

Adaptive Auction Mechanism Design and the Incorporation of Prior Knowledge

David Pardoe, Peter Stone

Department of Computer Sciences, University of Texas at Austin, Austin, Texas 78712
{dpardoe@cs.utexas.edu, pstone@cs.utexas.edu}

Maytal Saar-Tsechansky, Tayfun Keskin

McCombs School of Business, University of Texas at Austin, Austin, Texas 78712
{maytal@mail.utexas.edu, tayfun.keskin@mcombs.utexas.edu}

Kerem Tomak

Yahoo! Inc., Santa Clara, California 95954, ktomak@gmail.com

Electronic auction markets are economic information systems that facilitate transactions between buyers and sellers. Whereas auction design has traditionally been an analytic process that relies on theory-driven assumptions such as bidders' rationality, bidders often exhibit unknown and variable behaviors. In this paper we present a data-driven adaptive auction mechanism that capitalizes on key properties of electronic auction markets, such as the large transaction volume, access to information, and the ability to dynamically alter the mechanism's design to acquire information about the benefits from different designs and adapt the auction mechanism online in response to actual bidders' behaviors. Our auction mechanism does not require an explicit representation of bidder behavior to infer about design profitability—a key limitation of prior approaches when they address complex auction settings. Our adaptive mechanism can also incorporate prior general knowledge of bidder behavior to enhance the search for effective designs. The data-driven adaptation and the capacity to use prior knowledge render our mechanisms particularly useful when there is uncertainty regarding bidders' behaviors or when bidders' behaviors change over time. Extensive empirical evaluations demonstrate that the adaptive mechanism outperforms any single fixed mechanism design under a variety of settings, including when bidders' strategies evolve in response to the seller's adaptation; our mechanism's performance is also more robust than that of alternatives when prior general information about bidders' behaviors differs from the encountered behaviors.

Key words: auction design; online learning; data-driven modeling; information acquisition

History: Accepted by Alexander Tuzhilin, Area Editor for Knowledge and Data Management; received March 2008; revised February 2009, May 2009; accepted June 2009. Published online in *Articles in Advance* October 2, 2009.

1. Introduction

Recent years have seen the emergence of numerous electronic auction platforms that cater to a variety of markets, such as business-to-business procurement and consumer-to-consumer transactions. Electronic auctions can be characterized as economic information systems used to match buyers and sellers, facilitate transactions, and provide a regulatory infrastructure (Bapna et al. 2004). A key challenge for sellers is the identification of the best design for the encountered population and the item being sold. Central to this problem, rendering auction mechanism design particularly challenging, is the behavior exhibited by bidders; bidders often show heterogeneous, arbitrarily complex, and unknown strategies. Consider, for example, the challenges sellers face when auctioning off items through an auction service such as eBay. For each auction, a seller must set various parameters

defining the particular auction mechanism's rules of exchange, so as to maximize the seller's revenue. For instance, for eBay auctions, these parameters include the duration of the auction, the start price, reserve price, and possibly a Buy It Now (BIN) price. Different auction designs can lead to widely differing outcomes. The best design is determined by bidders' behaviors, which are in turn influenced by a variety of factors, such as bidders' valuations of the item and their levels of risk aversion. Buyers of rare coins and buyers of video games are likely to exhibit distinctly different approaches to bidding, and different auction mechanisms may be appropriate for each.

Auction design has been the focus of much recent research (Blum and Hartline 2005, Blum et al. 2003, McAfee and McMillan 1987, Myerson 1981, Rothkopf and Harstad 1994). In practice, auction design has traditionally been an analytic, incremental process,

requiring several live iterations to iron out wrinkles, with mixed results (Cramton 1997, Weber 1997). Central to this process is the need to explicitly represent bidders' behaviors, often relying on theories or assumptions regarding the principles and values governing bidders' strategies. For example, assumptions often pertain to the bidders' intrinsic properties, such as valuation and rationality, as well as the manner by which these properties are manifested in bidding strategies. Assumptions about bidder rationality are particularly salient (Parkes 2001). However, such theory-driven assumptions often do not reflect the complexities of many real auction settings; bidders also apply heuristic strategies that are opaque to the seller, certainly a priori and often even after the auction.

When extensive historical data on past auctions of identical items are available and when past patterns are likely to persist, one can use analytical tools, such as the tool proposed by Shmueli et al. (2006), to derive insights that suggest effective auction parameters. For instance, an understanding of how prices change during an auction (Bapna et al. 2008a) can help improve prediction of auction revenue and then the parameters that may promote a favorable winning bid. Bapna et al. (2004) suggest that data-driven induction can be used to capture properties of bidder behaviors from historical data and propose that the induced patterns may replace theory-driven models to facilitate the design of auction mechanisms.

However, even when bidders' behaviors are known, their representation is often not amenable to closed-form analysis. One possible approach is to use available information on bidder behavior in simulation to gain intuition into favorable auction parameters. For example, Bapna et al. (2003) find that the bid increment impacts bidder behavior in ascending multiunit auctions and offer a heuristic for choosing the increment. By combining such modeled behaviors with inference, it is possible to infer which one of a set of possible behaviors bidders exhibit and then to use these behaviors to set auction parameters. Bapna et al. (2008b) use this approach to classify individual bidder strategies and estimate their valuations in multiunit ascending auctions. They are able to dynamically set bid increments during the course of an auction to improve the seller's revenue. Rogers et al. (2005) apply a similar approach to setting bid increments in repeated English auctions by using Bayesian inference to estimate bidder parameters.

Such approaches are beneficial when the historical data used to model behaviors reflect the strategies of current bidders. We present an approach that aims to allow adapting the auction design to bidders when historical data that reflect these behaviors are not available and when behaviors change over time.

In the latter case, even when the theories or induced patterns about bidders can be successfully revised periodically, the process requires human input and is time consuming, undermining the efficiency with which changes can be made to the mechanism. In e-commerce settings, such as auctions on eBay or Google keyword auctions, a large number of auctions for similar goods are held within a short time frame, and this inefficiency may present a significant drawback. Importantly, the approach we propose also allows the seller to exploit prior information about bidders, if such exists; however, it does not exclusively rely on the availability or accuracy of such information.

Several recent papers have begun to explore adaptation of electronic auction mechanisms in response to bidder behavior from an empirical standpoint, using a variety of data-driven learning approaches in simulation (Cliff 2001, Byde 2003) to produce a *fixed* mechanism. Similar to analytical mechanism design, these approaches also rely on assumptions regarding bidder behavior such as goals, beliefs, and strategies, which may differ from the actual behaviors. In addition, simulation is used to produce fixed mechanisms; however, bidders' behaviors may change over time. Blum et al. (2003) and Blum and Hartline (2005) use an online learning approach that can revise the auction design; however, the method is tailored to a specific setting, similar to a posted price mechanism, and does not accommodate arbitrary bidders' strategies or auction settings. The method also aims to maximize the worst-case performance of the seller.

In this paper, we propose a fundamentally different approach that offers a response to the challenges above. The growing electronic auction market introduces opportunities for auctioneers to capitalize on incoming data on mechanism designs and their outcomes as soon as the data become available by using online learning for mechanism design. Thus, learning and auction adaptation can occur simultaneously, over the course of a sequence of auctions. We propose a general approach that aims to allow auctioneers to effectively adapt an arbitrary set of auction design parameters to an arbitrary bidder population over a sequence of auctions, with the objective of improving an observed outcome of interest in expectation. Our approach does not rely on an explicit representation of bidders' behaviors; however, it can exploit general information that may be available about bidders to enhance the adaptation. Finally, the adaptation process can be automated, thus avoiding the inefficiencies of human input.

Because of the sequential nature of online learning, the approach we propose is particularly beneficial to sellers of items that are available periodically and for sellers of large numbers of items that can be auctioned

sequentially over time. One relevant scenario that has received much recent attention is the keyword auctions run by search engines. Whereas the mechanism used for keyword auctions varies across search engines, the fundamental notion is the same: advertisers bid to display their advertisements with search results for a certain keyword; the order in which advertisements are displayed is also determined by the ranking of the bids. As an example, consider a search engine running a keyword auction in which an advertiser submits a bid indicating the amount it is willing to pay each time its ad is clicked, and where this bid may be revised at any time. Whenever a user searches for the keyword, the bids are ranked, and three ads are displayed in the order of the top three bids, so long as each bid exceeds a reserve price of 10 cents. If an ad is clicked, the advertiser is charged its bid. Search engines have historically used the same auction design for all keywords and rarely modified this design. In our example, the auction design parameters pertain to the choice of reserve price and the number of ads to display. However, because different keywords may attract different types of advertisers, it is likely that to maximize revenue, different parameters should be used for different keywords. As an alternative to a fixed design, we propose an adaptive mechanism that for each keyword revises the design parameters periodically, in response to observed bidder behaviors. For example, for the keyword “jelly beans” our search engine might switch to a reserve price of 15 cents and see its revenue increase. (In fact, some search engines do set different reserve prices for different keywords; see Jansen and Mullen 2008. For instance, in 2008 Yahoo! switched from using a 10-cent reserve on all keywords to setting each keyword’s reserve based on the level of competition, among other factors; see *Yahoo! Search Marketing Blog* 2008.) For colas, our search engine might switch to showing just one winner of the auction if it finds advertisers are willing to bid significantly higher when none of their competitors are displayed along with them. For each keyword, the challenge for the adaptive auction is how to efficiently explore the design space so as to identify the reserve price and number of ads that maximize the revenue.

The contributions of this article are as follows. First, we present an online adaptive auction mechanism that adapts in response to observed bidder behavior. The proposed adaptive mechanism does not require explicit representation of bidders’ strategies and thus does not rely on prior knowledge of bidder behavior to infer about the profitability of different auction designs; as we will see, this property enables us to learn effective auction designs for an arbitrary setting and arbitrarily complex bidder population. Second,

our adaptive mechanism allows the seller to benefit from general prior information on bidder behavior, when such is available, to enhance online learning while maintaining the capacity to adapt online to any encountered bidders’ behaviors. As we will see, this property enables the adaptive algorithm to exploit such information without exclusively relying on it, thus reducing the risk when the encountered bidders’ behaviors differ from those anticipated by the seller. Third, using the English (ascending, open-cry) auction as our test bed, we perform extensive empirical evaluations of our mechanism’s performance as compared to alternative approaches and under a variety of conditions. We find that the auctioneer’s revenue using the adaptive auction mechanism is consistently higher than that of a fixed auction mechanism that does not adapt. Because our approach is generic and aims to improve auction designs under the complexities of real auction settings, we extend our prior work (Pardoe et al. 2006) to study the robustness of our approach under a variety of conditions encountered in real auctions. Specifically, we explore the mechanism’s performance when it encounters heterogeneous bidder populations, when the seller can alter the auction design by varying multiple auction parameters simultaneously, when the seller encounters bidders whose strategies have evolved to respond specifically to the adaptive algorithm’s form of adaptation, and when bidder participation in the auction is dynamically revised in response to changes in bidder surplus. Because the proposed mechanism allows incorporating general prior knowledge to affect the exploration of designs, we also study the performance of the mechanism when it is given incorrect prior information on bidder behavior and when bidders’ strategies vary over time.

2. Data-Driven Adaptive Auction Mechanisms

In this section we present the settings we address for auction mechanism design and our approach to adapting the design online in response to bidders’ behaviors.

2.1. Framework

Consider an auctioneer who has n identical items to be auctioned off in a series of sequential auctions, henceforth referred to as an *episode*. The auctioneer’s objective is to maximize the total benefit over the course of an episode. The seller’s benefit may be any observable auction outcome of interest, such as the seller’s revenue or bidder participation. Henceforth, we assume the seller’s objective is to maximize revenue over the course of the episode. To improve her revenue, the auctioneer may revise the parameters of

an auction mechanism that define the mechanism's mode of operation, such as the auction duration, reserve price, and start price; we define an auction design as a unique set of auction parameters that are selected by the seller. We assume that no prior knowledge is available that would enable complete specification of the bidding strategies individual bidders will exercise. Later, we will present our approach to incorporate general prior knowledge about the bidder population. In the absence of an explicit formulation of bidders' behaviors, we infer the benefit to the seller from different designs via a data-driven mapping $M: x \rightarrow y$ from a set of k distinct auction designs $x \in D = \{d_1, d_2, \dots, d_k\}$ onto a real-value $y \in R$, where y refers to the observed auction outcome of interest. The seller "learns" the mapping over time by experimenting with different designs and incorporating new information on the profitability of these designs as it becomes available to revise the induced mapping M . We refer to this process as *data-driven adaptive mechanism design*: an online empirical process whereby the auctioneer obtains information through experimentation and adapts the mechanism's design to maximize an observed objective of interest.

2.1.1. A Desirable Exploration and Exploitation Schedule. To simplify the exposition of our approach we assume, for now, that the auction design is determined by a single design parameter—the auctioneer's reserve price. We also assume that the seller can choose one of k design choices for the reserve price at each auction. As such, our problem can be defined as selecting a sequence of auction designs that maximize the auctioneer's revenue over the course of an episode. This problem can be mapped to the classic multi-armed bandit problem introduced by Robbins (1952): a gambler has to decide which of k slot machines to play in a sequence of trials so as to maximize the overall reward. Each unique auction design selected in each auction during the episode corresponds to a slot machine in the bandit problem—in each auction our auctioneer must decide which design to select with the objective of maximizing her overall benefits over the entire episode. Because the expected profits of different designs are not known a priori, these profits can be estimated by the seller using different designs in auctions and obtaining their outcomes.

Given the seller's objective, consider the following two extreme strategies. The first is a greedy *exploitation* strategy, where at each auction the seller selects the design estimated at that time to be most profitable. If the seller's estimation of the profitability of designs is accurate, this strategy guarantees maximizing the seller's revenue. However, when the seller's estimation is imprecise, this strategy might select designs that yield lower profits than can be derived

from other designs. Importantly, an exploitation strategy will also not tend to improve the estimation of designs that are not currently considered profitable by the auctioneer. Yet when the seller's estimation is poor, one of these designs might have proved particularly advantageous had it been attempted by the seller. The other extreme approach is *exploration*. Similar to active learning (Saar-Tsechansky and Provost 2004), exploration pertains to selecting different designs during the course of the episode so as to acquire information about the outcomes of these designs and improve the seller's estimation of the profitability of these designs. When the auctioneer's estimation of the profitability of designs is poor, such sequential exploration may help to ultimately identify a particularly profitable design. However, this information acquisition strategy may also be costly to the seller, because until a profitable design is identified the seller may select many designs that yield poor profits, undermining the seller's overall revenue. Thus, when selecting the next auction's design the auctioneer faces a trade-off between exploring the design space to improve her estimation and exploiting the design currently estimated to be most profitable.

As discussed above, both extreme strategies have clear benefits and deficiencies, depending on the accuracy of the seller's estimation of the profitability of alternative designs—some exploration is necessary albeit costly when the seller's estimation is inaccurate, and exploitation is preferable when the seller's estimation is correct. Because the seller's estimation of design profitability is data driven and will gradually improve as the episode progresses, it would be desirable to prefer exploration of different designs initially and then to gradually increase the likelihood of selecting designs estimated to be profitable by the seller, as the seller's estimation improves.

2.1.2. Design Selection and the Bandit Adaptive Algorithm. To manage this trade-off effectively we use an adaptive algorithm to determine how the auction design should be selected in response to prior auction outcomes. We implement the adaptive algorithm using a method known as *softmax* action selection (Sutton and Barto 1998), which realizes a hybrid of the two extreme exploitation and exploration strategies discussed above. The fundamental notion underlying our adoption of the softmax action selection is that at each auction, the design suggested by the exploitation strategy as well as designs that align with the exploration strategy may be selected with some likelihood. Given all else is equal, as the seller's estimation of the benefit from different designs improves over the course of the episode, the probability of selecting designs estimated by the seller to be profitable increases, and designs considered to be less profitable are explored less frequently.

Formally, for each design i , avg_i^t denotes the expected revenue estimated for this design at time t during the episode, and $count_i^t$ denotes the number of times design i has been tried up to time t . In each auction t during an episode, the seller draws the design for the auction from a distribution D . Specifically, the probability of drawing design i at auction t is given by the following Boltzmann distribution $D_t(i) = e^{avg_i^t/\tau_t} / \sum_{j=1}^k e^{avg_j^t/\tau_t}$, where τ_t is a positive number referred to as *temperature*. Let us now discuss the two factors that determine the likelihood of selecting design i for auction t —the expected revenue from a given design and temperature. As shown in the equation above, the probability of selecting a given design is proportional to the estimated expected revenue from the corresponding design. Thus, the design estimated to be most profitable at any given time, and which would have been selected by an exploitation strategy, has the highest probability of being selected; however, each of the other designs can be selected as well, with a lower probability, enabling the seller to explore the space of designs. To see how temperature determines the extent to which exploitation trumps exploration, note that when temperature is set to a very large number, the distribution D will resemble a uniform distribution, which encourages exploration of a variety of designs. As such, a high temperature value discounts the effect that the seller’s estimations of design profitability have on the probability of drawing a design. This is indeed beneficial when the seller’s estimation of avg_i^t is poor. Similarly, setting a small temperature value increases the likelihood that a design estimated to be most profitable will be drawn, encouraging exploitation. The latter would be desirable when the seller’s estimation is highly accurate. We therefore gradually encourage exploitation as the episode progresses and the seller’s estimation improves by lowering the temperature parameter. Specifically, as shown in Algorithm 1, we set start and (lower) end temperatures, denoted τ_{start} and τ_{end} , respectively, and vary the temperature in each auction during the episode by interpolating linearly.

Algorithm 1 (Adaptive bandit)

- Input:** A finite set I of size k of auction designs, τ_{start} , τ_{end} , $\forall i \in I$: $count_i^0$ and avg_i^0 , the number of auctions to run A , current auction number $t = 0$
- 1 **While** episode has not terminated:
 - 2 $\tau_t = \tau_{start} + (\tau_{end} - \tau_{start}) \cdot \frac{t}{A}$
 - 3 $\forall i \in I$ set the probability of drawing design i as $D^t(i) = e^{avg_i^t/\tau_t} / \sum_{j=1}^k e^{avg_j^t/\tau_t}$
 - 4 Draw auction design j from the distribution D^t and run an auction to obtain outcome o
 - 5 $\forall i \in I$ set $avg_i^{t+1} = avg_i^t$ and $count_i^{t+1} = count_i^t$
 - 6 $avg_j^{t+1} = [(avg_j^t \cdot count_j^t) + o] / (count_j^t + 1)$

- 7 $count_j^{t+1} = count_j^t + 1$
- 8 $t++$
- 9 **End While**

The pseudocode for this form of adaptation mechanism, henceforth referred to as *bandit adaptation*, is presented in Algorithm 1. Note that the adaptive algorithm can be used to revise any number of auction design parameters simultaneously for any auction platform. In addition, whereas most often the Boltzmann (also known as Gibbs) distribution we use here is employed for the softmax action selection, this distribution is merely a vehicle to obtain the exploitation/exploration hybrid properties described above. In principle, other distribution functions may lead to these properties as well.

The adaptive algorithm uses parameters that need to be set at the outset. For example, before the auctioneer can estimate the profitability of different designs by acquired real auction outcomes through experimentation, it is necessary to decide what initial expected revenue values and counts to use for each choice of auction design and what initial and end temperature values to use. In §2.3 we propose an approach to setting these parameters that we evaluate subsequently in §3 as compared to an existing benchmark approach.

2.2. Information Sharing

In the bandit adaptive algorithm described above, we have implicitly assumed that the expected revenue of each choice is independent of the revenues of other designs. This entails that information on the benefit of a given design can only be acquired via experimenting with that design. However, the expected revenues of designs with similar reserve prices are likely to be similar; hence, it may be possible for experience to be profitably shared among choices.

We propose to capitalize on this notion of information sharing to infer the profitability of designs that have not been explicitly attempted by the auctioneer and to improve the auctioneer’s overall estimation. Whereas in the bandit adaptation the auctioneer maintains a mapping from each distinct auction design to an expected revenue, we now induce a mapping from the continuous domain of auction parameters $X \in \mathfrak{R}^N$ onto revenue: $M': X \in \mathfrak{R}^N \rightarrow \mathfrak{R}$. This mapping enables one to make predictions of auction outcomes for any auction parameters over this continuous domain and thus infer about a larger set of alternative designs for any given number of experiments and corresponding information acquisition costs. The revenue estimation for a given design may also improve as more auctions take place, even when the corresponding design has not been attempted again by the auctioneer.

In the empirical evaluations below, we induce this mapping from the auctioneer's prior experiences via locally weighted quadratic regression (LWQR) (Atkeson et al. 1997). This approach combines the simplicity of classical least-square regression with the flexibility of nonparametric modeling—LWQR allows us to induce a mapping for which no theoretical models exist. In addition, the induction can be automated because it does not require the human input often necessary in classical regression models. Indeed, no theory exists regarding the mapping of an arbitrary set of auction design parameters onto revenue, and this mapping can be arbitrarily complex. As such, LWQR is ideal for online, automated settings. In LWQR each experience is weighted by its distance d from the design for which profit is inferred using a Gaussian kernel of width w as follows: e^{d/w^2} .

As in the bandit adaptation, at the beginning of the auction and before actual auction outcomes can be used to induce the mapping M , it is necessary to have an initial mapping to be used by the auctioneer towards the initial induction. Similar to the bandit approach, we choose k distinct designs d_i and determine avg_i and $count_i$ for each design, representing prior experiences. These experiences populate the initial training set T with k tuples of the form: $\text{tuple}(d_i, avg_i, count_i)$, where avg_i refers an outcome that has been observed for design i , and $count_i$ refers to the number of times that outcome has been observed for this design. (Note that avg_i and $count_i$ need not be integers.) Once actual auction outcomes become available during the episode, both prior experiences and actual experiences obtained during the course of an episode are used to induce the mapping M . The pseudocode for the adaptive sharing algorithm is presented below.

Algorithm 2 (Adaptive sharing)

Input: A finite set I of size k of auction designs, τ_{start} , τ_{end} , w , an initial data set T , the number of auctions to run A , and current auction number $t = 0$

- 1 **While** episode has not terminated:
- 2 Apply LWQR to data set T to induce a mapping $M: X \rightarrow O$
- 3 $\tau_t = \tau_{\text{start}} + (\tau_{\text{end}} - \tau_{\text{start}}) \cdot \frac{t}{A}$
- 4 $\forall i \in I$ set the probability of drawing design i as $D^t(i) = e^{M(i)/\tau_t} / \sum_{j \in I} (e^{M(j)/\tau_t})$
- 5 Draw auction design i from the distribution D^t and run an auction to obtain outcome o
- 6 Augment data set T with the tuple $(i, o, 1)$:
 $T \leftarrow T \cup (i, o, 1)$
- 7 $t++$
- 8 **End While**

2.3. The Incorporation of Prior Knowledge via Metalearning

When historical data are not available for sellers to derive the particular distributions that characterize bidder properties, sellers often have more general information, such as the range in which the mean valuation may lie. Consider, for example, a seller who has a large number of copies of a newly published book to be sold sequentially through a series of auctions. Although there are no previous auction outcomes for this book to guide the choice of auction design, the seller is not completely ignorant—auction results are available for books by similar authors or on similar topics and that have likely attracted similar bidder populations. The seller may also receive marketing information suggesting valuations that buyers might have for the book, perhaps along with bidding strategies that bidders tend to use in such auctions. Such information may be consistent with many different plausible bidder populations. In this section we posit that it is useful to consider how such general information can be effectively used by sellers while maintaining the ability to adapt when this information is incorrect.

A potential approach that we consider in §3 is to use general information to set the auction design parameters so as to derive a fixed design that is beneficial *in expectation* over all plausible populations that the seller considers. This approach avoids the costly exploration that is necessary for data-driven adaptation to identify advantageous designs. However, a fixed design may be susceptible to incorrect prior information. Even when the seller's information is accurate, a design selected to perform well for a variety of plausible populations in expectation may not be best for the particular population the seller encounters.

Alternatively, we propose a principled approach, which we refer to as *metalearning*, and that aims to incorporate general prior information into the data-driven adaptive algorithm itself. Metalearning's objective is to enable the seller to benefit from prior knowledge to adapt more efficiently given the plausible populations, such as by reducing costly exploration of unfavorable designs early in the episode while continuing to use real-time auction outcomes to allow adapting the design online to any encountered population. As we will see, maintaining the ability to adapt is fundamental to robust performance in difficult cases, such as when the population varies over time. Specifically, both the bandit and shared adaptive algorithms use *adaptation* parameters—the expected revenue avg_i^0 and number of attempts $count_i^0$ for each design in the bandit adaptive algorithm, the corresponding parameters avg_i and $count_i$ in the sharing adaptive algorithm, the initial and end temperatures

τ_{start} and τ_{end} , the kernel width w for the adaptive sharing algorithm, and the design space discretization parameter k . Because these parameters affect the manner by which the adaptive algorithm adapts the design to the encountered population, they offer an opportunity to incorporate prior knowledge to adapt more effectively in expectation given the plausible bidder populations consistent with prior knowledge. In particular, it is possible to identify the adaptation parameters for which the adaptive algorithm yields the highest revenue per episode, in expectation, over all plausible bidder populations. Whereas the adaptation parameters will benefit adaptation for the populations considered plausible, the continued use of real auction outcomes to revise the design allows the algorithm to adapt to any encountered population, even if it differs from those consistent with prior knowledge.

To illustrate how metalearning can achieve these goals, consider Figure 1, which presents the frequency with which different designs are optimal over 10,000 different plausible populations of bidders, as well as the expected revenue over all populations for each design, in a scenario we will describe below. Two key observations can be made. First, as shown by the frequency in which different designs are optimal, almost every design is optimal for some population. Because the seller’s general information gives rise to many plausible populations and because prior information may be incorrect, it is beneficial to use the adaptive algorithm to adapt online to the particular population encountered by the seller. Second, whereas most designs are optimal for some populations, for many different populations the optimal reserve price is in one of two small regions. This suggests that general information that gives rise to a variety of plausible populations can be beneficial to focus the adaptive algorithm’s exploration on a few promising designs early on in the episode. This can be achieved through the adaptation parameters, such as by assigning high initial expected revenue to

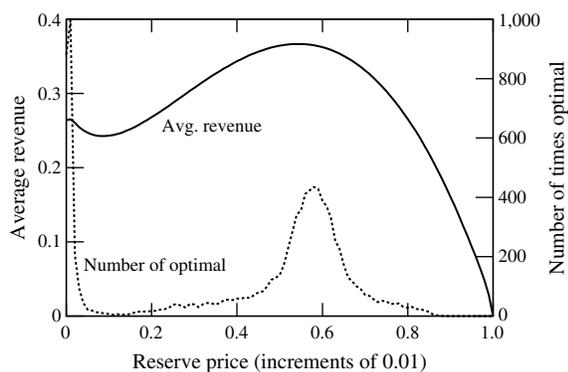


Figure 1 Revenue and Frequency with Which Designs Are Optimal for 10,000 Different Populations

such promising designs, yielding a higher probability of selecting these designs early in the episode. In general, other adaptation parameters can also affect other aspects of the adaptive algorithm to render it more efficient given the plausible populations. For example, a higher initial temperature can be beneficial when the revenues produced for the same design vary largely for the plausible populations, increasing initial exploration of designs to help identify profitable designs. Similarly, when the revenue from a given design exhibits high variance, it would also be beneficial to assign a higher value for $count_i^0$ in the bandit adaptive algorithm so as to require substantial exploration to change the expected revenue estimation of designs, avg_i^t . Importantly, although incorporating prior knowledge into the adaptation algorithm itself will allow more efficient adaptation when the seller’s information is correct, the seller still maintains its ability to adapt online to any encountered population, even if it is different than those anticipated.

The potential complexity and stochastic nature of the auction setting also affects our choice of method to search for the adaptation parameters. In particular, because of the many possible factors affecting the relationship between the adaptation parameters and the objective function, we are restricted to sampling the results obtained from any particular design rather than analytically determining expected results. Our search problem can thus be viewed as a stochastic optimization problem, entailing an optimization method that does not rely on direct measurements of any derivatives of the objective function. In the empirical evaluations we present here, we use simultaneous perturbation stochastic approximation (SPSA) (Spall 1998) to identify the best adaptation parameters. SPSA is based on gradient approximation, where at each step two estimates of the expected episode revenue are taken for slight perturbations of the current parameters (the same, randomly chosen bidder population is used for each estimate), a gradient approximation is found, and the parameters are updated in the direction of the gradient. Note that SPSA requires an estimate of the expected revenue. To accommodate arbitrarily complex auction settings and stochastic bidder behaviors, it is possible to estimate the revenue in expectation via simulation. Given prior knowledge of the distribution over different possible populations, the parameters of a specific bidder population are drawn from this distribution; the adaptive mechanism with a given set of parameters is then applied to this simulated population to estimate the seller’s revenue over an entire episode. We will later discuss how the seller constructs distributions that are consistent with prior knowledge. Note that bidder behaviors that are consistent with prior knowledge and used for stochastic optimization may indeed differ from those of the

encountered population. In §3 we study the performance of the adaptive algorithm with metalearning as compared to alternatives when the seller's prior information is incorrect and also when bidders' strategies change over time. It is important to note that metalearning is performed only once, offline, prior to the actual episode and before online adaptation takes its course.

3. Empirical Evaluations

To study our approach 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. The seller has n identical items. For now, we assume the seller auctions the items one at a time to bidders from the population through a series of auctions. We present results for multi-item auctions in §3.6.

To facilitate the analysis of the performance of our approach in the initial set of experiments we assume that the seller can set the reserve price for each auction. In §§3.4–3.6, we evaluate our approach for auction designs determined by additional parameters. The seller's goal is to set the reserve price for each auction so that the total revenue obtained from all the auctions during an episode is maximized.

Our motivation to explore the use of a data-driven, adaptive auction stems from the complex and often irrational behaviors exhibited by bidders. In the initial set of evaluations we study the adaptive mechanism's performance when the encountered population consists of loss-averse bidders (Dodonova and Khoroshilov 2005), shown to explain empirical results that are inconsistent with bidder's rationality. We provide a detailed description of the equilibrium strategy in §A of the Online Supplement, available at <http://joc.pubs.informs.org/ecompanion.html>. In §§3.3–3.6, we report the performance of our mechanism when bidders exhibit other conceivable strategies that have been reported in prior work, as well as when the population includes bidders exhibiting a variety of strategies simultaneously. A loss-averse bidder is characterized by an independent, private value v for the sold item drawn at random from a normal distribution within the range $[0, 1]$ and a level of loss aversion α drawn at random from a normal distribution within the range $[1, 2.5]$. We assume that a given bidder assigns the same value to any one of the identical items sold during an episode. For now, we also assume that the population does not change during an episode. In §3.4 we explore the mechanism's performance when the bidder population changes during an episode.

3.1. Comparison to Alternative Learning Approaches

We begin by evaluating the performance of the bandit and sharing adaptive algorithms compared to the best fixed design. For this experiment, we evaluated the seller's performance over 10,000 episodes, with a different, randomly drawn bidder population in each episode. To identify the best fixed design, we selected the reserve price yielding the highest average revenue over all episodes.

For now, we let the populations assumed plausible by the seller be the same as the actual distribution of bidder populations; thus, although the seller does not know which particular population will be encountered, the seller does have correct general information regarding the populations that might be encountered, their probabilities of occurring, and the strategies that bidders will use. We also derive the best fixed design using the same information. Later we study the performance of our approach when these assumptions are violated. Specifically, we first assume that the seller knows that bidders' mean valuations are in the range $[v_{\min} = 0, v_{\max} = 1]$, respectively, and that the mean level of loss aversion is within $[\alpha_{\min} = 1, \alpha_{\max} = 2.5]$. To perform stochastic optimization during metalearning, for each episode being simulated, the seller randomly generates an "arbitrary" distribution for valuations by taking a Gaussian with a mean drawn uniformly from $[0, 1]$ and a variance of 10^x with x drawn uniformly from $[-2, 1]$. The distributions are then normalized so as to produce a distribution over the range $[0, 1]$. A distribution for the level of loss aversion is generated in the same way, drawing variance as before and using a range of $[1, 2.5]$ for both the mean and the entire distribution.

We also consider an alternative approach to metalearning for setting the adaptation parameters for the bandit and sharing adaptive algorithms. In doing so, our objective is to examine whether the adaptation parameters have a significant impact on the adaptive algorithm's performance and, if so, whether metalearning is advantageous as compared to common practices for setting online learning algorithm's parameters. We use an approach known as *optimistic initialization* (Sutton and Barto 1998), commonly used to set the initial expected values of alternative actions and the corresponding counts used in the softmax action selection we adopt here. Optimistic initialization aims to encourage early exploration by setting all initial averages to the same value, which is higher than the estimated value of the largest possible value. In our setting this approach corresponds to the assumption that all designs are equally likely to be beneficial, and thus each design is likely to be selected at least once near the beginning of the episode. In addition,

each count is assigned the same small but nonzero value; thus, several negative experiences with a design are required to significantly reduce its probability of being selected, $D^i(i)$, during the softmax action selection step. To our knowledge, there are no alternative principled approaches to set the remaining adaptation parameters. To complement optimistic initialization, we followed common practices and performed limited experimentation using plausible bidder populations to set all other adaptation parameters. Using optimistic initialization, we set each avg_i to a high value of 0.6 and each $count_i$ to 1, and we use $\tau_{start} = 0.1$, $\tau_{end} = 0.01$, and $w = 0.1$. We set $k = 13$ and $k = 11$ for the bandit and sharing adaptive algorithms, respectively.

Finally, in addition to comparing the adaptive algorithm to the best fixed design, we explore the relative contributions of metalearning and of information sharing to the seller’s performance. We therefore report the performance of the bandit and sharing adaptive algorithms when the adaptation parameters are determined by either optimistic initialization or metalearning. We refer to information sharing with metalearning and with optimistic initialization as *learned sharing* and *optimistic sharing*, respectively, and to the bandit approaches as *learned bandit* and *optimistic bandit*.

Figure 2(a) shows the average revenue per auction over the course of an episode for each approach. The difference in the average revenue between each of the adaptive methods and the fixed design is statistically significant ($p < 0.01$), according to paired t -tests. As shown, the average revenues obtained by all the adaptive mechanisms are higher than the revenue produced with the best fixed design for the plausible populations. Of the four adaptive mechanisms, the seller’s average revenue is higher when *information sharing* is used as compared to the bandit approach and even higher when metalearning is used together with information sharing.

Figure 2(b) presents the average revenue in each auction over the course of an episode obtained with each approach. For a given initialization method, either optimistic or learned, information sharing con-

sistently produces higher revenue for any given number of experiences. In addition, when metalearning is used with information sharing (learned sharing), the adaptive mechanism requires fewer than 20 auctions to learn a more profitable auction design than the fixed auction design. The cumulative revenue produced by the learned sharing algorithm exceeds that of the optimal fixed design after 35 auctions.

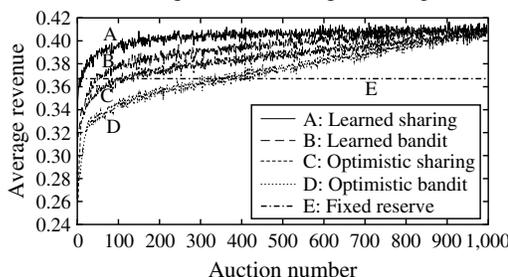
The initial average revenue and count values produced during the metalearning phase for the information sharing approach are displayed visually in Figure 2(c). Each of the reserve prices considered towards the initial training data T is represented by a circle; the area of the circle is proportional to the initial $count_i$ parameter, and the y -axis shows the initial average revenue value. As shown, some reserve prices have both high average revenue and large counts, leading to higher initial probability of being selected. As shown in Figure 2(c), metalearning identified several promising designs to select and evaluate early in the episode, which lead to higher revenues in the early auctions as compared to the optimistic approach.

In summary, our results suggest that online adaptation to the encountered population during the course of an episode can outperform the best fixed design. In addition, by sharing auction outcomes among designs, the sharing adaptive algorithm better estimates the expected revenue of alternative designs and thereby significantly improves the adaptive algorithm’s choices of designs throughout the episode. Finally, the adaptation parameters significantly affect the performance of the adaptive algorithm; metalearning can effectively exploit general prior information to identify advantageous adaptation parameters to improve online adaptation and to yield higher revenues than can be obtained with optimistic initialization. Next, we present studies that explore the trade-offs from using general prior information for mechanism design and evaluate the robustness of our approach under different settings. Henceforth, we use only information sharing for adaptation, and refer to the adaptive algorithm with information sharing simply as the adaptive algorithm.

(a) Average revenue per auction

Mechanism	Average revenue
Fixed design	0.367
Optimistic bandit	0.374
Optimistic sharing	0.385
Learned bandit	0.394
Learned sharing	0.405

(b) Average revenue throughout an episode



(c) Initial averages with metalearning

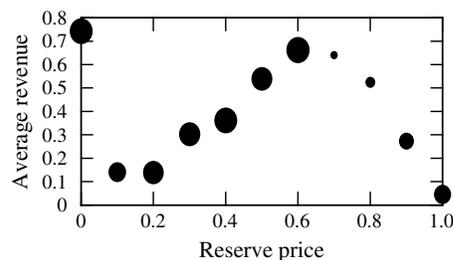


Figure 2 Comparison of Alternative Approaches

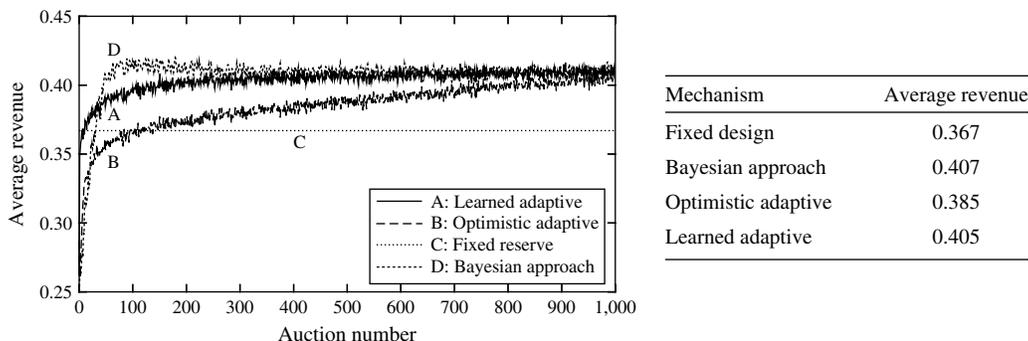


Figure 3 Average Seller Revenue Produced in Each Auction During an Episode

3.2. A Comparison with a Bayesian Approach

The adaptive algorithm's online exploration of designs during the course of the episode is guided by a data-driven mapping between the auction design and the outcome of interest, such as revenue. In this section, we consider an alternative form of online adaptation. Specifically, in place of a data-driven mapping to guide the choice of alternative designs, we let the adaptive algorithm rely on prior knowledge of the exact strategies bidders will use to estimate the expected revenue of possible designs; data produced from online interactions with bidders are used only to estimate, via Bayesian parameter estimation, the parameters governing bidder's strategies. While in the absence of historical data, knowledge of bidders' strategies may not be a realistic assumption, in the following sections we contrast this form of adaptation with our approach to study the trade-off from the use of prior knowledge under different conditions and the robustness of each approach when prior information is incorrect. Details of the Bayesian approach are provided in §B of the Online Supplement.

Compared with the optimistic and learned adaptive algorithms that do not rely on knowledge of bidders' strategies and plausible populations to adapt to the encountered population, the *Bayesian* adaptive algorithm's estimation of auction outcomes has lower variance and higher bias (Friedman 1997)—the Bayesian approach's estimations of expected revenue can only yield values that are feasible given the assumed bidder's strategies and plausible bidder populations; by contrast, improvements in the adaptive algorithm's estimation are entirely data driven and do not use this bias. Thus, when the seller's prior information and thereby the Bayesian approach's bias are correct, the Bayesian approach is likely to exhibit faster improvement in performance earlier in the episode compared with the adaptive algorithm, which uses only limited real auction outcomes to improve its estimation. At the same time, the Bayesian approach's estimation bias may undermine its performance when

the seller's information on bidder's strategies is incorrect, and its bias is wrong. We next explore the trade-off from the use of prior information by the Bayesian approach and the adaptive algorithm.

We begin our evaluation with a seller that has correct prior information on bidders' strategies and plausible bidder populations. We later revisit this comparison for settings in which certain properties of buyers' bidding behaviors differ from those assumed by the seller. Figure 3 shows the average revenues generated from each auction in the episode over 10,000 episodes using the Bayesian approach (using the same temperatures as the learned adaptive approach), the adaptive mechanism with learned and optimistic parameters, and the best fixed design. Let us first compare the Bayesian approach (curve D) with the optimistic adaptive approach (curve B), which uses only real auction outcomes to improve its estimation of the profitability of designs. From Figure 3 we see that the performance of the Bayesian approach improves more quickly because it uses accurate prior knowledge on bidders' strategies and plausible parameter distributions to estimate which population is being encountered and subsequently to infer the expected revenue of each design. The average revenue per auction produced over the course of the entire episode with the Bayesian approach is 0.407 versus 0.385 with the adaptive approach with optimistic parameters. (Note a slight decrease in the performance of the Bayesian approach after around 100 auctions. Because we consider a discretized space of bidder populations, the Bayesian approach's estimation may converge to slightly suboptimal parameters.)

The learned adaptive approach (curve A) exploits the seller's prior knowledge in a metalearning phase to set the adaptation parameters prior to the online adaptation. As can be seen in Figure 3, it does so effectively and exhibits higher revenue early in the episode. Similar to the optimistic adaptive approach, during online adaptation the learned adaptive algorithm's performance can improve only with improvements in its estimation of the expected revenue from

each design; because these improvements in estimation are data driven and rely exclusively on real auction outcomes, the learned adaptive algorithm exhibits slower improvements compared with the Bayesian approach. Yet the learned adaptive algorithm’s average revenue over an entire episode is only slightly lower than that obtained with the Bayesian approach. (The difference is statistically significant at $p < 0.01$, according to paired t -tests.) Although the learned adaptive algorithm relies exclusively on real auction outcomes to revise the design selection distribution D , its ability to exploit prior information through metalearning enables it to achieve performance that is close to that obtained with the Bayesian approach, which depends on knowledge of bidding strategies and plausible parameter distributions to adapt. Next, we will explore the robustness of these approaches when the seller encounters unexpected bidder behavior.

3.3. The Case of Incorrect Prior Information

We now explore the performance of each mechanism when the encountered bidder population differs from the populations considered plausible by the auctioneer. Recall that the adaptive algorithm relies on information on bidders’ strategies and plausible parameter distributions to adapt. We therefore consider two scenarios in which these elements of bidder behavior are altered so as to differentiate the encountered bidder population from those anticipated by the seller. We first consider the case where the bidding strategies differ from the expected ones—while the auctioneer assumes a loss-averse equilibrium bidding strategy, she encounters a mixed population of loss-averse bidders and of bidders that simply bid a value drawn uniformly at random from the range $[0, 1.2]$. Bidders follow this alternate strategy with probability p and follow the loss-averse equilibrium strategy with probability $1 - p$. To evaluate the effect of unexpected behavior on the performance of each approach, we change the rate of unexpected behavior by increasing the proportion p of bidders who exhibit this unexpected behavior.

Figure 4 shows the revenue generated by each approach for different proportions of bidders with unexpected behavior. Because the unexpected strategy is more profitable for the seller, as the proportion of bidders using this strategy increases the auctions become more profitable. The challenge is to take full advantage of the increased revenue potential by identifying early on the best design for the encountered population. As shown in Figure 4, the Bayesian approach and the optimistic adaptive approach each perform worse under some conditions: the Bayesian approach leads to poor selection of designs as more bidders use the unexpected strategy. This is because

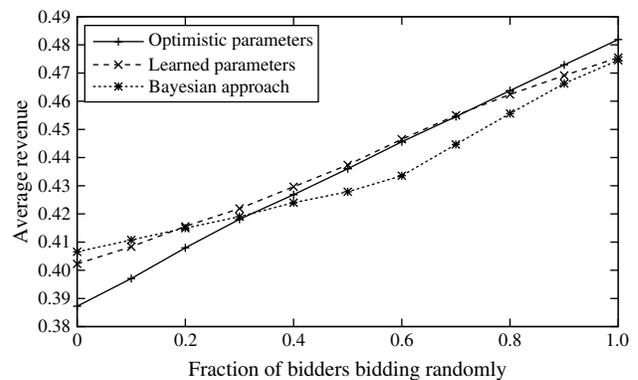


Figure 4 Average Seller Revenue When Bidders Exhibit Unexpected Strategies

it assumes a different behavior than the encountered one to estimate the profitability of alternative designs. As before, the adaptive approach with optimistic parameters yields lower average revenue when accurate prior knowledge is available about bidders’ strategies because it does not use this knowledge to inform its exploration of designs; contrary to these two approaches, the adaptive algorithm with learned parameters via metalearning exhibits a more robust performance and is never the worst approach. It is also never more than 0.0065 behind the best approach, whereas the other approaches fall more than twice as far behind the best approach for some settings. The learned adaptive approach strikes a useful balance between the use of prior information and data-driven estimation. Through metalearning, the learned adaptive approach can benefit from prior information when the information is correct; however, it also continues to use real auction outcomes to improve its estimation of the expected benefits of designs, and is thus able to adapt to any encountered population, even when the population differs from those considered plausible.

Next, we let all bidders employ the expected loss-averse equilibrium bidding strategy; however, the distributions from which valuations and bidders’ loss aversion levels are drawn differ from the expected ones. Specifically, the auctioneer assumes the ranges of valuations and of bidders’ loss aversion levels are $[0, 1]$ and $[1, 2.5]$, respectively, although the true ranges are $[0.3, 0.7]$ and $[1.5, 2]$, respectively. (Such a situation could arise if the seller uses liberal estimates of these ranges when developing and biasing the adaptive algorithm.)

Figures 5(a) and 5(b) show the average revenue in each auction and average revenue over a complete episode produced by each approach. As shown, the performance of the Bayesian approach suffers even more in this setting than when bidders’ bidding strategy is unexpected. The estimated expected profitability of designs inferred by the Bayesian approach

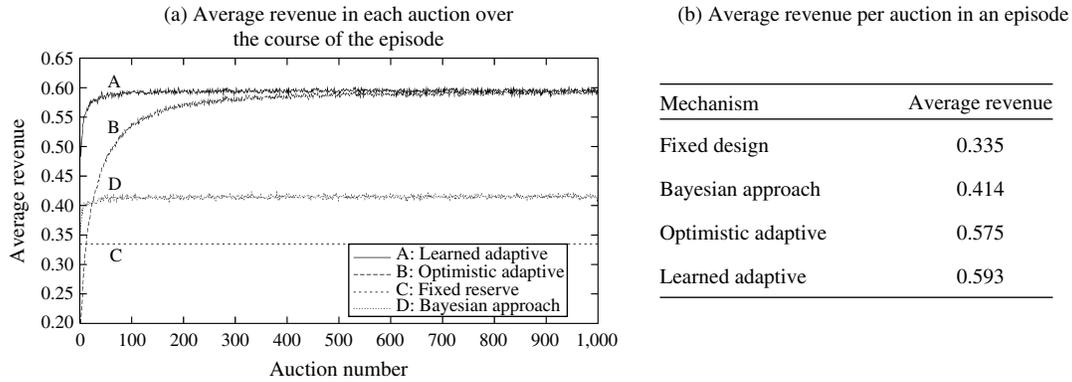


Figure 5 Revenue When the Seller's Prior Information Is Inaccurate

is derived from assumed behaviors that substantially differ from those actually encountered by the seller. By contrast, the performance of the adaptive approaches is significantly better. In this case, the learned adaptive approach also produced higher revenue than the optimistic adaptive approach—as shown in Figure 5(a), prior knowledge enabled the learned adaptive method to focus on more profitable designs early than designs selected by the optimistic adaptive algorithm and which are drawn uniformly at random. The improvements shown by the adaptive mechanisms over the Bayesian approach demonstrate that designs that prior knowledge suggest are profitable were found by the data-driven mapping to be profitable, and thus selected by the adaptive mechanisms. Finally, the overall average revenue of the learned adaptive algorithm exceeds that of the best fixed design derived from prior knowledge after the very first auction.

Our results show that there is a clear trade-off in the use of prior information about potential bidder behaviors and that this trade-off can be managed effectively. An approach that relies on knowledge of bidders' behaviors to adapt identifies profitable designs faster than purely data-driven approaches when the seller's prior information is correct but is also highly susceptible to incorrect prior information. A purely data-driven approach is unable to exploit prior information to reduce costly explorations of designs; however, this is advantageous when prior information is incorrect. Finally, the learned adaptive algorithm's ability to exploit prior knowledge to obtain performance that is comparable to that of the Bayesian approach when prior information is correct, as well as its use of data-driven inference to adapt online to the encountered population, yields a robust performance across different settings.

3.4. Nonstationary Bidder Behavior

We now explore the robustness of our mechanism when bidder behavior is nonstationary and changes

over time. Shifts in bidder behavior may occur because changes in the preferences of bidders, such as decreasing or increasing valuations over time. Bidders' strategies may also change as they learn from past experiences. In principle, the adaptive mechanism we propose can be used in this setting to adjust to the changes exhibited by the encountered population over time. Thus we first evaluate the performance of the learned adaptive algorithm as compared to the best fixed design in this setting. In addition, we propose several modifications that aim to improve performance specifically under changing bidder behavior. We describe each modification below and evaluate the revenues obtained by the adaptive mechanism with each.

Our primary modification draws from the notion that older auction outcomes provide less-relevant information on the current profitability of different design if the population changes over time. It may be therefore beneficial to weight recent experiences more heavily than older and less-relevant ones. We implement this modification using exponential smoothing—an approach used widely in many time series forecasting models, such as autoregressive integrated moving average (ARIMA) (Mills 1990), and that assigns weights to time-stamped observations so as to decrease the effect of observations in the more-distant past on the model's predictions. Specifically, recall that the locally weighted quadratic regression (LWQR) we use for information sharing assigns a weight to each past experience, given by e^{d/w^2} , which is proportional to the corresponding design's distance d from the design for which profit is predicted. To discount older outcomes, the weight assigned to an auction outcome which occurred at time l is given by $(e^{d/w^2} \gamma^{t-l})$, where $0 < \gamma < 1$ is the decay factor, and t is the current auction number. Because observations with lower weights will have less impact on LWQR's prediction, these weights cause older experiences to have a lower effect on the auctioneer's estimations of expected revenue and subsequently also on the adaptive algorithm's choice likelihood of design i , $D^t(i)$.

Recall that metalearning produces initial training data T populated with the initial expected revenues given by avg_i for each design i . In the empirical evaluations that follow, we consider two variations for decaying past experiences when metalearning is used. In the first version, actual auction outcomes are discounted over time; however, the prior experiences generated in simulation during the metalearning phase are not discounted. For designs with little recent actual experience, the predicted revenue would then be based largely on the revenues estimated through simulation (because these outcomes will have larger weights). We also evaluate the performance of our mechanism when the prior experiences generated during the metalearning phase are discounted in the same manner as actual experiences with bidders. Finally, in addition to discounting less-recent auction outcomes, if the seller has information on the process by which the population changes, the seller can also use this prior information during metalearning to improve its selection of adaptation parameters.

In the empirical evaluation we introduce changes to the bidders' behaviors over time by letting each of the parameters defining bidder behavior—the mean and variance of the valuations, and the mean and variance of level of bidder loss aversion—vary according to a random walk: after each auction, each parameter is either increased or decreased by 1% of the total range for that parameter. For the decay function, we used a conservative decay rate of $\gamma = 0.99$. In principle, because the seller has no historical data to fit the best decay rate, the stochastic optimization during metalearning could also search for an adequate decay rate.

Table 1 presents the average revenues and the number of auctions required for the adaptive algorithm

Table 1 Average Revenue When Bidders' Behaviors Change Over Time

Auction design method	Average revenue	No. of auctions to exceed optimal fixed design
Best fixed design	0.367	
Original learned adaptive design*	0.379	51
Decay of past experiences		
Learned adaptive: Decay of actual experiences*	0.383	50
Learned adaptive: Decay of actual and simulated experiences, metalearning with changing bidders*	0.389	50
Learned adaptive: Decay of only actual experiences, metalearning with changing bidders*	0.394	49

*The difference from the best fixed mechanism is statistically significant according to a paired t -test ($p < 0.01$).

to exceed the revenue of fixed designs. We present the performances obtained by the original learned adaptive mechanism with metalearning, which does not incorporate the adaptation proposed, the modified adaptive mechanisms, and the best fixed mechanism. As shown in Table 1, the original adaptive mechanism with the learned parameters, as well as all the modified adaptive mechanisms, yield higher revenues than the best fixed design. All the adaptive algorithms require between 49 and 51 auctions to exceed the cumulative revenue of the optimal fixed design.

Compared with the original mechanism with no modification, decaying the effect of past outcomes and incorporating information on the process by which bidders' behaviors change via metalearning are each effective, further improving the auctioneer's profitability. Although decaying actual experiences is advantageous, there is no benefit from decaying the effect of experiences obtained in the metalearning phase during simulations. As we note above, designs that have become less profitable over time because of the changing population, and for which the auctioneer has little actual recent experiences, are more likely to be explored again if they are found to be beneficial in expectation in the metalearning phase, and these simulated experiences are not discounted. Encouraging reexploration of designs that have not been profitable in less-recent auctions is clearly beneficial.

3.5. Learning Bidders

Thus far, we let bidders behave according to a known equilibrium. In this section, we let bidder strategies evolve through repeated participation in auctions. As before, we consider an English auction in which bidders have independent, private valuations. However, bidders may arrive at any time during an auction, and both the bidders and the seller are time sensitive. Buyers incur a fixed participation cost (actual or perceived) per unit time in which they remain active; the seller also incurs a fixed cost per unit time, motivating the seller to set a low BIN price to shorten the auctions. The seller revises the auction's BIN price, which can affect the outcome in the new setting we explore here. A BIN price that is set too low can reduce the seller's revenue from the auction, whereas setting it correctly may induce impatient bidders to accept a BIN price higher than the price they might pay if the auction ran to completion. Finally, a new auction begins as soon as the previous one has completed, and the seller has 100 items to sell.

To enable representation of arbitrary strategies, we represent strategies via a neural network—for any given design and conditions that the bidder may encounter, the network takes as input information available to the bidder and outputs the bid to place

and, unless the bidder is the current winner, whether the bidder will stop participating. (Note that because bidders are time sensitive, it may be beneficial for a bidder to leave if she is not likely to win the auction. The inputs to the network represent the current winning price, whether the bidder is currently winning, the number of periods left, whether this is the last period, the total number of periods in which the bidder has participated, the BIN price, the winning price of the previous auction, the average and standard deviation of the last five winning prices, the bidder's valuation and cost per period, and the auction number currently running during the episode, between 1 and 100.) Thus, the strategy captured by a given network consists of different actions for different conditions the bidder may encounter within an auction and throughout an episode. To evolve a network we use the process of competitive coevolution, where strategies compete against each other as they evolve. Coevolution has been used successfully to evolve bidder strategies when closed-form equilibrium is infeasible (Cliff 2001, Byde 2003). We evolve networks using NeuroEvolution of Augmenting Topologies (Stanley and Miikkulainen 2002), a genetic algorithm designed to evolve both the network topology and its weights. In the interest of space, complete details of this setting including the populations' parameters are provided in §C of the Online Supplement to allow reproduction of our results.

In this study we consider bidders who may participate in several auctions during an episode. This setting offers incentives to evolve behaviors that take the seller's adaptation over the course of an episode into account. For example, bidders may find that refusing to accept a high BIN price in early auctions induces the adaptive algorithm to offer lower BIN prices in later auctions. In practice, this scenario emerges when bidders remain active for several auctions until winning or dropping off, or when bidders experience recurring demand. We define the utility derived by a bidder in a given auction by $U = (w \cdot (v - p) - c \cdot L)$, where v is the bidder's valuation, p is the winning bid, w is one if the bidder won the auction or zero otherwise, c is the cost per time unit incurred by the bidder, and L is the duration in which the bidder was active. The fitness of each network depends on the average utility it yields over an episode in which it encounters the adaptive algorithm, and when the network is used by bidders drawn at random from the population. Because coevolution requires a substantial number of experiences to produce effective actions in response to the seller, more than are available during an episode, we first let networks evolve while encountering the adaptive algorithm over 500 generations so as to evolve responses to the seller's

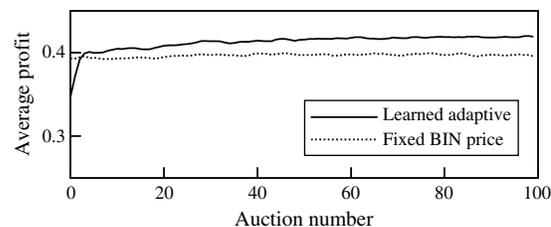


Figure 6 Results for Bidder Strategies Evolved via Competitive Coevolution

adaptation. We then evaluate the adaptive algorithm's performance over a single episode as before.

Our use of competitive coevolution aims to evolve bidders' responses specifically to the proposed adaptive algorithm's form of adaptation. Some recent prior work has offered insights into bidders' behaviors and other aspects of sequential auctions (Zeithammer 2006, Juda and Parkes 2006, Ginsburgh and van Ours 2007) or examined stylized sequential auctions to derive equilibrium bidding strategies analytically (Jeitschko 1998). Interestingly, although Zeithammer (2006) finds that data on eBay bidders suggest bidders are taking future auctions into account, Juda and Parkes (2006) show that eBay bidders tend to bid inefficiently in sequential auctions. This irrationality may manifest itself differently in different bidder populations and auction settings; however, our adaptive algorithm is particularly designed to accommodate such uncertainty.

Figure 6 shows the average profit per auction for the best fixed design and the adaptive algorithm with learned parameters over 10,000 bidder populations. As shown, the adaptive algorithm identifies auction designs that are more profitable than the alternative best fixed mechanism. The adaptive algorithm with learned parameters exceeds the total profit obtained by the best fixed design after 14 auctions and obtains an overall average profit over the entire episode of 41.35, as compared to 39.67 with the best fixed design. The differences in performance are statistically significant ($p < 0.01$). Importantly, note that our empirical results consider an asymmetric scenario, where bidders have more opportunities to evolve to the adaptive algorithm's form of adaptation than the number of auctions in a single episode used by the adaptive algorithm to adapt. Overall, our results suggest that the adaptive mechanism is not likely to be undermined by individual bidder learning.

3.6. The Generality and Sustainability of the Adaptive Algorithm

The adaptive mechanism we propose is generic and can be applied in any setting in which the objective for improvement can be observed by the seller after each auction. The empirical evaluations presented so

far have focused on the robustness of our approach as compared to alternatives under a variety of settings. In this section we present studies that explore the generality of the approach—its ability to adapt effectively to a variety of complex bidder populations as well as to remain advantageous under auction settings of particular practical interest.

To facilitate the analysis of our approach, in the previous sections auction mechanisms were characterized by a single design parameter (reserve price or BIN price). However, our approach can accommodate adaptation of multiple parameters simultaneously. In addition, the adaptive algorithm can adapt to an arbitrarily complex and heterogeneous population of bidders (Bapna et al. 2004). We explore the adaptive mechanism’s performance when it adapts to a heterogeneous population consisting of three distinct bidding strategies described separately in prior work and increase the complexity of the design space by simultaneously revising two parameters—the auctioned item’s start price and the BIN price. Specifically, the population of bidders consists of exponential discounters whose utility is given by $U(v, b, \alpha) = (1/(1 + \alpha t))(v - b)$, $\alpha > 0$, where v is the valuation, b refers to the bid, and α captures the level of time sensitivity; hyperbolic discounters (Ainslie 1992) who have a utility given by $U(v, b, \delta) = e^{-\delta t}(v - b)$, where t is the current time during the auction and δ is a discount factor; and snipers who wait to submit their bids until very close to the end of the auction (Roth and Ockenfels 2002). For an auction that ends at time T , a sniper’s probability of joining the auction at time t is given by $0.7/e^{\beta(T-t)}$ such that the probability increases closer to the end of the auction. At each auction 40% of bidders are exponential discounters, 40% are hyperbolic discounters, and the remaining 20% are snipers. We assume that the exponential and hyperbolic discounters have a fixed probability (0.7) of joining the auction.

For each bidder we draw a valuation uniformly at random from the range $[0, 1.2]$, and for each hyperbolic discounter we draw α from a Weibull distribution given by $W(3, 0.1, 0.04)$. We used a fixed hyperbolic discount factor $\delta = 0.94$ and assigned $\beta = 1.2$ for all snipers. When a buyer joins the auction, we draw a bid uniformly at random from the range $(0, v)$, where v is the bidder’s valuation for the item. If the bid is higher than the current winning bid, the bidder becomes the leading bidder. If the bid is equal to or exceeds the BIN price, we assume the bidder executes the BIN option, and the item is sold to the bidder at the BIN price. We also assume the seller does not have any prior information. Thus, we do not use metalearning and assume initially that it is uniformly desirable to explore all start and BIN prices.

We compare the adaptive algorithm’s capacity to identify effective designs with the performance of a fixed auction mechanism. For the fixed design we use parameter values derived in prior work for each parameter separately, and assume exponential discounters. This is because no analytical results are available to suggest the best multiple parameters for a mixed population of bidders. We set the start price at zero (Ariely and Simonson 2003, Dholakia and Simonson 2005), and because the BIN price is a ceiling price, we use a high BIN price (Budish and Takeyama 2001). Last, we set a low reserve price of zero, which was suggested in prior work to be advantageous (McAfee 1993, Peters and Severinov 1997, Peters 1997). We apply our adaptive algorithm to simultaneously adapt the auction’s start price and BIN price.

Figure 7(a) shows the average revenue per auction in the episode over 100 independent experiments, produced by the fixed auction design and the adaptive algorithm. As shown, the advantages of the adaptive mechanism generalize well to the complex bidder population and design space we explore here. The adaptive auction yields statistically significant and considerably higher revenue than the fixed auction mechanism. As we note in §1, because the adaptive algorithm does not require explicit representation

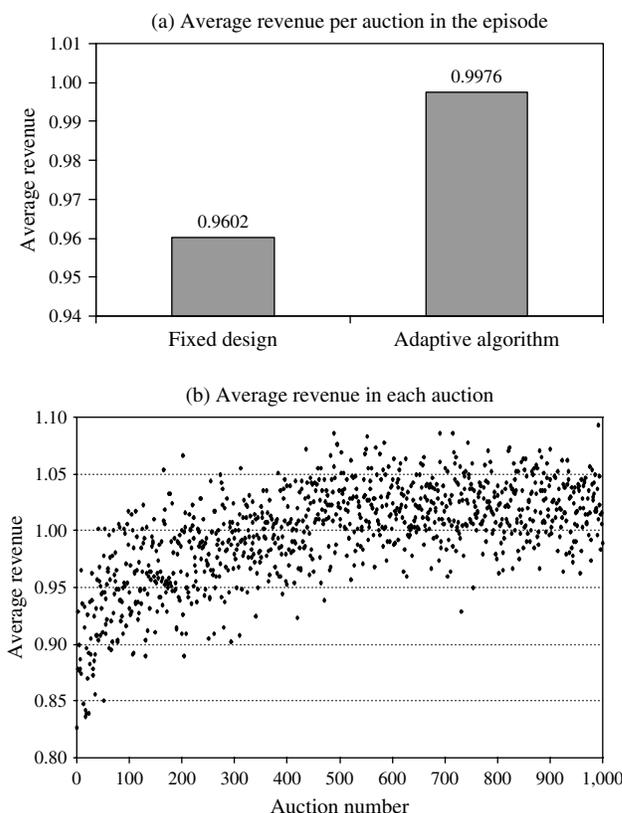


Figure 7 Revenue with Heterogeneous Population and Multiple-Parameter Mechanism Design

of bidders' strategies, it is particularly advantageous in complex settings such as those explored here. As shown in Figure 7(b), the learning rate of the adaptive algorithm in this setting is slower than that shown in Figure 2. This is because the algorithm now adjusts more parameters simultaneously; the larger number of stochastic factors affecting the auction outcome also contributes to the lower learning rate.

A practical concern when a large number of items can be auctioned sequentially is the possible need to accommodate sellers auctioning multiple items in each auction so as to accelerate the selling of items. In this setting, the data-driven mapping, used to construct the distribution D from which auction designs are drawn, will capture a mapping between a design and the revenue for the corresponding multiple-item auction. As an example, for a seller auctioning five items in each auction, Figure 8(a) shows the average revenue obtained by the adaptive algorithm as compared to a fixed design; Figure 8(b) presents the average revenue throughout the episode. Note that in this setting only 200 auctions are required to auction 1,000 items. As shown, in this setting as well the adaptive algorithm effectively adapts the auction design, exceeding the revenue obtained with a fixed design.

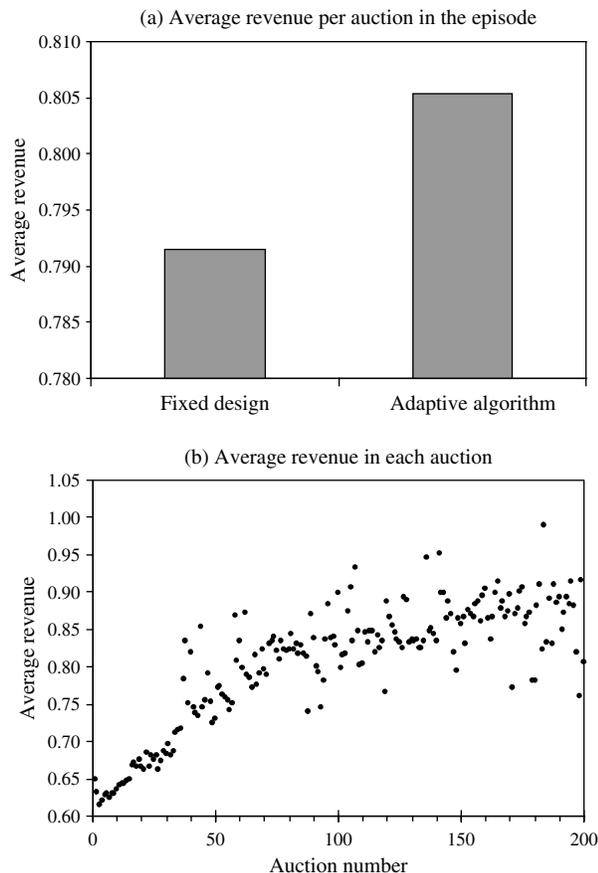


Figure 8 Adaptive Algorithm's Performance in a Multi-Item Auction Setting

3.6.1. Bidders' Surplus and Participation. In §3.4 we examined the adaptive algorithm's robustness when bidder behavior changes over time. One special case of interest is when each bidder's likelihood to participate in the auction decreases in response to a recent decrease in bidder surplus. Although the seller cannot directly observe consumer surplus and thus cannot employ the data-driven approach to benefit consumer surplus directly, this setting is interesting because it entails the need for sufficient bidder surplus to sustain bidder participation and avoid revenue loss. Recall that the adaptive algorithm's choice of designs is determined by a data-driven mapping from a design onto its expected benefits. After each auction, the adaptive algorithm revises this mapping and subsequently the likelihood of selecting these designs in the future. Thus, designs that produce lower benefits over time, such as those that result in lower bidder participation, are also less likely to continue to be selected by the adaptive algorithm.

We examine the performance of the adaptive algorithm in this setting for the heterogeneous population of hyperbolic and exponential discounters, and we let the likelihood to participate in any given auction decrease if the average bidder surplus over the most recent two auctions has decreased. In our empirical evaluations we subtract the decrease in surplus from the likelihood to participate in the auction.

Figure 9(a) presents the adaptive algorithm's increasing revenue over the course of the auction. As the figure panel shows, the adaptive algorithm effectively selects designs that ultimately yield improvement in the seller's revenue. As shown in Figure 9(b), we also find that the adaptive algorithm yields higher average revenue over an episode than the revenue obtained by the fixed design. Overall, the adaptive algorithm is able to maintain sufficient level of participation to remain profitable over time.

Also of interest in this investigation is the level of consumer surplus obtained with the adaptive algorithm when bidder surplus influences participation as compared to the surplus when participation is unaffected by bidder surplus. We find that when a bidder's likelihood to participate changes in response to bidder surplus, the adaptive algorithm's choice of designs yields an average bidder surplus of 0.15, higher than the average surplus of 0.12 obtained when bidder surplus does not affect participation. This result suggests that when bidder participation depends on the surplus bidders are able to derive, the adaptive algorithm selects designs that yield higher bidder surplus to sustain sufficient participation, as compared to when bidders are unaffected by their surplus. As we discuss above, these effective choices of design are enabled by the continuous revision

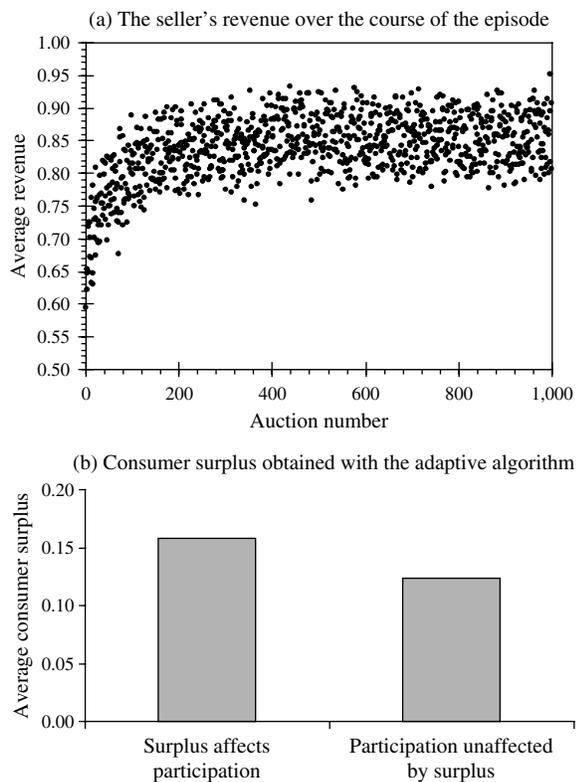


Figure 9 Revenue and Consumer Surplus When Low Surplus Undermines Future Bidder Participation

of the adaptive algorithm's mapping used to estimate the profitability of designs and the likelihood of selecting them.

4. Conclusions and Future Work

The growing popularity of electronic auctions presents new opportunities to improve the mechanisms sellers use to govern the exchange. In this paper, we present a new approach to auction design that capitalizes on key properties of electronic markets, such as the large volume of auctions and the possibility to dynamically alter the mechanism's rules of exchange. We present an approach to adapt the auction design online to an arbitrary bidder population in response to real auction outcomes, as well as a principled approach to incorporate prior knowledge to enhance the adaptation to the encountered bidders.

Central to the seller's adaptation is an effective balance between exploration of possible designs and exploitation—the selection of designs the seller estimates to be particularly profitable at any given time. Because information acquisition through exploration is costly but necessary we study how the seller's estimation of the profitability of designs can be accelerated to reduce the exploration of suboptimal designs. We propose two approaches to reduce costly explorations, information sharing and metalearning, which

are shown to be effective under a variety of conditions. The incorporation of prior knowledge into data-driven auction design is a particularly important venue for future research. We discuss how the trade-off between learning bias and variance affect the performance of metalearning and of other alternatives. In future work, it would be beneficial to further improve our understanding of effective mechanisms for incorporating prior knowledge into data-driven adaptation, perhaps by alternating among different mechanisms during the online adaptation.

Extensive empirical evaluations demonstrate that our adaptive mechanism effectively adapts the auction design to a variety of encountered bidder populations and consistently yields higher average revenues per episode when compared to the best fixed mechanism, which does not adapt. When prior information does not reflect current bidder behavior and when bidders' strategies vary over time, our evaluations explore the bias versus variance trade-off of alternative approaches and suggest the balance offered by adaptive auctions with metalearning yields significantly more robust performance. Also, we show that the adaptive algorithm generalizes well to settings in which the seller encounters heterogeneous populations and where multiple auction design parameters are revised simultaneously. For the keyword auction problem discussed in §1, the adaptive algorithm can be used in place of a fixed design to revise multiple design parameters, such as the reserve price and number of ads to display, and can identify different advantageous designs for different keywords. Furthermore, as demonstrated by our empirical results, our adaptive mechanism is likely to be particularly beneficial for keywords pertaining to new products or services, where the populations of bidders may not be well understood prior to the auction; the mechanism is also likely to be particularly useful for keywords that, because of changes in the competitive landscape, may induce different behaviors from advertisers over time. Finally, we demonstrate that the mechanism we propose is sustainable even when bidders' strategies evolve over time to respond specifically to the adaptive algorithm's form of adaptation and when bidders' participation is affected by their surplus.

Existing electronic auction platforms are relatively simple. They allow sellers to set certain auction parameters, but they do not provide intelligent mechanisms to advise sellers how to effectively acquire information through interactions with bidders to adapt to the encountered population. The mechanism we present here can be implemented by individual sellers as well, independently of the platform they use; however, our results suggest that auction platforms such as eBay should consider making such technologies available to sellers. The demonstrated

benefits to sellers are likely to encourage participation and contribute to the platform's overall popularity and profitability.

Acknowledgments

The authors thank two anonymous referees and the associate editor for their thoughtful suggestions. The authors are responsible for any errors. Some of this work has taken place in the Learning Agents Research Group (LARG) at the Artificial Intelligence Laboratory, the University of Texas at Austin. LARG research is supported in part by grants from the National Science Foundation (CNS-0615104), DARPA (FA8750-05-2-0283 and FA8650-08-C-7812), the Federal Highway Administration (DTFH61-07-H-00030), and General Motors.

References

- Ainslie, G. 1992. *Picoeconomics: The Strategic Interaction of Successive Motivational States Within the Person*. Cambridge University Press, New York.
- Ariely, D., I. Simonson. 2003. Buying, bidding, playing, or competing? Value assessment and decision dynamics in online auctions. *J. Consumer Psych.* **13**(1–2) 113–123.
- Atkeson, C. G., A. W. Moore, S. Schaal. 1997. Locally weighted learning. *Artificial Intelligence Rev.* **11**(15) 11–73.
- Bapna, R., P. Goes, A. Gupta. 2003. Analysis and design of business-to-consumer online auctions. *Management Sci.* **49**(1) 85–101.
- Bapna, R., W. Jank, G. Shmueli. 2008a. Price formation and its dynamics in online auctions. *Decision Support Systems* **44**(3) 641–656.
- Bapna, R., P. Goes, A. Gupta, Y. Jin. 2004. User heterogeneity and its impact on electronic auction market design: An empirical exploration. *MIS Quart.* **28** (1) 21–43.
- Bapna, R., P. Goes, A. Gupta, G. Karuga. 2008b. Predicting bidders' willingness to pay in online multiunit ascending auctions: Analytical and empirical insights. *INFORMS J. Comput.* **20**(3) 345–355.
- Blum, A., J. D. Hartline. 2005. Near-optimal online auctions. *SODA '05: Proc. 16th Annual ACM-SIAM Sympos. Discrete Algorithms*, Society for Industrial and Applied Mathematics, Philadelphia, 1156–1163.
- Blum, A., V. Kumar, A. Rudra, F. Wu. 2003. Online learning in online auctions. *SODA '03: Proc. 14th Annual ACM-SIAM Sympos. Discrete Algorithms*, Society for Industrial and Applied Mathematics, Philadelphia, 137–146.
- Budish, E. B., L. N. Takeyama. 2001. Buy prices in online auctions: Irrationality on the Internet? *Econom. Lett.* **72**(3) 325–333.
- Byde, A. 2003. Applying evolutionary game theory to auction mechanism design. *Proc. 4th ACM Conf. Electronic Commerce*, ACM Press, New York, 192–193.
- Cliff, D. 2001. Evolution of market mechanism through a continuous space of auction types. Technical Report HPL-2001-326. HP Labs, Bristol, UK.
- Cramton, P. C. 1997. The FCC spectrum auctions: An early assessment. *J. Econom. Management Strategy* **6**(3) 431–495.
- Dodonova, A., Y. Khoroshilov. 2005. Optimal auction design when bidders are loss averse. Working paper, University of Ottawa, Ottawa.
- Dholakia, U. M., I. Simonson. 2005. The effect of explicit reference points on consumer choice and online bidding behavior. *Marketing Sci.* **24**(2) 206–217.
- Friedman, J. H. 1997. On bias, variance, 0/1-loss, and the curse-of-dimensionality. *Knowledge Discovery Data Mining* **1**(1) 55–77.
- Ginsburgh, V., J. C. van Ours. 2007. On organizing a sequential auction: Results from a natural experiment by Christie's. *Oxford Econom. Papers* **59**(1) 1–15.
- Jansen, B. J., T. Mullen. 2008. Sponsored search: An overview of the concept, history, and technology. *Internat. J. Electronic Bus.* **6**(2) 114–131.
- Jeitschko, T. 1998. Learning in sequential auctions. *Southern Econom. J.* **65**(1) 98–112.
- Juda, A. I., D. C. Parkes. 2006. The sequential auction problem on eBay: An empirical analysis and a solution. *Proc. 7th ACM Conf. Electronic Commerce*, ACM, New York, 180–189.
- McAfee, R. P. 1993. Mechanism design by competing sellers. *Econometrica* **61**(6) 1281–1312.
- McAfee, R. P., J. McMillan. 1987. Auctions and bidding. *J. Econom. Literature* **25**(2) 699–738.
- Mills, T. C. 1990. *Time Series Techniques for Economists*. Cambridge University Press, Cambridge, UK.
- Myerson, R. B. 1981. Optimal auction design. *Math. Oper. Res.* **6**(1) 58–73.
- Pardoe, D., P. Stone, M. Saar-Tsechansky, K. Tomak. 2006. Adaptive mechanism design: A metalearning approach. *Proc. Eighth Internat. Conf. Electronic Commerce*, ACM, New York, 92–102.
- Parkes, D. C. 2001. Iterative combinatorial auctions: Achieving economic and computational efficiency. Ph.D. thesis, Department of Computer and Information Science, University of Pennsylvania, Philadelphia.
- Peters, M. 1997. A competitive distribution of auctions. *Rev. Econom. Stud.* **64**(1) 97–123.
- Peters, M., S. Severinov. 1997. Competition among sellers who offer auction instead of prices. *J. Econom. Theory* **75**(1) 141–179.
- Robbins, H. 1952. Some aspects of the sequential design of experiments. *Bull. Amer. Math. Soc.* **58**(5) 527–535.
- Rogers, A., E. David, J. Schiff, S. Kraus, N. R. Jennings. 2005. Learning environmental parameters for the design of optimal English auctions with discrete bid levels. *AAMAS 2005 Workshop Agent Mediated Electronic Commerce VII*, IFAAMAS, Utrecht, The Netherlands.
- Roth, A. S., A. Ockenfels. 2002. Last-minute bidding and the rules for ending second-price auctions: Evidence from eBay and Amazon auctions on the Internet. *Amer. Econom. Rev.* **92**(4) 1093–1103.
- Rothkopf, M. H., R. M. Harstad. 1994. Modeling competitive bidding: A critical essay. *Management Sci.* **40**(3) 364–384.
- Saar-Tsechansky, M., F. Provost. 2004. Active sampling for class probability estimation and ranking. *Machine Learn.* **54**(2) 153–178.
- Shmueli, G., G. Jank, A. Aris, C. Plaisant, B. Shneiderman. 2006. Exploring auction databases through interactive visualization. *Decision Support Systems* **42**(3) 1521–1538.
- Spall, J. C. 1998. An overview of the simultaneous perturbation method for efficient optimization. *Johns Hopkins APL Technical Digest* **19**(4) 482–492.
- Stanley, K. O., R. Miikkulainen. 2002. Evolving neural networks through augmenting topologies. *Evolutionary Comput.* **10**(2) 99–127.
- Sutton, R. S., A. G. Barto. 1998. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA.
- Weber, R. J. 1997. Making more from less: Strategic demand reduction in the FCC spectrum auctions. *J. Econom. Management Strategy* **6**(3) 529–548.
- Yahoo! Search Marketing Blog. 2008. Reserve prices. Posted February 26; accessed September 10, 2009. <http://www.ysmblog.com/blog/2008/02/26/minimum-bids>.
- Zeithammer, R. 2006. Forward-looking bidding in online auctions. *J. Marketing Res.* **43**(3) 462–476.