In *Proceedings of the International Conference on Electronic Commerce* (ICEC 06), Fredericton, Canada, August 2006.

Adaptive Mechanism Design: A Metalearning Approach

David Pardoe Peter Stone
Department of Computer Sciences
The University of Texas at Austin
{dpardoe, pstone}@cs.utexas.edu

Maytal Saar-Tsechansky Kerem Tomak
McCombs School of Business
The University of Texas at Austin
{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 often exhibit unknown and variable behavior, making them difficult to model and complicating the design process. To address this challenge, we explore the use of 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. In addition, we show how predictions of possible bidder behavior can be incorporated into the adaptive mechanism through a *metalearning* process. The approach is fully implemented and tested. Results indicate that the adaptive mechanism is able to outperform any single fixed mechanism, and that the addition of metalearning improves performance substantially.

Categories and Subject Descriptors

I.2.6 [**Computing Methods**]: Artificial Intelligence—*Learning*

General Terms

Algorithms, Economics

Keywords

auctions, mechanism design, machine learning, metalearning

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. Many different types of auction *mechanisms* defining the rules of exchange may be used for such purposes. Varying parameters of the auction mechanism, such as auctioneer fees, minimum bid increments, and reserve prices, can lead to widely differing results depending on factors such as bidder strategies and product types. This paper considers *learning* auction parameters so as to maximize an objective function, such as 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 [12]. 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 [6, 20]. 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, it is known that the expected revenue of first-price and second-price sealed-bid auctions are the same, assuming risk-neutral, fully rational bidders. In general, 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 revised based on their past behavior, the process requires human input and is time consuming, undermining the efficiency with which changes can be made to the mechanism. In the case of the FCC spectrum auctions, months or years elapse between each iteration, leaving time for the experts to reconvene and update the mechanism. But 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 eBay or Google keyword auctions, this inefficiency is a serious drawback. Perhaps the biggest challenge results from the fact that, in practice, bidders are not able to attain full rationality in complex, real-world settings [9]. Rather, they employ heuristic strategies that are in general opaque to the seller, certainly a priori, and often even after the auction.

One method of addressing these challenges that has received recent attention is the use of machine learning algorithms to revise auction parameters in response to observed bidder behavior. For instance, [2] and [3] consider the problem of maximizing seller revenue in online auctions through the use of online learning algorithms for combining "expert" advice. Such approaches differ significantly from the traditional approach to mechanism design in that few or no assumptions are made about bidders, and auction mechanisms are evaluated based on worst-case performance.

In this paper, we present somewhat of an intermediate approach. As with the methods described in the previous paragraph, we consider an *adaptive mechanism* that changes in response to observed bidder behavior through the use of a learning algorithm. However, we assume that reasonable predictions about a *range of possible bidder behaviors* can be made, and we choose the learning algorithm in such a way that *expected* performance is optimized with respect to these predictions.

In particular, we present an adaptive mechanism in which a continuous parameter defining the mechanism is dynamically adjusted over time, 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*, a particular form of bidder irrationality that has been observed in empirical studies. 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. Experimental results are presented in Section 5 for this scenario, along with results for modified scenarios designed to illustrate the robustness of our method. We discuss how our approach compares to related work in Section 6, and Section 7 concludes.

2. ADAPTIVE MECHANISM DESIGN

To motivate the problem under consideration, we begin with an illustrative example that we will refer to throughout this section. Consider the challenge faced by a seller auctioning off items through an auction service such as eBay. For each auction, the seller must set various parameters defining the particular auction mechanism to be used. (In the case of eBay auctions, these parameters include the start time and duration, start price and reserve price, and possibly a buy-it-now price.) The behavior of bidders depends heavily on the type of item being sold; for instance, buyers of rare coins and buyers of video games are likely distinct populations with very different approaches to bidding. As a result, the mechanism that works best for one type of item may not be appropriate for a different item.

The goal of the seller is to identify the auction parameters that will result in the highest revenue for each item sold. When extensive historical data on past auctions of identical items is available (as is the case with eBay), it may be possible for the seller to estimate the optimal parameters by analyzing this data (e.g. [16]). This approach is not always possible, however. If the seller is introducing a new item to the market, no such data will be available. Alternatively, if there is a sudden change in demand for an item, past data may not accurately reflect the behavior of current bidders. In such cases, the seller must guess which auction parameters will work best. If the seller has multiple identical items to sell, however, then it may be possible for the seller to learn from its own experience about the effectiveness of various parameter settings.

This situation illustrates the general problem considered in this paper. While the effectiveness of an auction mechanism can vary drastically as a function of the behavior of bidders, this behavior is often difficult to predict when choosing the mechanism. In settings in which a large number of similar auctions are held, it may be reasonable to assume that although bidders' behaviors may change, they remain somewhat consistent for a period of time, 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 suggest the use of what we call *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.)

Our view of adaptive mechanism design is illustrated in Figure 1. A parameterized mechanism is defined such that the seller can use an *adaptive algorithm* to revise parameters in response to observed results of previous auctions, choosing the most promising parameters to be used in future 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 adaptive algorithm 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 [21].)

The adaptive algorithm is essentially an online machine learning algorithm aiming to characterize the function from mechanism parameters to expected revenue (or any other objective function). Because the seller can select its own training examples (by choosing which set of auction parameters to try next), and because the target output is, in general, continuous, the problem is an active learning [15] 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 [13]. 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 somewhat consistent (e.g. bidders associated with a particular industry tend to bid similarly) for a sufficient period of time to allow for prediction of auction results as a function

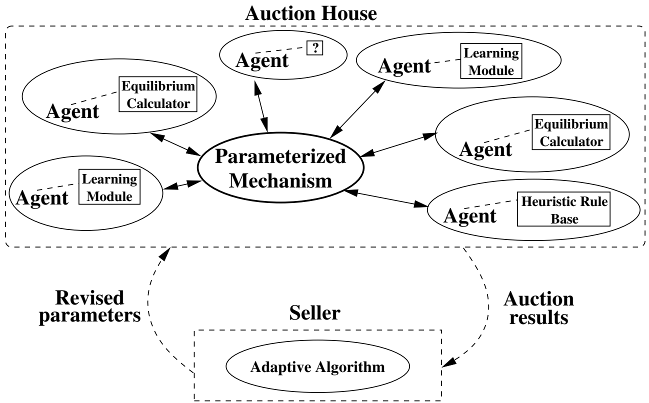

Figure 1: A high-level illustration of the concept of adaptive mechanisms. From the point of view of the seller, the bidder behaviors are unknown aspects of the environment.

of the mechanism, at least in expectation.

The use of an adaptive mechanism provides the possibility of identifying effective auction parameters even without explicitly modeling the bidders. However, if predictions can be made about the types of behavior to be expected, it makes sense to choose the adaptive algorithm in a way such that the adaptive mechanism performs well when bidder behavior conforms to these predictions.

To illustrate, suppose that the seller from our previous example has a large number of copies of a newly published book to be sold through a series of auctions[1]. As the books are new, there are no previous auction results available to guide the choice of auction parameters. The seller is not completely in the dark, however, because auction results for books by similar authors or on similar topics are available, and these books likely attract bidder populations with characteristics similar to those of potential bidders on the new book. If the seller wanted to choose a single set of auction parameters, a reasonable choice might be the parameters that were best for the most similar book, or an average of the best parameters for several similar books. If instead the seller wishes to use an adaptive mechanism and needs to choose the adaptive algorithm for it, it makes similar sense for the seller to choose an adaptive algorithm that *would have worked well* for these similar books. To determine how well an adaptive algorithm would have worked for a particular book, the seller could attempt to simulate the behavior of the population of bidders for that book using the past auction data, and apply the adaptive algorithm to these simulated bidders.

Alternatively, suppose that the seller receives marketing information suggesting valuations that buyers might have for the book, along with information suggesting bidding strategies that bidders tend to use in such auctions. Again, if the seller wishes to use an adaptive mechanism, it makes sense to choose an adaptive algorithm that would work well if this information is correct. The seller could again evaluate an adaptive algorithm by simulating a number of bidder populations that are plausible given the information available and applying the algorithm to each population.

In both cases, we have a seller that wishes to choose an adaptive algorithm in order to implement an adaptive auction mechanism, and the seller is able to make predictions concerning possible bidder behavior that allow it to simulate a number of possible bidder populations. The goal of the seller is to choose an adaptive algorithm that performs well with the actual bidder population it encounters, but as this population is unknown in advance, the seller attempts to identify an adaptive algorithm that performs well under simulation but maintains the ability to adapt in the case of unexpected bidder behavior. (For clarity, we will use the term *encountered* population in the following text whenever referring to the actual encountered population, and not a simulated population.)

In this paper, we explore a *metalearning* [19] approach to choosing the adaptive algorithm. Specifically, we choose an adaptive algorithm that is itself parameterized, and then search for the parameters that result in the best performance under expected bidder behavior as implemented in simulation. 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 algorithm, and the family of related tasks is the set of different bidder populations generated during simulation.

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 adaptive algorithm and its parameters.

4. Search the space of parameters of the adaptive algorithm 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

In this section we provide a concrete auction scenario in which adaptive mechanism design can be put to good use. In particular, we consider bidder populations with a specific form of irrationality that has been observed empirically: loss aversion. 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

[1]This scenario suggests the possibility of selling multiple items simultaneously, a topic explored in previous work on adaptive mechanisms [11]. In this paper, however, we restrict our attention to sequential, single item auctions, which may be most appropriate or even required in some settings.

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 [10]; however, a reserve price of 0 is often seen in practice.

Dodonova and Khoroshilov explain this phenomenon by bidders' *loss aversion* [8]. 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. Thus, in practice, a loss-averse bidder bids more aggressively after having had the highest bid at any point during the auction.

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 and Khoroshilov derive 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 shown 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* (characterized by distributions over valuations and α).

The seller has n identical items that are sold one at a time to bidders from the population through a series of n English auctions. To allow the use of the equilibrium strategy as presented in the previous section, we assume that exactly two bidders participate in each auction, and that these bidders participate in no other auctions.[2] 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, we assume that the seller does not have such a complete model, for instance because it is a new item, or even just a new set of bidders. Thus, 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 $[\alpha_{min}, \alpha_{max}]$), 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 (unless otherwise stated). Thus, the behavior exhibited by bidders are 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 behavior of the encountered bidder population, it is still possible to take advantage of the partial knowledge that is available to guide the selection of the adaptive algorithm. In order to do so, the seller needs a method of generating plausible bidder behavior to allow evaluation of the adaptive mechanism in simulation. We now describe the method of generating bidder behavior that will be used by the seller in the experiments of this paper. First, however, we emphasize that while the seller may make a number of simplifying assumptions about bidders, such assumptions have no direct impact on the implementation of the learning approach that will be described in the next section. From the standpoint of the learning approach, the seller's simulation of bidders simply represents a black box used to produce auction results, and the bidders simulated may be arbitrarily complex.

To define the scenario, we choose the following values and make them available to the seller: $n = 1000$, $v_{min} = 0$, $v_{max} = 1$, $\alpha_{min} = 1$, 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 the seller does as

[2] In principle, our learning approach should extend naturally to auctions with more than two bidders and bidders that participate in multiple auctions. However, doing so would require a much more complicated bidder model. Since it is the learning approach that is the focus of the paper, and not the bidding strategies, we restrict the bidding scenario for the sake of simplicity. Evaluating our methods in more complex scenarios is an area of ongoing work.

follows. For each episode the seller wishes to simulate, it first randomly generates 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 normalizes the distribution so that the portion over the range [0, 1] represents a PDF. The seller then generates a distribution for α in the same way, choosing variance as before and using a range of [1, 2.5] for both the mean and the entire distribution.

The seller simulates 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, the seller draws two values from the valuation distribution, draws a single α from the α distribution, randomly assigns one bidder to be the initial bidder, and then has 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, the seller is addressing its uncertainty about bidder behavior by training the mechanism to adapt to a variety of bidder populations. It is important to note that the seller's belief in the distribution over populations need not be accurate. 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.

Figure 2 helps illustrate the task faced by the seller. The left side shows the average revenue as a function of reserve price for three different populations. (The mean of v, variance of v, mean of α, and variance of α are 0.8, 0.01, 1, and 0.1 respectively for population A; 0.3, 1, 1.75, and 1 for population B; and 0.4, 0.1, 1.5, and 0.01 for population C.) In each case, the function has a different shape and a different maximum. On the right side of Figure 2, average results are shown for 10,000 bidder populations, generated as described above. The solid line represents the average revenue for each choice of reserve. 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 choose this price. However, we know that for each individual bidder population there is a distinct choice of reserve that yields the highest average revenue. In particular, the dotted line shows the number of times that each reserve (tested at intervals of 0.01 between 0 and 1) 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 adaptive algorithm is to take advantage of the first observation. Specifically, we aim to use the knowledge represented by the first observation to focus the mechanism's exploration, biasing it towards reserve prices that are more likely to be best.

4. ADAPTIVE MECHANISM IMPLEMENTATION

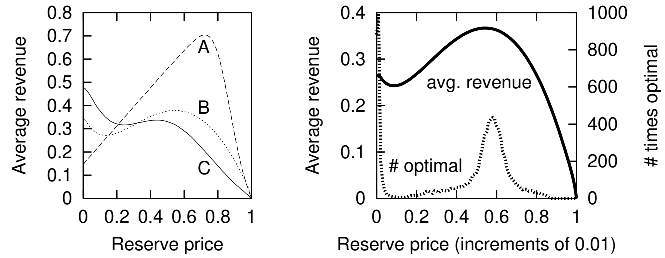

Figure 2: The graph on the left shows the average revenue for each reserve price for three different populations. The graph on the right shows average results for 10,000 populations. The solid line represents the average revenue for each reserve price, while the dotted line represents the number of times each price was optimal.

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.

4.1 Adaptive algorithm

The goal of using an adaptive mechanism is to identify the reserve price that maximizes revenue for the encountered bidder population. Recall that the adaptive algorithm attempts to estimate the expected revenue as a function of reserve price over the range of possible reserve prices, similar to one of the functions depicted on the left side of Figure 2. In this section, we describe how the adaptive algorithm obtains this estimate and how the estimate is used to select the reserve price for each auction in an episode.

For clarity, we begin by describing a somewhat simplified version of the adaptive algorithm 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 [18]. 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 of reserve price, 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 maintain a record of the average revenue for each reserve price, for each price we must track both the average revenue so far, avg_i, and the number of times that price 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* [18] 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 of reserve price 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 *prior experience* (so named because it serves a role similar to a Bayesian prior). The choice of prior 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 (one for each initial avg_i and $count_i$, plus τ_{start} and τ_{end}). These parameters can be observed in Figure 3.

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) [1], 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 prior 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 prior 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 parameterized adaptive algorithm 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 reference, all parameters are described in Table 1.) For any given set of parameter values, 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) [17], a 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 for values of k between 5 and 20.

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

5. RESULTS

In this section we present the results of experiments with our adaptive auction mechanism. We first show the results of the parameter searches and discuss the performance of the resulting adaptive mechanisms. Next, we give a comparison between our approach and a Bayesian approach to adapting

Parameter	Description
k	number of discrete choices of reserve
$count_i$	weight of prior experience at choice i
avg_i	average initial revenue at choice i
τ_{start}	initial temperature
τ_{end}	final temperature
$kernel\ width$	kernel width for regression

Table 1: Parameters of the adaptive algorithm.

to the same auction scenario. We then explore how results are impacted if the encountered bidder population does not behave exactly as expected. Finally, we consider the situation in which the bidder population is allowed to change over time.

5.1 Comparison of approaches

To evaluate our adaptive algorithms, 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. Prior experience is displayed visually by plotting a circle for each avg_i with area proportional to $count_i$. Both sets of prior 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 found the average revenue per episode for both approaches using both the initial (hand chosen) and the learned parameters. For this experiment and all remaining experiments in this paper, each approach was evaluated on 10,000 bidder populations for one episode per population. For each approach, the random number generator used in generating all bidder parameters received the same seed, ensuring identical bidder behavior. In this experiment, we drew the populations for testing (the encountered populations) from the same distribution used by the seller to generate simulated populations, meaning that these results describe the situation in which the seller has correctly predicted the types of populations that might be encountered and their probabilities of occurring. The average revenues per auction are shown in Table 2, 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. In fact, all but the bandit approach with initial parameters require far fewer than 1000 auctions to be beneficial (i.e. obtain a higher revenue sum than the fixed reserve). 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

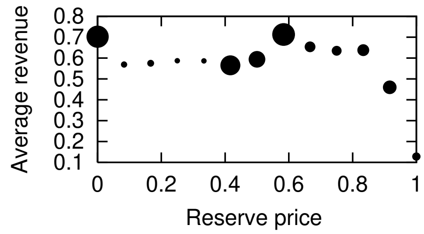

Figure 3: Learned parameters for the bandit approach. Circles represent prior experience. Each circle's height represents the initial value of *avg* at its position, while the size represents the initial value of *count*. τ_{start} = .0423, τ_{end} = .0077

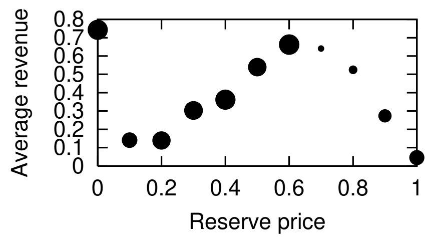

Figure 4: Learned parameters for the regression approach. Circles represent prior experience. Each circle's height represents the initial value of *avg* at its position, while the size represents the initial value of *count*. $\tau_{start} = .0081$, τ_{end} = .0013, kernel width = .138

Adaptive algorithm	Avg. 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 2: Average revenue per auction for each adaptive algorithm over 10,000 populations. Differences are statistically significant at the 99% confidence level according to paired t-tests.

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 prior experience to permit a higher initial degree of exploitation, or both.

This experiment, representing the central result of this paper, demonstrates a scenario in which an adaptive mechanism can outperform a fixed mechanism, and in which the use of metalearning can produce an adaptive algorithm better suited to the encountered population. We now present supporting experiments to further explore the properties and robustness of our approach.

5.2 A Bayesian alternative

As a result of the method described in Section 3.3 for simulating bidders, the behavior of a bidder population is specified completely by four parameters: the mean and variance

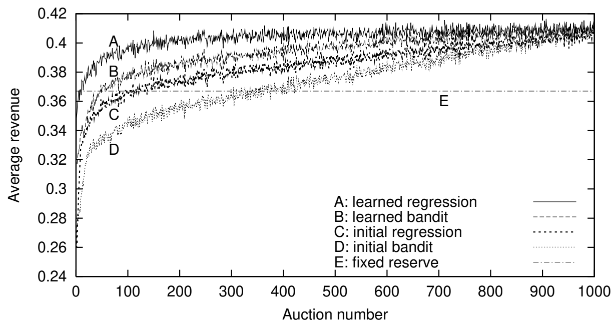

Figure 5: Average revenue per auction over the course of an episode for each method.

of the valuations, and the mean and variance of α values. If these four parameters could be determined, it would be possible to directly calculate the optimal reserve price for any bidder population encountered. This fact suggests an alternative means of adaptation: using Bayesian inference to maintain a joint probability distribution over these parameters, representing our current beliefs about those parameters as a result of observed auction outcomes.

It should be emphasized that such an approach is applicable only if a complete model of bidder behavior is available. Because a primary motivation for the use of adaptive auction mechanisms is the fact that bidders often do not behave as predicted by standard models, Bayesian inference will not generally be a suitable method of adaptation for adaptive mechanisms. We explore its use here in order to provide a comparison with the adaptive method presented in this paper, which attempts only to characterize the relationship between auction parameters and outcomes, and not to draw inferences directly about bidders. In theory, the use of Bayesian inference should allow the maximum possible use of the information contained in auction outcomes, providing a strong benchmark for comparison.

In order to perform Bayesian inference, we need to be able to determine the probability of a particular auction outcome for any setting of the four bidder population parameters. Due to the complexity of the bidder strategy described in Section 3.1, instead of determining these probabilities analytically, we precompute approximations to these probabilities by discretizing the range of possible population parameter settings, reserve prices, and outcomes, and use simulation to determine expected outcomes. 10 uniformly spaced values for each population parameter (for a total of $10,000$ possible populations) are used, along with 30 uniformly spaced values for both reserve prices and outcomes. For each choice of population and reserve price, the probability that the resulting revenue is closest to each outcome value is estimated through simulation of auctions. Increasing the granularity of discretization beyond this level does not appear to improve the results of applying these approximate probabilities.

We represent our beliefs about the bidder population as a probability distribution over all $(10,000)$ possible populations, initialized uniformly in keeping with the method of generating a population described in Section 3.3 (thus representing the prior). After each auction, the nearest of the 30 discrete reserve and outcome values are identified, the current probability for each bidder population is multiplied by the probability of observing that outcome given the reserve price, and then the distribution over populations is normalized to bring the total weight to 1, all in accordance with Bayes' rule.

To allow for the most direct possible comparison between Bayesian inference and the learning approach presented in this paper, we choose the reserve price at each step as before, by applying softmax action selection to the predicted revenues. The expected revenue for each choice of reserve price can be computed by summing the products of the probability of each bidder population and the expected revenue for that population and reserve price, also precomputed. The same temperatures are used as with the regression approach with learned parameters. Thus, the primary comparison between the two approaches involves how quickly the predictions of revenue become accurate.

The Bayesian approach results in an average revenue per auction of 0.407, slightly better than the 0.405 obtained when using the regression approach with learned parameters. (The difference is again statistically significant with 99% confidence according to paired t-tests.) Considering that the Bayesian approach relies on much stronger assumptions about bidder behavior than our approach, we feel that this small difference in performance is encouraging. Since the use of the regression approach without learned parameters results in an average revenue of 0.385, it appears that by using metalearning we are able to take nearly full advantage of the information available to us concerning the bidder population without explicitly attempting to model the population.

5.3 Unexpected bidder behavior

Although we have proposed adaptive auction mechanisms as a means of dealing with unknown bidder behavior, so far in all experiments we have evaluated adaptive algorithms on the same distribution of bidder populations assumed in developing the algorithms. We now investigate the effects when the encountered bidder population differs from these assumed populations.

We first consider the case where the bidders' strategy differs from the expected strategy. As an alternative to the equilibrium bidding strategy described in Section 3.1, consider a strategy in which a bidder simply chooses a value uniformly at random from the range $[0, x]$ and bids this value regardless of any other factors. This behavior is clearly different from that present in any of the possible bidder populations considered previously. If bidders using such a strategy were to participate in auctions based on an adaptive mechanism developed in anticipation of the previously assumed bidder population, the adaptive mechanism might not produce the results expected.

To explore the effect of unexpected bidder behavior, we consider populations in which bidders follow this alternate strategy with probability p, and follow the previous equilibrium strategy with probability $1 - p$. Three adaptive approaches are used: the regression approach with initial parameters, the regression approach with learned parameters, and the Bayesian approach of the previous section. We consider values of p ranging from 0 to 1 in order to measure the effect of increasingly unexpected behavior.

Figure 6 shows the results when x, the maximum bid under the alternate strategy, is set to 0.6, and Figure 7 shows the results when x is set to 1.2. The results are similar in both cases. (The differing slopes are due to the fact that the alternate strategy results in increased revenue in one case and decreased revenue in the other.) As the probability of bidders using the alternate strategy increases, the performance of the regression approach using the initial parameters improves relative to the other two methods, and the Bayesian approach is worst in many cases. This result can be explained by the fact that the regression approach with initial parameters makes use of no assumptions about bidder behavior, while the Bayesian approach makes very strong assumptions. With a high value of p, the Bayesian approach is never able to identify the optimal reserve price because it is unable to match the population's behavior with one of the populations it believes to be possible. The regression approach with learned parameters is eventually able to identify the optimal reserve price, but often much later than the learning approach with initial parameters, due to the fact that it is concentrating its early exploration in areas that are now less likely to contain the optimal value.

While the performance of the Bayesian approach falls below that of the regression approach with initial parameters around $p = 0.3$ for both choices of x, the performance of the regression approach with learned parameters is more sensitive to the value of x. The likely reason is contained in Figure 2. For $p = 1$, when $x = 0.6$ the optimal reserve price is 0.3, a value that is rarely optimal under the expected bidder populations, but when $x = 1.2$ the optimal reserve price is 0.6, a commonly optimal value. The regression approach with learned parameters should be expected to perform well when the results of auctions are similar to those it expects, even if the underlying bidder behavior that leads to those results is different.

The lesson from this experiment is that there is a tradeoff in attempting to use information available about potential bidder behavior. If the behavior of the encountered bidder population is significantly different than expected, an approach that makes use of such information, such as the metalearning approach presented here, may actually cause performance to suffer. Improving the ability of an adaptive auction mechanism to identify and handle unexpected behavior is an important area for future work.

Next, we consider the case where the strategy used by the bidders remains the same as before (the equilibrium strategy), but the distributions from which we draw valuations and α values are changed. Recall that these two bidder parameters were previously assumed to fall within the ranges $v \in [0, 1]$ and $\alpha \in [1, 2.5]$. We now use the narrower ranges of $v \in [.3, .7]$ and $\alpha \in [1.5, 2]$ when generating an encountered population (but not the simulated populations used in developing the adaptive algorithms). Distributions within these ranges are generated as before (see Section 3.3). Such a situation could arise if the seller was unsure of the actual ranges of v and α and intentionally used conservative estimates of these ranges when developing the adaptive algorithm.

Once again, we compare the regression approach with initial parameters, the regression approach with learned parameters, and the Bayesian approach. Table 3 shows the results. In this case, the performance of the regression approach is better when learned parameters are used, suggesting that the behavior of bidders remains sufficiently similar to make metalearning useful in this scenario. The performance of the Bayesian approach suffers even more than in the previous experiment, however. Again, this approach is

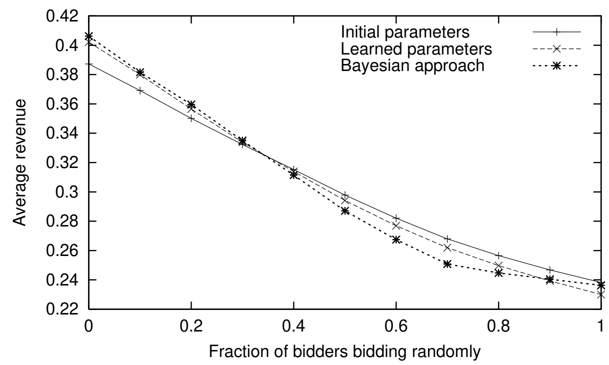

Figure 6: Results when some fraction of bidders choose a bid randomly in [0, 0.6].

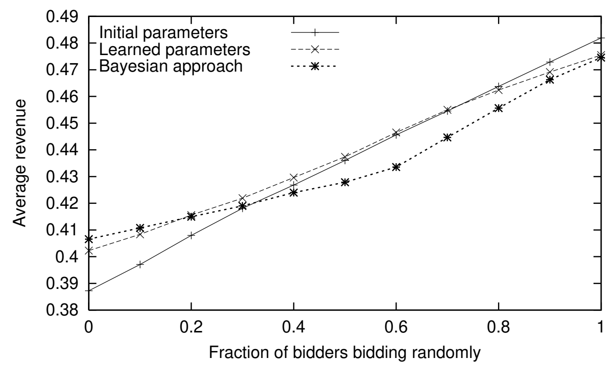

Figure 7: Results when some fraction of bidders choose a bid randomly in [0, 1.2].

often unable to match the encountered population's behavior with one of the populations it believes to be possible. This problem appears to be particularly severe when the optimal reserve price is near zero – in many such cases the population's behavior most closely resembles the behavior of populations for which a much higher reserve price is optimal, and so the Bayesian approach consistently tries a suboptimal reserve price.

5.4 Non-stationary bidder populations

In the scenarios considered previously, the behavior of the bidder population remained constant over the series of auctions. Such consistency is unlikely in the real world. Possible causes of shifts in bidder behavior include changes in the preferences of bidders and changes in the strategies employed by bidders. Bidder preferences may change gradually, such as when valuations decrease over time for outdated products, or they may change quickly, such as when a new type of bidder suddenly enters the marketplace. Changes in bidder strategies may be the result of the bidders themselves adapting by learning from past experiences.

In this section we consider modifications to our adaptive algorithm to cope with a gradually changing bidder population. Dealing with more sudden changes would likely require more extensive modifications, such as attempting to determine precisely when a change has occurred, and we leave this to future work. Recall that the behavior of our simulated bidder population is controlled by four parameters: the mean and variance of the valuations, and the mean and variance of α values. We now allow each of these parameters

Adaptive method	Avg. revenue
Bayesian approach	0.414
initial parameters	0.575
learned parameters	0.593

Table 3: Average revenue per auction for each adaptive method when bidder parameters are drawn from modified distributions. Differences are statistically significant at the 99% confidence level according to paired t-tests.

to vary according to a random walk: after each auction, each parameter is either increased or decreased by one percent of the total range for that parameter, within the bounds specified in Section 3.3. Bidder populations are initialized as before.

Under this scenario, the probability of a particular set of bidder parameter values occurring during an auction is essentially the same as before, and therefore the optimal fixed reserve remains 0.54, with an average revenue of 0.367. Using the regression-based learning approach with the parameters learned in Section 5.1 results in an average revenue of 0.379, still a significant improvement over the fixed reserve but well below the 0.405 obtained previously.

One way of dealing with a changing bidder population is to weight recent experience more heavily than older (and possibly more inaccurate) experience. We implement this modification by choosing a decay rate by which the weight of all past experience (including the prior experience) is multiplied at each step. When regression is performed, these reduced weights cause older data points to have less impact on the resulting function. When using a decay rate of 0.99, chosen manually after limited experimentation, the average revenue per auction increases to 0.383, a modest improvement over the previous 0.379.

Another possible means of improving performance with a non-stationary population is to increase the temperature τ used in selecting reserve prices for each auction. Recall that τ controls the degree of exploration when using softmax action selection, and that in the previous experiments τ was gradually decreased over time to encourage increased exploitation once reasonable estimates of auction outcomes had been learned. Because these estimates may become inaccurate as the bidder population changes, maintaining a higher degree of exploration by decreasing τ more slowly might be a way to improve performance of the adaptive mechanism. For this particular scenario, however, experiments show that the average revenue remains almost unchanged when τ is decreased more slowly (with or without the decay of experience), suggesting that the gains from increased exploration are offset by the losses from reduced exploitation. (With the stationary bidder population considered previously, decreasing τ more slowly results in a significant reduction in average revenue, so it is not simply the case that τ is unimportant.)

If the seller is aware of the specific way in which the population can change, then the seller can make use of this information by performing the metalearning step using the non-stationary population in simulation. Although the learned parameters that result (prior experience, temperatures, kernel width, and decay rate) are fairly similar to those used previously – the most notable changes are a decrease in the decay rate to 0.985 and an increase in kernel width to 0.154 – the average revenue increases to 0.389.

Adaptive method	Avg. revenue
best fixed reserve price (0.54)	0.367
previous learning parameters	0.379
previous parameters, experience decay	0.383
new metalearned parameters, exp. decay	0.389
new parameters, no initial exp. decay	0.394

Table 4: Average revenue per auction for each adaptive method with a non-stationary population. Differences are statistically significant at the 99% confidence level according to paired t-tests.

The method of decaying past experience described above includes the decay of the prior experience (the values stored before actual experience is obtained). The prior experience could instead be handled separately and have its weight remain unchanged. For reserve prices with little recent experience, the predicted revenue would then be based largely on the initial values, potentially encouraging periodic re-exploration of reserve prices previously determined to be suboptimal. For a non-stationary population, such re-exploration could possibly be beneficial. When using this approach and applying metalearning with full knowledge of how the population changes, the average revenue obtained reaches 0.394.

The results of this section, summarized in Table 4, demonstrate that our adaptive method can be applied in situations in which bidder populations vary over time, at least when such change is gradual. When knowledge of how the population might change is available, this knowledge can be usefully incorporated into the process of metalearning optimal learning parameters.

6. RELATED WORK

In addition to the theoretical work described previously ([2] and [3]), a few recent articles have begun to explore the subject of adapting auction mechanisms in response to bidder behavior from an empirical standpoint using a variety of learning approaches in simulation. In this section, we briefly survey that work and relate it to our own.

Cliff [5] 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. [13] also address continuous double auctions, using genetic programming to co-evolve buyer and seller strategies and auction rules from scratch.

Byde [4] 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.

Rogers et al. [14] provide an example of using Bayesian inference to determine optimal auction parameters. As discussed previously, such an approach is suitable when it is known that the behavior of bidders can be fully described by a number of parameters. In this case, the auction parameters to be determined are the discrete bid levels in an English auction, and the bidder parameters to be estimated based on auction outcomes are the number of bidders participating and their valuation distribution.

Dittrich et al. [7] 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.

7. CONCLUSIONS AND FUTURE WORK

In this paper, we have presented an approach to creating self-adapting mechanisms that adjust 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. The key contribution of this paper is the presentation of a metalearning technique with which information about potential bidder behavior can be used to 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. The choice of learning algorithm itself is a possible area for improvement. Instead of using softmax action selection with revenue estimates, a more sophisticated approach might involve estimating a distribution over possible outcomes and choosing parameters based on an estimate of the value of the information to be obtained in improving this estimate.

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.

8. REFERENCES

[1] C. G. Atkeson, A. W. Moore, and S. Schaal. Locally weighted learning. *Artificial Intelligence Review*, 11(1-5):11–73, 1997.

[2] A. Blum and J. D. Hartline. Near-optimal online auctions. In *SODA '05: Proceedings of the sixteenth annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1156–1163, 2005.

[3] A. Blum, V. Kumar, A. Rudra, and F. Wu. Online learning in online auctions. In *SODA '03: Proceedings of the fourteenth annual ACM-SIAM Symposium on Discrete Algorithms*, 2003.

[4] A. Byde. Applying evolutionary game theory to auction mechanism design. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 192–193. ACM Press, 2003.

[5] D. Cliff. Evolution of market mechanism through a continuous space of auction types. Technical Report HPL-2001-326, HP Labs, 2001.

[6] P. C. Cramton. The FCC spectrum auctions: An early assessment. *Journal of Economics and Management Strategy*, 6(3):431–495, 1997.

[7] D. A. V. Dittrich, W. Guth, M. Kocher, and P. Pezanis-Christou. Loss aversion and learning to bid. Technical Report 2005-03, Max Planck Institute for Research into Economic Systems, Strategic Interaction Group, 2005.

[8] A. Dodonova and Y. Khoroshilov. Optimal auction design when bidders are loss averse. Working Paper. University of Ottawa.

[9] D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.

[10] R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, feb 1981.

[11] D. Pardoe and P. Stone. Developing adaptive auction mechanisms. *SIGecom Exchanges*, 5(3):1–10, 2005.

[12] D. C. Parkes. *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency*. PhD thesis, Department of Computer and Information Science, University of Pennsylvania, May 2001.

[13] S. Phelps, P. Mc Burnley, S. Parsons, and E. Sklar. Co-evolutionary auction mechanism design. In *Agent Mediated Electronic Commerce IV*, volume 2531 of *Lecture Notes in Artificial Intelligence*. Springer Verlag, 2002.

[14] A. Rogers, E. David, J. Schiff, S. Kraus, and N. R. Jennings. Learning environmental parameters for the design of optimal english auctions with discrete bid levels. In *AAMAS 2005 Workshop on Agent Mediated Electronic Commerce VII*, 2005.

[15] M. Saar-Tsechansky and F. Provost. Active learning for class probability estimation and ranking. *Machine Learning*, 2004.

[16] G. Shmueli, W. Jank, A. Aris, C. Plaisant, and B. Shneiderman. Exploring auction databases through interactive visualization. In *Decision Support Systems*, 2005. To appear.

[17] J. C. Spall. An overview of the simultaneous perturbation method for efficient optimization. *Johns Hopkins APL Technical Digest*, 19:482–492, 1998.

[18] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 1998.

[19] R. Vilalta and Y. Drissi. A perspective view and survey of meta-learning. *Artificial Intelligence Review*, 18(2):77–95, 2002.

[20] R. J. Weber. Making more from less: Strategic demand reduction in the FCC spectrum auctions. *Journal of Economics and Management Strategy*, 6(3):529–548, 1997.

[21] P. R. Wurman, M. P. Wellman, and W. E. Walsh. A parameterization of the auction design space. *Journal of Games of Economic Behavior*, 35:304–338, 2001.