown aspects of the environment.

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 may 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 [8]. However, our approach does not depend on repeated interactions with the same bidders. The only required assumption about the 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.

In the following sections, we present an illustrative application of this approach to a particular auction scenario.

3 An Auction Scenario

We now describe the auction scenario that we will use to demonstrate our approach to choosing an adaptive mechanism. First we introduce the concept of loss averse bidders, then we describe the auction scenario and provide a means of simulating bidder behavior under this scenario.

3.1 Loss averse bidders

We consider an English (ascending, open-cry) auction in which the bidders have independent, private (i.e., unknown to other bidders) values for the goods being sold. Bidders submit ascending bids until no incremental bids are made above the winning bid. We assume that the seller may set a *reserve price* indicating the minimum acceptable bid. In the absence of any bid higher than the reserve price, no transaction occurs. It has been shown that when bidders are rational, the optimal reserve price should be higher than the seller's valuation of the item [9]; however, a reserve price of 0 is often seen in practice.

Dodonova explains this phenomenon by bidders' *loss aversion* [10]. Loss aversion violates the rationality assumption because the utility from a gain is lower than the disutility from a loss 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 $\alpha > 1$. A bidder considers that it is "losing" an item if it was the high bidder at some point in the auction, but then does not win the item.

We assume that the bidders are, to varying degrees, loss averse. Note that if $\alpha = 1$ we arrive at the traditional loss neutral bidders as a degenerate case. Under these assumptions and model setup, Dodonova derives the equilibrium as follows. Assuming two loss averse bidders, a first mover submits a bid in the beginning of the auction if his valuation is higher than the reserve price. The second bidder responds by submitting an increment above the current winning bid only if by doing so the bidder can guarantee a positive expected utility. In particular, this will be the case only if

$$\int_r^{v_2} (v_2 - \alpha v_1) f(v_1) dv_1 > 0$$

where r is the reserve price, v_2 is the second bidder's valuation, and f is the (known) probability distribution function over valuations. Intuitively, the second bidder recognizes that loss aversion may lead him to pay more than his valuation if the first bidder's valuation is sufficiently high, making him more reluctant to enter the auction. With only one active bidder, the auctioneer's revenue is decreased.

If all bidders participate, the auction continues as a standard ascending price English auction until a bidder's marginal utility from losing the object is less than the (potential) winning bid, i.e, the losing bidder will bid up to α times his valuation and then drop out. This equilibrium can cause the seller's optimal reserve price to be 0 under certain conditions. For instance, if f is a uniform distribution, a reserve of 0 will maximize the seller's revenue for values of α above 1.3. The equilibrium 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 2. 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.

3.2 Scenario description

We consider a scenario in which a seller interacts repeatedly with bidders drawn from a fixed population. In particular, the seller has n identical items that will be sold one at a time through a series of English auctions. For the sake of simplicity, we assume that two bidders participate in each auction. The seller sets a reserve price for each auction, thus restricting the possible bids available to the bidders and 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 by solving for the reserve price maximizing expected revenue under this model. 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. The seller is also able to estimate the ranges of values for bidders' valuations and degrees of loss aversion ($[v_{min}, v_{max}]$ and $[\alpha_{min}, \alpha_{max}]$, respectively), but does not know the actual distributions from which these values are drawn, or the strategies bidders will employ.

We assume that a given bidder assigns the same value to any one of the items sold. In addition, 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.

3.3 Bidder simulation

As described in Section 2, although the seller does not have a complete model of the bidder behavior, it is still possible to take advantage of the partial knowledge that is available to guide the selection of the adaptive mechanism. In order to do so, we need a method of generating plausible bidder behavior so that we can evaluate the adaptive mechanism in simulation.

To represent the information available to the seller, we choose the following values: $n = 1000$, $v_{min} = 0$, $v_{max} = 1$, $\alpha_{min} = 1$, and $\alpha_{max} = 2.5$. The distributions from which v and α are drawn, and the strategies that take these values as inputs, are unknown. In order to simulate a set of n auctions, which we will refer to as an *episode*, these unknowns must be specified, which we do as follows. For each episode we wish to simulate, we first randomly generate an "arbitrary" distribution for valuations by taking a Gaussian with a mean chosen uniformly from [0, 1] and a variance of 10^x with x chosen uniformly from [-2, 1], and then normalize the function so that the portion over the range [0, 1] represents a PDF. We then generate a distribution for α in the same way, only with a range of [1, 2.5] for both the mean and the function.

We simulate bidder behavior by having bidders follow the equilibrium strategy given in Section 3.1 under the assumption that the other bidder has the same α (because this is the situation to which the equilibrium solution applies). Thus for each auction in an episode, we draw two values from the valuation distribution, draw a single α from the α distribution, randomly assign one bidder to be the initial bidder, and then have both bidders bid as specified in the equilibrium strategy.

This approach to simulating bidder behavior could be viewed as specifying a probability distribution over bidder *populations*, and drawing a population from this distribution for each episode to be simulated. Essentially, we are addressing the seller's uncertainty about bidder behavior by training the mechanism to adapt to a variety of bidder populations. It is important to note that the distribution over populations need not be accurate or known to the seller. Furthermore, this distribution need not be expressed explicitly as a function – it may be any algorithm that can generate bidder behavior, such as a learning algorithm. All that matters is that the seller be able to to generate experience with a variety of different representative bidder populations.

To illustrate the task faced by the seller, we generated 10,000 bidder populations as described, and found the average revenue over all populations for each reserve price between 0 and 1 at intervals of 0.01. The average revenue for each choice of reserve is shown by the solid line in Figure 2. 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 2 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, while our goal in learning the parameters of the mechanism is to take advantage of the first observation by using the knowledge it represents to focus the mechanism's exploration.

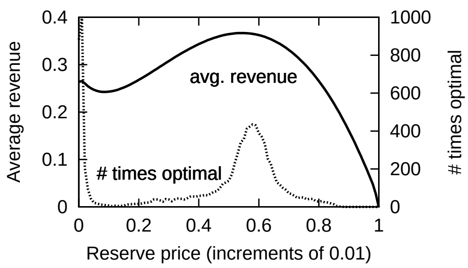

Fig. 2.

4 Implementation and Results

As specified at the end of Section 2, for the auction scenario with the specific goal of maximizing revenue over n auctions, we have now 1) chosen the auction parameterization (the reserve price represents a single, continuous parameter), and 2) described 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

For clarity, we begin by describing a somewhat simplified version of the adaptive method we will implement. In this approach, we discretize the problem by restricting the seller to choosing one of k choices for the reserve price at each step, where the ith choice is a price of $(i-1)/(k-1)$. The resulting problem can be viewed as an instance of the k-armed bandit problem, a classic reinforcement learning problem [11]. In k-armed bandit 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 choice'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 weight of each choice. 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* [11] 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}).

The approach just described, which we will call the *bandit approach*, has one significant limitation: the assumption that the expected revenue of each choice is independent. Because the choices we are considering represent points chosen along a continuous range of reserve prices, it is likely that the expected revenues of nearby choices will be similar, and thus experience could be profitably shared between choices. To address this issue, we now introduce an enhanced approach we will call the *regression approach*. As the name suggests, we perform regression over past auction results to derive a function mapping the reserve price to the expected revenue. In particular, we perform locally weighted quadratic regression (LWQR) [12], a form of instance-based regression. To predict the expected revenue for a given reserve price, the weight of each existing data point is determined by taking its distance from the given price and applying a Gaussian weighting function. Parameters are then found specifying the quadratic that minimizes the weighted sum of squared errors. This process is repeated for each price for which we want an estimate of expected revenue.

Because we can now predict the expected revenue of any reserve price, even if we have no experience at that price, we are no longer restricted to considering a finite number of choices as in the bandit approach. We continue to discretize the range of prices for computational reasons — doing so allows us to implement an incremental version of LWQR and also to use softmax action selection without modification. However, we are able to effectively use much finer discretizations than before. In fact, we observed no benefit from increasing beyond 100 choices, so we treat the degree of discretization as a fixed parameter for the regression approach, and reinterpret k as described below.

The parameters for the regression approach are almost the same as those of the bandit approach. We allow the temperature to vary as before, and the concept of initial experience remains similar. We still use k pairs of parameters avg_i and $count_i$, with each pair now representing a data point for reserve price $(i-1)/(k-1)$ and revenue avg_i that will be used during regression as if it represented $count_i$ such data points. It should be noted that in the regression approach, k is used only to specify the number of points used as initial experience, and is independent of whatever degree of discretization is used for selection of reserve prices. The only additional parameter is the kernel width used in the weighting function. We use a single kernel width, and ignore for now the possibility of having the kernel width vary as a function of the reserve price.

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 as described in Section 3.3 and running an episode using those parameters. This estimate 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) [13], a popular method of stochastic optimization based on gradient approximation. At each step, two estimates of the expected episode revenue are taken for slight perturbations of the current parameters (the same bidder population is used for each estimate), a gradient approximation is found, and the parameters are updated in the direction of the gradient.

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, and a kernel width of 0.1 is used. The search results appear stable in that repeated runs result in parameters that are fairly similar and provide nearly identical expected revenue per episode. Modest changes to the initial parameters do not affect the quality of the outcome.

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 appear to be the best values after running searches with several values of k.

4.3 Results

To evaluate our adaptive methods, we first searched for the best possible set of parameters, including k, as described above, for both the bandit and regression approaches. For the bandit approach, a value of 13 was optimal for k, while increasing k beyond 11 gave no apparent benefit in the regression case. The learned parameters are presented in Figures 3 and 4. Initial experience is displayed visually by plotting a circle for each avg_i with area proportional to $count_i$. Both sets of initial experience appear reasonable given Figure 2. For the bandit approach, 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 in the beginning of an episode. For the regression approach, the values of *count* are similar in most cases, but the values of *avg* are higher in the more promising regions, again encouraging initial exploration. The reasons for such small *count* values at 0.7 and 0.8 are not immediately clear.

We next generated a set of 10,000 bidder populations, and found the average revenue per episode for both approaches using both the initial and the learned parameters. The average revenues per auction are shown in Table 1, while a plot of the average revenue for each auction over an entire episode is shown in Figure 5. 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 5 we can see that while all methods approach the same revenue by the last auction in an episode, using learned parameters leads to much higher revenues during the early part of an episode, especially with the regression approach. For instance, the average revenue reached on the 100th auction by the regression approach with learned parameters is not reached until after at least 500 auctions with other approaches. 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. In this section, we briefly survey that work and relate it to our own.

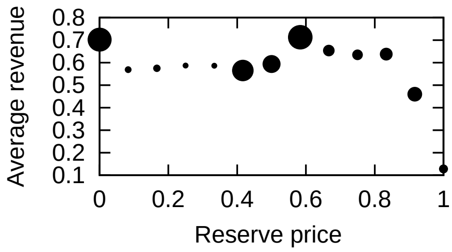

Fig. 3. Learned parameters – bandit $\tau_{start} = .0423$, $\tau_{end} = .0077$

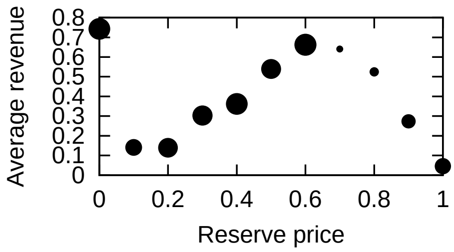

Fig. 4. Learned parameters – regression $\tau_{start} = .0081$, $\tau_{end} = .0013$, k. width = .138

Adaptive method	Total Revenue
best fixed reserve price (0.54)	0.367
bandit, initial parameters	0.374
bandit, learned parameters	0.394
regression, initial parameters	0.385
regression, learned parameters	0.405

Table 1. Average revenue per auction for each adaptive method. Differences are statistically significant at the 99% confidence level according to paired t-tests.

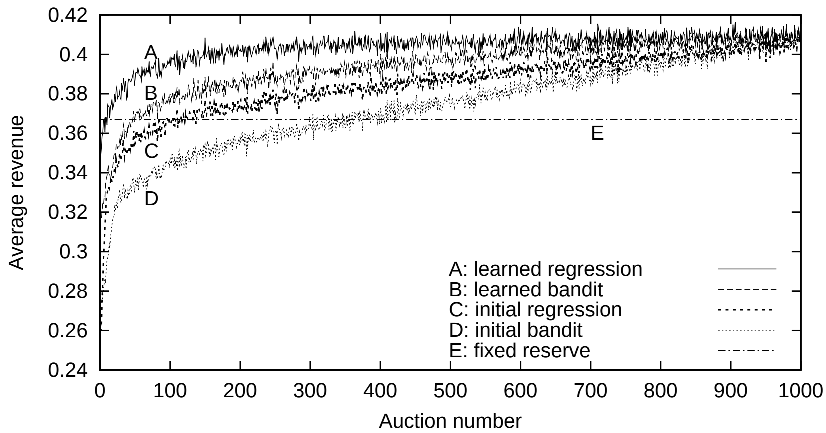

Fig. 5. Average revenue per auction over the course of an episode for each method.

Cliff [14] explores a continuous space of auction mechanisms defined by a parameterized 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. [8] also address continuous double auctions, using genetic programming to co-evolve buyer and seller strategies and auction rules from scratch.

Byde [15] 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 use simulation to produce fixed mechanisms, while our aim is to develop mechanisms that are self-adapting in an online setting. (The methods used to learn bidder strategies, however, could possibly be applied in our approach to generate the bidder behavior needed during the search for optimal adaptive parameters.) 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 problem may arise as with analytical mechanism design: bidders' goals, beliefs, and strategies may be different from those assumed, leading to unexpected results. While 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.

Dittrich et al. [16] present a different take on adaptation involving loss averse bidders, analyzing the effect that loss aversion has on the learning dynamics exhibited by bidders adapting in response to experience.

The process of identifying the parameters of the adaptive mechanism can be viewed as an instance of *metalearning* [17]. 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. Cramton, P.C.: The FCC spectrum auctions: An early assessment. Journal of Economics and Management Strategy **6** (1997) 431–495
3. Weber, R.J.: Making more from less: Strategic demand reduction in the FCC spectrum auctions. Journal of Economics and Management Strategy **6** (1997) 529–548
4. Kahneman, D., Tversky, A.: Prospect theory: An analysis of decision under risk. Econometrica **47** (1979) 263–291
5. Pardoe, D., Stone, P.: Developing adaptive auction mechanisms. SIGecom Exchanges **5** (2005) 1–10
6. Wurman, P.R., Wellman, M.P., Walsh, W.E.: A parameterization of the auction design space. Journal of Games of Economic Behavior **35** (2001) 304–338
7. Saar-Tsechansky, M., Provost, F.: Active learning for class probability estimation and ranking. Machine Learning (2004)
8. 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)
9. Myerson, R.B.: Optimal auction design. Mathematics of Operations Research **6** (1981) 58–73
10. Dodonova, A., Khoroshilov, Y.: Optimal auction design when bidders are loss averse. Working Paper. University of Ottawa. (2004)
11. Sutton, R.S., Barto, A.G.: Reinforcement Learning: An Introduction. MIT Press, Cambridge, MA (1998)
12. Atkeson, C.G., Moore, A.W., Schaal, S.: Locally weighted learning. Artificial Intelligence Review **11** (1997) 11–73
13. Spall, J.C.: An overview of the simultaneous perturbation method for efficient optimization. Johns Hopkins APL Technical Digest **19** (1998) 482–492
14. Cliff, D.: Evolution of market mechanism through a continuous space of auction types. Technical Report HPL-2001-326, HP Labs (2001)
15. 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
16. Dittrich, D.A.V., Guth, W., Kocher, M., Pezanis-Christou, P.: Loss aversion and learning to bid. Technical Report 2005-03, Max Planck Institute for Research into Economic Systems, Strategic Interaction Group (2005)
17. Vilalta, R., Drissi, Y.: A perspective view and survey of meta-learning. Artificial Intelligence Review **18** (2002) 77–95