

Designing Intelligent Software Agents for Auctions with Limited Information Feedback

Gediminas Adomavicius, Alok Gupta

Carlson School of Management, University of Minnesota, Minneapolis, Minnesota 55455
{gedas@umn.edu, alok@umn.edu}

Dmitry Zhdanov

School of Business, University of Connecticut, Storrs, Connecticut 06269,
dmitry.zhdanov@business.uconn.edu

This paper presents analytical, computational, and empirical analyses of strategies for intelligent bid formulations in online auctions. We present results related to a weighted-average ascending price auction mechanism that is designed to provide opaque feedback information to bidders and presents a challenge in formulating appropriate bids. Using limited information provided by the mechanism, we design strategies for software agents to make bids intelligently. In particular, we derive analytical results for the important characteristics of the auction, which allow estimation of the key parameters; we then use these theoretical results to design several bidding strategies. We demonstrate the validity of designed strategies using a discrete event simulation model that resembles the mechanisms used in treasury bills auctions, business-to-consumer (B2C) auctions, and auctions for environmental emission allowances. In addition, using the data generated by the simulation model, we show that intelligent strategies can provide a high probability of winning an auction without significant loss in surplus.

Key words: online auctions; intelligent agents; software agents; limited information feedback; bidding strategies; discrete event simulation; heuristics; parameter estimation

History: Sanjeev Dewan, Senior Editor; H. Raghav Rao, Associate Editor. This paper was received on September 2, 2004, and was with the authors 1 year and 3 months for 5 revisions. Published online in *Articles in Advance*.

1. Introduction and Motivation

A variety of online auction sites such as eBay (www.ebay.com) and uBid (www.ubid.com) have become an integral part of mercantile processes on the Internet. The advantages of using auctions for selling commodities are different for sellers and buyers. While sellers try to reach a wider geographical market in order to obtain a potentially better price, buyers hope to purchase the item of interest at a bargain. The auction mechanisms used by different auctioneers show tremendous variability. For example, while the majority of auctions on eBay sell a single unit of an item with 4–6 bidders bidding in most auctions (Kauffman and Wood 2006), uBid often sells multiple units of an item with participant count in hundreds (Bapna et al. 2004). Compared to their traditional (i.e., offline) counterparts, online auctions run significantly longer,

with most auctions being conducted over a period of 24 hours to a week.

Because of a variety of interesting issues, such as auction design mechanisms, incentives, participant strategies, and the potential for monitoring and using information via software agents, online auctions are attracting researchers from a wide array of disciplines. In particular, roles and applications of *intelligent agents* in auctions are growing for a number of reasons. For example, Chang (2002) reports that, in a double-auction experiment at IBM Watson Research Center, bidding agents programmed with simple strategies were much faster than human players and were 10% more profitable. However, there is a potential to use bidding agents in more complex environments that are currently being developed. In this paper, we explore how theoretical properties of a

mechanism can be used to design software agents that provide strategic advantage to their users.

While many conventional auction settings have been studied in detail in economic literature (e.g., see Milgrom 1989), various modifications made by online auctioneers often change the incentive structures and create opportunities for bidding strategies and evaluation metrics that are not analyzed in conventional treatments (Bapna et al. 2003a). Traditional game-theoretic solutions typically used to analyze auction mechanisms generally require restrictive assumptions to make the analysis tractable and rarely consider the real-world operating environment (Rothkopf and Park 2001). As Klemperer (2002) points out, extensive theoretical literature on auctions is not fully applicable to the practical auction design.

Therefore, computer science researchers develop algorithms and heuristics that use ad hoc approaches for resource allocation in reasonable time instead of trying to formally solve general cases of complex real-world implementations of auction mechanisms. For example, the ongoing Trading Agent Competition (Stone and Greenwald 2005) simulates auctions in which multiple software agents try to maximize their utilities in the travel industry and supply chain management settings. However, these agents usually try ad hoc strategies that are rarely designed based on economic principles (Ketter et al. 2006).

Information systems (IS) researchers have also made significant contributions to online auction research by examining issues such as non-game-theoretic characterization of online auctions (Bapna et al. 2003a), empirical investigation of factors affecting the results of online auctions (Bapna et al. 2002, Kauffman and Wood 2006), discovering and monitoring bidder behaviors (Bapna et al. 2004), and creation of test beds for exploring auction design and bidder strategies that cannot be explored analytically (Bapna et al. 2003b, Mehta and Bhattacharya 2006). One of the key research areas explored by computer science and IS researchers is the role of *software agents* in online auctions (Adomavicius and Gupta 2005, Mehta and Bhattacharya 2006). Because the auction durations of 24 hours or more make it impractical for human bidders to monitor such auctions continuously, software agents facilitate monitoring of these auctions

and place bids on behalf of the user in predefined ways.

Online auctioneers use many different auction mechanisms that range from traditional English and Dutch auctions to multiunit discriminatory and uniform-price auctions. Even with seemingly similar auction formats, online auctioneers use a variety of strategies to set controllable auction parameters resulting in different *rules of engagement*. Typically, the control of an auction parameter involves two decisions: (i) choosing a particular value for the parameter, such as the reservation price, and (ii) determining whether to disclose the chosen value to the bidders. While Milgrom and Weber (1982) have shown that revealing all information raises the expected revenue in most auctions where bidders have independent private values, there are multiple arguments about the need to conceal some of this information from bidders in other types of auctions. For example, Tu (2005) indicates that in iterative sequential auctions, the information release about bid history should be selective. Specifically, they show that the only type of information disclosure that improves seller's revenue is a report of winning bids and only for the first-price and Dutch auctions.

The richness-of-information factors in online auctions combined with the sellers' incentives to hide some of the information make intelligent bidding a challenge. Klusch (2000) specifies that information agents should carry out at least one of the following functions: information acquisition and management, information synthesis and presentation, and intelligent user assistance. Furthermore, Greenwald et al. (2003) argue that software agents have at least three advantages over humans: faster operation, no distraction, and flawless logic. Bicharra Garcia et al. (2001) point out that agents are present for the entire duration of an auction and can exploit the information gathered during an auction. Overall, intelligent agents are important for both the theory and practice of IS because they facilitate coping with the variety and volume of data in constantly changing environments (March et al. 2000).

However, it is often not clear how the flawless logic can be generated from the perspective of benefiting the user, and, more generally, how the presence of agents can benefit the bidder. One of the ways to

make the agents directly serve the interests of bidders is to assist bidders in constructing and placing intelligent bids. Intelligent bids can be broadly understood as the ones which align with bidders' goals and strategies, e.g., maximizing surplus when bidding on a commonly available item, or guaranteeing a win when bidding on a rare item. However, most current software agents primarily play a role of passive participators. According to the categorization of online bidding agents by Karuga and Maganti (2004), existing bidding agents do not have any intelligence built into them. Most bidding agents, such as *Bid Butler* (<http://www.ubid.com/help/topic10.asp>), simply ask the bidder for a maximum bid amount and then keep placing incremental bids up to this prespecified amount when a bidder is no longer winning. Very little attention has been paid to the creation of intelligent agents for B2C auctions, where a bidder participates only occasionally and does not usually have an opportunity to learn from multiperiod data. Based on the literature cited above, we can identify the following three core capabilities that intelligent bidding agents should have:

1. *Inference of auction parameters.* As discussed earlier, there are multiple decisions about the auction parameters that the auctioneer has to make and there are incentives to hide some of these parameters from the bidder. An intelligent agent can infer the hidden parameters—e.g., the number of auction participants—to use a more informed strategy in subsequent bidding on behalf of its user.

2. *Estimation of the current auction state.* As Wurman et al. (2001) argue, there are three major functions of an auction: accept bids, clear the market, and provide intermediate information. The first two functions are mandatory, while the last is common, but not required. Intermediate information typically conveys the current state of the auction to the participating bidders. The current state of an auction can sometimes also be used to judge the status of a bidder's own bid. For example, in an English auction, the intermediate information is the highest bid at a given point in time, which tells every bidder whether they are currently winning. However, intermediate information may not allow the computation of exact allocation easily. For example, in combinatorial auctions (see Kwasnica et al. 2005 and Adomavicius and

Gupta 2005) various approximations often need to be used (e.g., Sandholm 2002). An intelligent agent can use complex mathematical algorithms and approximations to compute useful information to its users.

3. *Prediction of the future auction state.* Armed with the knowledge of auction parameters and its current state, an intelligent agent must use that knowledge to benefit the user, e.g., by maximizing the bidder's expected surplus. This decision problem may be partitioned into three subproblems (Parkes 2000): metadeliberation, valuation, and bidding. Metadeliberation is a decision about when to place a bid; valuation refers to ascertaining the value of an item when it is not certain, and bidding is the formulation of an exact bid based on the previous two decisions. It is clear that all three decisions depend on the information derived using the parameter inference and the capability to estimate the current auction state. For example, if there are three active bidders for three available items, a bidder may not want to update her bid until another bidder arrives.

In this paper, we demonstrate how theoretical properties of a mechanism can be used to develop agents that have all three desired capabilities discussed above. Therefore, we chose an auction mechanism that has a high degree of information opaqueness, and where the computational requirements for an agent to develop the three core capabilities are nontrivial. As a result, we concentrate on developing an intelligent agent for the *weighted average* auction, which is inspired in part by the Spanish auction (Alvarez et al. 2003). We first provide a theoretical analysis of a general form of the weighted-average auction. Then, we derive intelligent strategies for bidding agents to help make bidding decisions. We demonstrate the efficacy of our approach using a comprehensive simulation-based analysis. This paper will demonstrate and argue that economic principles and analysis of incentives can be used to derive strategies for the effective design of intelligent agents (Ba et al. 2001). Because the optimal strategy of bidders is to maximize their surplus, which is positive only if a bidder wins an auction, the primary strategy of an agent should be to win at a reasonable price. Hence, in our context, the economic incentives of an agent are to develop a strategy that provides a high likelihood

of winning at a reasonable price below the agent's private valuation.

The rest of this paper is organized as follows. In the next section, we present our analytical model and results characterizing the structure of auctions that involve reporting of the weighted-bid average. In §3, we develop strategies for an intelligent agent. In §4, we present the simulation model to test and validate the agent bidding strategies. Finally, the conclusions and directions for future research are presented in §5.

2. Analytical Model

In this section, we provide the details of the weighted-average auction mechanism. We analyze the auction process and derive theoretical results about its key parameters. We also discuss the links between the analyzed auction design and some of the more familiar auction designs.

2.1. Information Properties of Auction Mechanisms

Practical implementation of optimal auction design is complex, even for auctions that may look simple. For example, Banks et al. (2003) analyze multiple settings that might be suitable for FCC spectrum license auctions. They argue that, in individual private-value settings, progressive English auctions may not maximize revenue if bidders are risk averse; instead, first-price sealed-bid auctions generate higher revenue. In contrast, in common-value settings, iterative English auctions may generate higher revenue. However, these conclusions are based on several assumptions that can be challenged, and the equilibrium for such auctions is very hard to derive. In multiunit auctions, the problem becomes even more complex and analytically intractable (Nautz and Wolfstetter 1997, Bapna et al. 2003a).

Therefore, researchers (e.g., Banks et al. 2003) have been focusing on the complexity of auctions and induced bidding strategies to control the properties of the mechanisms via policies rather than focusing on equilibrium analysis. In the same vein, we focus on deriving strategies based on auction rules and the resulting complexity of several different variants of weighted-average auctions. There are three dimensions of complexity in auctions that define information content and drive bidder strategies. First, an

auctioneer needs to decide whether to use an *iterative* mechanism because there are revenue-based arguments both for and against these types of auctions. Moreover, in iterative mechanisms the duration of the auction also has to be considered—for example, auctions that are limited in duration by time or number of rounds induce sniping (Roth and Ockenfels 2002). Second, *multiunit* (as opposed to single-unit) auctions have an additional set of strategic issues for bidders related to possible partial fulfillment of bids. Therefore, creating an agent to bid on multiple items is more complex than bidding for a single item, potentially requiring more sophisticated strategies. Finally, a decision about the extent of *information revelation* needs to be made because it has an impact on the revenue as well (Tu 2005).

A variety of auction configurations can be constructed by choosing different joint configurations of number of items and iterations—from a single-item single-round Vickery auction to iterative multiunit auctions. However, the final dimension of information transparency and feedback given to the bidders during the auction has not been explored in theory and practice as comprehensively as other dimensions. The majority of commonly used auctions use either full-transparency or no-transparency design. A promising set of auctions with *partial transparency* (e.g., muni auctions for selling municipal bonds, such as at grantstreet.com) have not received much attention in the academic literature, although they are used in practice. In this paper, we analyze a family of partial-transparency auctions that may be used to conceal the bidding information from bidders and influence bidder behavior. These auctions use the bidder-feedback scheme based on a generic form of the weighted average of winning bids, inspired in part by Spanish treasury auctions. As discussed later, while considering higher complexity in information transparency, we do not sacrifice much complexity in the dimensions of duration and number of items available.

The original design of Spanish auctions is used by the Treasury of Spain to sell government securities (Alvarez et al. 2003). A similar design is also used in electricity auctions in the United Kingdom (Fabra et al. 2002) and California (Thomas et al. 2002). Spanish auctions have the following characteristics (e.g.,

Alvarez et al. 2003): divisible goods, allowing for partial fulfillment of bidders' orders; the price paid by bidders is a combination of discriminatory and uniform pricing schemes, where those bidders who win below the weighted average of winning bids pay their own bid amount, while those who win above the weighted average pay that average; the weights used in averaging of winning bids are equal to the bundle size of bids; and auctions are typically sealed-bid with each bidder placing a single bid.

From the modeling perspective, Spanish auctions may be considered multiunit auctions with arbitrarily small unit size. The allocation rule in Spanish auctions provides *some* incentives for bidders to bid their true valuations because a winner either pays below average or an average price below their bid retaining some surplus. However, because the weighted average depends on the winners' bids, the auction is not incentive compatible in a true sense and bidders tend to overbid thinking that they will win at lower prices. The properties of such auctions are not yet well studied, but the mechanism itself seems promising, e.g., Abbink et al. (2006) show that the Spanish format generates higher revenues for the seller in the common-value sealed-bid setting with repeated auctions as compared to an English auction. Alvarez et al. (2003) arrive at similar results using a theoretical model and simulations with two bidders making two bids.

Auctions implementing different forms of the weighted-average mechanism have been used by various trading entities throughout the world. For FCC spectrum auctions, the Auctions Division of the Federal Communications Commission defined bid increment for each license based on a weighted average of the activity on that license in the most recently completed round as compared to the activity in the previous rounds.¹ Chicago Climate Exchange auctions greenhouse gas emission allowances in a mixed-auction scheme, with about 20% of allowances sold using average-price auctions. In these auctions, the current weighted average is defined based on prices from a preceding discriminatory sealed-bid auction, which sells the other 80% of allowances.² Such

mechanisms have also been used by B2C auction sites, e.g., dealspin.com.³

In this paper, we consider a weighted-average auction mechanism in which the bidders are allowed to place *multiple* bids over time. The feedback provided to the bidders is restricted to a *generalized weighted average* of winning bids. In other words, bidders are not given information about other bids; instead, bidders are provided a certain aggregate metric—a weighted average—computed from current winning bids. In addition, each bidder is restricted to a prespecified number of bids. At the end of an auction, bidders are notified about whether their bid was among the winning bids. The formation of a bid in such an environment becomes a challenging task because a bidder does not know the bids that are currently winning. In addition, because of the restriction on the number of bids a bidder can place, it becomes important to make intelligent bids and support bidders in evaluating their own bids. Note that auctions with a variety of different bidding restrictions have been investigated in the research literature (sealed-bid auctions being one of the simpler examples). However, the reporting of a specialized weighted average of winning bids combined with the restriction on the number of bids constitutes a novel, interesting, and plausible way of information masking, particularly if some of the information involved in the computation of the weighted average is not known to bidders. Specifically, our model has the following properties: It is an iterative auction allowing bidders to place multiple bids; it uses a general form of weighted-average computation; bidders' final price is equal to their own bid, which discourages the strategy of placing high bids; and auctions are for multiple units of the same item.

While our main focus is on specific multiunit auction settings where each bidder can bid only on a single item, we also provide a theoretical and empirical exploration of bidder strategies for bidding on multiple items. Next, we describe a business setting in which such an auction may be useful, provide a formal statement of the auction problem, and set up our analytical model.

¹ <http://wireless.fcc.gov/auctions/data/papersAndStudies/SmMethFactSheet.pdf>.

² http://www.chicagoclimateexchange.com/news/auction_intro.html.

³ The site now seems to be defunct pending investigation of complaints about irregular credit-card charges.

2.2. Business Model

Several online sites, such as uBid.com, buy certain consumer items in bulk and resell them to the end-consumer. They buy a certain number of units to receive price discounts from the manufacturer or an intermediary. Then, they use auctions to attract the buyers and/or to identify potential buyers. Therefore, in every auction only a small portion of overall stock is sold. It is common to hold several auctions for the same or similar items over time. We define the total number of units that the reseller wants to sell as *deal group*, N . Based on the deal-group size, the site owner negotiates a discounted price with their supplier. This price is termed as *deal price*, P .

In every auction, several units of the item are sold simultaneously. Each bidder can only bid on a single unit and is allowed to place only a prespecified number of bids. In other words, bidding on multiple items is not allowed, and the subsequent higher bids from the same bidder replace the earlier bids from the same bidder. Suppose that the auctioneer is selling n units ($n < N$) in a given auction; then, the bidders with top n bids win the auction. The entire stock of the product then is sold through repeated auctions or postauction solicitation. In this way, the auctioneer tries to estimate the relationship between prices and respective quantities for sale. In addition, such information is useful in estimating and negotiating for volume discounts with manufacturers and/or distributors. We next describe the auction process and derive some key results about the information available to the bidders during this auction process.

2.3. Characterization of the Auction Process

As mentioned earlier, deal price P and deal-group size N are defined before the auction begins and remain constant for its duration. These two parameters are not known to the bidders during the auction. Suppose that the auctioneer decides to sell n items per auction ($n < N$). Furthermore, assume that all bidders are allowed to place at most k bids during the auction. If $k = 1$, then we have a conventional sealed-bid auction, which has been extensively studied in the literature. Thus, we will focus on auctions with $k \geq 2$.⁴ The additional key rules of the auctions are:

- Before their *first* bid, bidders observe a minimum required bid which remains fixed for the duration of the auction, i.e., it does not change as the auction progresses.
- After the first bid is placed, bidders are *continuously* provided with the *weighted average*, that is calculated from the currently winning bids and deal price (we formally define this weighted average below). This is the only piece of information provided to the bidders with respect to the competitiveness of their bids.
 - As the auction progresses, bidders who have placed at least one bid may continue to observe the change in weighted average after each bid. Bidders can place up to $k - 1$ more bids at any time during the auction. A bidder is not allowed to revise her own bid downwards.
 - At the end of the auction, the n highest bidders are declared winners and can receive the items by paying the price equal to their own final (i.e., highest) bid.

Note that in this auction mechanism, the bidders do not automatically know whether their bid is winning; the only information they have is the weighted average, which is computed from the values of winning bids and the deal price. Therefore, one important question is whether there is a way for the bidders to determine whether their bid is winning at a given point in time.

First, let us define the weighted average formally. Denote WIN_t as the set of winning bids at auction state t (i.e., after t bids have been submitted). Then, $|WIN_t| = t$, when $t < n$ and $|WIN_t| = n$, when $t \geq n$. We will also use the $\min(WIN_t)$ notation to denote the smallest winning bid at any auction state t . Then, the weighted average at time t , denoted as A_t , can be expressed as

$$A_t = (1/N) \left(\sum_{b \in WIN_t} b + (N - |WIN_t|)P \right). \quad (1)$$

We next explore the properties of this weighted average A_t .

⁴ While the computational properties developed in this paper can be used to make better bids when $k = 2$, our approach is most

effective for cases when $k \geq 3$. We omit the analysis of the $k = 2$ case because of space limitations.

2.4. Properties of Weighted-Bid Average A_t

Let A be a finite collection (i.e., a multiset) of real numbers, i.e., $A = \{a_1, \dots, a_m\}$, where $a_i \in R$ ($i = 1, \dots, m$). Define the *average* of A as $avg(A) = (1/|A|) \sum_{a \in A} a$.

LEMMA 1. *Let A and B be finite collections of real numbers, where $avg(A) \leq avg(B)$. Then $avg(A) \leq avg(A \cup B) \leq avg(B)$. (The proof is based on straightforward arithmetic manipulations.)*

Now we can derive several properties of weighted average A_t .

THEOREM 1. *At any auction state t , the relationship between the true average of currently winning bids $avg(WIN_t)$, the current weighted average A_t , and the deal price P can be characterized as:*

- (i) *If $avg(WIN_t) \leq P$, then $avg(WIN_t) \leq A_t \leq P$;*
- (ii) *If $P \leq avg(WIN_t)$, then $P \leq A_t \leq avg(WIN_t)$.*

PROOF. Immediate from Lemma 1, by taking $A = WIN_t$ and B to be a collection of $N - |WIN_t|$ numbers, each equal to P , and noting that $avg(A) = avg(WIN_t)$, $avg(B) = P$, and $avg(A \cup B) = A_t$. \square

In summary, Theorem 1 states that weighted average A_t always stays between deal price P and the true average of winning bids, $avg(WIN_t)$. Intuition suggests that over time, as the bids increase, weighted average A_t should also increase; however, in some cases the weighted average may decrease. Theorem 2 and Corollary 2A describe the behavior of A_t over time.

THEOREM 2. *For any t such that $t \geq n$, we have $A_t \leq A_{t+1}$.*

PROOF. Because $t \geq n$, at least n bids have been submitted so far and, therefore, $|WIN_t| = n$. There are two possible scenarios: (1) If $b_{t+1} \leq \min(WIN_t)$, then $WIN_t = WIN_{t+1}$ and, consequently, $A_t = A_{t+1}$. (2) If $b_{t+1} > \min(WIN_t)$, then $b_{t+1} \in WIN_{t+1}$, i.e., it displaces bid $\min(WIN_t)$ ⁵ in the winning list. Based on the definition of A_t , we have $A_{t+1} = A_t + b_{t+1} - \min(WIN_t)$, and thus, $A_t < A_{t+1}$. \square

Assuming $t \geq n$, the following relationships are true based on Theorem 2:

$$b_{t+1} \notin WIN_{t+1} \Leftrightarrow A_t = A_{t+1} \quad \text{and}$$

$$b_{t+1} \in WIN_{t+1} \Leftrightarrow A_t < A_{t+1}.$$

⁵ Or the previous bid by the same bidder, if it is currently winning.

COROLLARY 2A. *If $A_t > A_{t+1}$, then $t < n$.*

PROOF. Immediate from Theorem 2 (by taking the logical negation of the implication). \square

Next, we further explore the dynamics of weighted average A_t to derive intelligent bidding strategies. From here on, we assume that the bidding has reached the stage when there are as many bidders as there are items ($t \geq n$).⁶ Also, recall that deal-group size is greater than the number of items for sale in a given auction ($N > n$). Furthermore, note that a bidder can use the observed weighted average information to formulate her subsequent bids. To illustrate the relationship between the weighted average and the winning bids, suppose that at some time t the observed weighted average is A_t . The following result explores the situation when the bidder strategically bids the current weighted average, i.e., the bidder places a new bid $b_{t+1} = A_t$.

THEOREM 3. *Assume that at auction state $t \geq n$, the following new bid is submitted: $b_{t+1} = A_t$. If $A_t = A_{t+1}$ (i.e., $b_{t+1} \notin WIN_{t+1}$), then for all winning bids $b \in WIN_t$, we have that $b \geq P$.*

PROOF. Assume otherwise, i.e., that $\min(WIN_t) < P$. When $t \geq n$, by the definition of A_t , we have that $A_t = (\sum_{b \in WIN_t} b + (N - n)P)/N$, and $\min(WIN_t)$ represents the smallest of N values (that are being averaged) in the numerator. Therefore, $\min(WIN_t) < A_t$.⁷ Because $b_{t+1} = A_t$, we have that $\min(WIN_t) < b_{t+1}$, and, consequently, b_{t+1} would displace $\min(WIN_t)$ in the winning list at auction state $t + 1$. This would result in $A_{t+1} > A_t$ —a contradiction. \square

In other words, Theorem 3 implies that (at any point in time $t \geq n$) if bidding the current weighted average does not cause the weighted average to change, then all the winning bids are equal to or higher than the deal price.

COROLLARY 3A. *Given auction state t , if there exists a currently winning bid $b \in WIN_t$ such that $b < P$, then new bid $b_{t+1} = A_t$ is guaranteed to be winning, i.e., $b_{t+1} \in WIN_{t+1}$.*

⁶ The situation where $t < n$ is not very interesting because any next bid, however small, will be a winning bid.

⁷ This is true because $N > n$, otherwise the relationship could possibly be $\min(WIN_t) \leq A_t$.

PROOF. Immediate from Theorem 3. \square

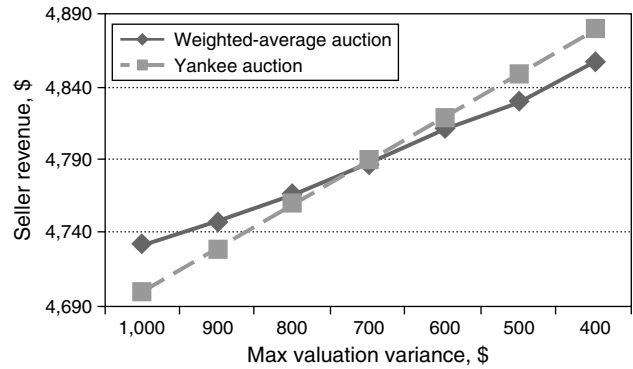
As stated in Theorem 1, A_t always lies between $avg(WIN_t)$ and P . Note that if $avg(WIN_t) < P$, at least one winning bid is guaranteed to be less than deal price P . In this case, bidding the current weighted average, i.e., $b_{t+1} = A_t$, will definitely place bid b_{t+1} in the winning bid list. However, when $avg(WIN_t) > P$, bidding the current weighted average A_t will not place the bidder in the winning list unless $\min(WIN_t) < A_t$.

By applying the results of Theorems 1–3 and strategically placing bids, it is possible for a bidder to compute whether her bid is among the winning bids. In addition, by exploiting the mathematical properties of weighted average A_t , bidding strategies can be devised to increase the chances of winning in an auction. In §3, we derive and describe these strategies. However, first we briefly discuss how our auction formulation may be applied in some related settings.

2.5. Generality and Applicability of the Weighted-Average Auction

Let us explore the circumstances that would lead a seller to choose the weighted-average auction format. As Banks et al. (2003) point out, transparency of information has an impact on bidder strategies and seller revenues in auctions.⁸ We conjecture that the use of the weighted-average auction can increase seller's revenues in cases when there is high uncertainty about the valuations of the bidders, for example, when new products are introduced to the market or when an auction is used to sell used commodities. In these cases, individual valuations can be highly variable. Because in most traditional auctions prices tend to coalesce around the marginal bid, a higher variance in valuation implies that, for the same number of units on sale, prices might be lower because marginal bid will typically be based on lower valuation (assuming that maximum valuation is stable and set by market information). To test our conjecture, we conducted a simulation-based analysis comparing the seller's revenue from multiunit English (Yankee)

Figure 1 Comparison of Seller Revenue for Yankee and Weighted-Average Auctions



auctions and weighted-average auctions, while identically varying uncertainty about the valuations of the bidders.

In particular, we used seven uniform distributions of valuations with increasing variance (as measured by the difference between minimum and maximum possible bids) to draw 10 sets of valuations from each distribution. Then, we simulated 100 pairs of Yankee and weighted-average auctions using these valuations and computed the average revenue for each distribution. All the auctions had five items for sale with a maximum possible individual valuation of \$1,000. We found that the weighted-average auction outperforms the Yankee auction with respect to the seller's revenue when the degree of uncertainty in bidders' valuations is relatively high. (See Figure 1.) We also found that relatively high revenues generated by the weighted-average auction are not caused by the excessive bidder aggressiveness because the bidding strategies used for the simulated bidders were risk averse.

As discussed earlier, this mechanism design considers auctions in the space of iterative bidding, while providing the opaqueness of feedback via the weighted-averaging technique. Next, we address the generality of the proposed design by showing that other conventional averaging techniques can be viewed as special cases of our general averaging mechanism.

Recall the definition of weighted average A_t in (1): $A_t = (\sum_{b \in WIN_t} b + (N - |WIN_t|)P)/N$. Suppose that the auctioneer decides to sell her entire stock in a single auction, thus making the deal-group size equal to the number of winners, i.e., $N = n$. In this case,

⁸ For example, Wes Shepherd, CEO of Channel Velocity, mentions that he tries to withhold the inventory information from auction participants. Channel Velocity sells excess inventory of consumer goods using eBay auctions. (See <http://channelvelocity.blogspot.com/2006/08/thoughts-on-auction-psychology.html> for more details.)

the second term in the above summation is equal to zero, and A_t becomes an arithmetic mean of winning bids. Because of the simplicity of calculating $A_t = \text{avg}(WIN_t)$, presenting such A_t as feedback to bidders would not provide the auctioneer with any protection against information derivation strategies described in the next section. However, it still could be used as a basis for allocation rule in the auction, for example, as in Spanish auctions.

It is also easy to see that our form of the weighted average encompasses the case of *bundle-size* weights used in Spanish treasury auctions. To see that, we need to make the same assumption $N = n$, and allow bidders to bid for multiple units. In this case, the bid of, say, \$300 for three units together, can be entered into calculation of the weighted average as three bids of \$100 for one unit. Such repeated counting of single-unit bids creates the effect of averaging by bid bundle size. One related question remains, however: Will the behavior of the bidders bidding for multiple units be any different than in the case of single-unit bids? We will answer this question after the bidding strategies for the iterative weighted-average auction are analyzed in the next two sections.

3. Deriving Intelligent Bidding Strategies

The properties of average bid metric A_t , described in §2, can be used to address a number of questions related to quality of bidder response in terms of formulating intelligent bids, such as: Is a given bid currently winning? Can we develop bidding strategies based on the available information that will increase bidder's chances to win in a given auction?

To apply the results of Theorems 1–3 for assessing whether a given bid is on the winning bid list, at some auction state t the bidder needs to obtain information about weighted average A_t that can be used to estimate important auction parameters. This task can be accomplished as follows:

- Assume that the strategic bidder placed the required minimum bid b_{\min} (her first bid) early in the auction to be able to continuously observe the weighted average.
- At auction state t ($t \geq n$), the bidder can observe the current weighted average A_t and then place bid

$b_{t+1} = A_t$ (her second bid overall) and observe A_{t+1} . By doing so, the bidder can check whether bidding has gotten above the deal price (by applying Theorem 3). Note that we are proposing the design of a software agent and, therefore, we assume that the agent will be able to participate in the auction early enough to place this second bid (and assess strategic information, as will be described later on) before all winning bids go above deal price P . With this bid, the bidder prepares for future bidding by trying to estimate deal-group size N and other important auction parameters, as discussed next.

3.1. Estimating the Auction Parameters by Strategic Bidding

Let us rewrite Equation (1) by representing the lowest of currently winning bids, $\min(WIN_t)$, at time t by α for convenience

$$A_t = \left(\alpha + \sum_{b \in WIN_t, b \neq \min(WIN_t)} b + (N - n)P \right) / N. \quad (2)$$

When the bidder places the bid at time $t + 1$ valued at A_t (her second bid overall), she displaces⁹ $\min(WIN_t)$ by overbidding it by some amount x , which we call *safety padding*. Therefore, we can write her second bid $b_{t+1} = A_t = \alpha + x$. We can then express weighted average A_{t+1} as

$$A_{t+1} = \left(\alpha + x + \sum_{b \in WIN_t, b \neq \min(WIN_t)} b + (N - n)P \right) / N. \quad (3)$$

Note that the difference between A_{t+1} and A_t is equal to x/N . To estimate safety padding x and deal-group size N , the bidder needs to place an immediate strategic bid $b_{t+2} = A_{t+1} + s = \alpha + x + s$, where s is a *strategic increment*. Because b_{t+2} is a bid by the same bidder (her third bid overall), it replaces b_{t+1} . One of the key issues in placing b_{t+2} is that it should be placed as soon as possible after placing b_{t+1} . As Greenwald et al. (2003) mention, software agents are best suited for such applications because the agents can place the bids in rapid succession (while recording the observed resulting weighted average). Assuming that the second and third bids can be placed in

⁹ Note that, while we are definitely displacing the lowest bid, we do not know how our bid fares compared to other winning bids. The new bid can be lowest as well as highest or midranked winning bid.

rapid succession without another bid affecting the weighted average, Theorem 4 provides the results about the auction parameters of interest.

THEOREM 4. *Given strategic bids b_{t+1} , b_{t+2} and observed weighted averages A_t , A_{t+1} , and A_{t+2} , the deal-group size, safety padding, and the value of the smallest winning bid at time t are given by*

$$\begin{cases} N = s/(A_{t+2} - A_{t+1}), \\ x = s(A_{t+1} - A_t)/(A_{t+2} - A_{t+1}), \\ \alpha = A_t - s(A_{t+1} - A_t)/(A_{t+2} - A_{t+1}). \end{cases} \quad (4)$$

PROOF. After the third bid b_{t+2} is placed, the weighted average can be represented as

$$A_{t+2} = \left(\alpha + x + s + \sum_{b \in \text{WIN}_t, b \neq \min(\text{WIN}_t)} b + (N - n)P \right) / N. \quad (5)$$

Obviously, $A_{t+2} - A_{t+1} = s/N$. We compute N , x , and α from the following system of equations:

$$\begin{aligned} b_{t+1} &= \alpha + x, & A_{t+1} - A_t &= x/N, \\ A_{t+2} - A_{t+1} &= s/N. \end{aligned} \quad (6)$$

Solving this system of equations, we obtain the expressions represented in (4). \square

Note that the actual value of strategic increment s must be determined carefully. On the one hand, it should not be too high (in order not to overbid unnecessarily). On the other hand, the value of s determines the ability to derive deal-group size, e.g., if A_t is reported with rounding of \$0.10, s should allow one dollar for every 10 members of the deal group. Thus, after strategically placing three bids, a bidder can observe average bids A_t , A_{t+1} , and A_{t+2} and also estimate the lowest winning bid α at time t , “safety padding” above the lowest winning bid x at time t , and deal-group size N .

Before we demonstrate the utility of the parameter estimation procedure in placing final bids for potentially winning the auction, we would like to address the issues related to the possible effect of any additional bids (i.e., bids by other auction participants) on the weighted average before the intended strategic bid b_{t+2} is placed.

3.2. Sensitivity of the Parameter Estimation Procedure to Bid Interference

After placing the first bid, the bidder can monitor the average bid continuously. Generally, it will change as other users place their bids. As discussed in §3.1, the parameter estimation procedure is based on the following observed weighted averages: A_t , A_{t+1} , and A_{t+2} . Thus, there are two time intervals of interest from the perspective of changes in the weighted average: $(t, t + 1)$ and $(t + 1, t + 2)$. Also, recall that every new bid that becomes a winner increases the weighted average.

Consider the interval $(t, t + 1)$. After t bids have been submitted, according to our parameter estimation procedure, the bidder is supposed to bid $b_{t+1} = A_t$. If someone else bids instead and becomes a winner, the actual weighted average at time $t + 1$ will change to A_t^* . Therefore, the bidder using our strategy is forced to bid A_t^* to be on the winning bid list. However, this change in weighted average does not matter because A_t^* is observed before making a bid.

On the other hand, any change in weighted average during the interval $(t + 1, t + 2)$ is more important. Ideally, we want to hold everything fixed, so that after placing the strategic bid b_{t+2} we know with certainty that we are displacing our own winning bid b_{t+1} and, therefore, can estimate the parameter values precisely (using Theorem 4). However, in practice this may not hold because of bids by other people interfering with our strategic bidder’s bids. We next analyze the impact of these potential interferences on the estimated parameter values. Assume that the weighted average changes to A_{t+1}^* ($A_{t+1}^* > A_{t+1}$) because of another bid during the time between observed A_{t+1} and reception of our strategic bid b_{t+2} . (For the moment, let us assume that there is exactly one such interfering bid.) The interference by another bid or bids can be classified in two broad cases based on whether bid b_{t+1} remains winning after the interfering bid was placed.

Case 1. Assume that the strategic bidder’s previous bid, b_{t+1} , is still among the winning bids. The net result of the interfering bid is that it increases the value of the winning bids by some amount δ as compared to the situation where there was no interfering

bid. More formally,

$$A_{t+1}^* = \left(\alpha + x + \delta + \sum_{b \in WIN_t, b \neq \min(WIN_t)} b + (N - n)P \right) / N. \quad (7)$$

Equation (7) implies that the difference between A_{t+1}^* and A_{t+1} is equal to δ/N . If deal-group size N is large or if the “interference” bid increment is small compared to N , the distortion will not be significant. Nonetheless, the impact of the additional bid would be that a larger than expected difference between A_{t+1} and A_{t+2} will be estimated with the net result of underestimating deal-group size N . Note that the formula for A_{t+1}^* in (7) applies in more general conditions. For example, if instead of just one there were several (up to $n - 1$) bids that became winners, then δ represents the cumulative interference bid increase and can be interpreted in the same way as before.

Case 2. Suppose that the strategic bidder’s previous bid, b_{t+1} , is displaced from the winning bid list because of another bid. In this situation, the new bid must replace the current lowest winning bid instead of the previous bid by the same bidder. Therefore, δ in (7) will also include the amount by which earlier strategic bid b_{t+1} is smaller than the current lowest winning bid. Thus, the estimation of the deal-group size would be even smaller than the situation described in Case 1.

In summary, if we look at Theorem 4 and the systems of Equations (4)–(6) and evaluate the impact of using A_{t+1}^* instead of A_{t+1} , we can make the following three observations:

- Because the “distorted” value of the weighted average is greater than the “undistorted” value and it is placed in the denominator of the formula for N , deal-group size will be underestimated.
- Similarly, safety padding x will be underestimated.
- Because the formula for the lowest currently winning bid α involves x with the negative sign, and because the latter is underestimated, α will be overestimated.

Thus, the interference of other bidders when placing strategic bids leads to an inflated estimate of the weighted average attributed to the strategic bid b_{t+2} . That, in turn, results in underestimation of the deal-group size and safety padding and overestimation of the lowest winning bid. On the one hand, thinking

that the lowest winning bid was higher than it really was will force the strategic bidder to place higher bids in the final stage of the auction (as we discuss in the next section). Therefore, the bidder will have higher chances of winning the auction. On the other hand, the bidder will be able to retain smaller surplus from the transaction.

However, as discussed earlier, using software agents for placing strategic bids can greatly reduce the chances of making these estimation errors because of the ability to place several bids in rapid succession and to obtain true weighted averages just before placing the strategic bids. In addition, depending on the number of bids that a bidder is allowed to bid, an agent can use an improved strategy by placing multiple bids in succession and estimating the parameters with respect to two successive bids repeatedly. If the two (or more) consecutive assessments are the same, then it is less likely that there was interference from another bid. If the estimates are different, then a larger estimate (of the deal-group size) is better to use because the smaller estimate was likely affected by interference.

Having provided the basic information extraction strategies and their vulnerabilities, in the next subsection we discuss strategies for winning the auction and associated trade-offs.


3.3. Strategies for Winning the Auction

We first provide a general result for the guaranteed winning bid at any time during the auction. We then argue that a surplus-maximizing bidder might like to win with a smaller bid than the guaranteed winning bid and provide a strategy for intelligent bidding that can facilitate that.

Suppose that the bidder placed a winning bid $b_{t+1} = A_t$ at time $t + 1$ (assuming $A_t < P$, as mentioned in Corollary 3A). If at some later time $t + y$ we have that $A_{t+y} > A_t$, then Theorem 5 provides the result about the guaranteed winning bid at that point in time.

THEOREM 5. *Suppose that the winning bid at time $t + 1$ was $b_{t+1} = A_t$ (where $A_t < P$) and, by time $t + y$ the weighted average was $A_{t+y} > A_t$. Then, a guaranteed winning bid at time $t + y$ is given by*

$$b_{t+y}^* = A_t + (A_{t+y} - A_t)(N/n). \quad (8)$$

Figure 2 Growth Curves for Weighted Average and True Average of Winning Bids

PROOF. From (1), we have that $A_t = (n/N)avg(WIN_t) + P(1 - (n/N))$ and, similarly, $A_{t+y} = (n/N)avg(WIN_{t+y}) + P(1 - (n/N))$. Thus, $A_{t+y} - A_t = (n/N)(avg(WIN_{t+y}) - avg(WIN_t))$. By rearranging this equality, we obtain $avg(WIN_{t+y}) = avg(WIN_t) + (A_{t+y} - A_t)(N/n)$. Because $A_t < P$, then $avg(WIN_t) < A_t$ (by Theorem 1). Hence, $avg(WIN_{t+y}) < A_t + (A_{t+y} - A_t)(N/n)$. Because $avg(WIN_{t+y})$ is the average of winning bids at auction state $t + y$, there must be a winning bid that is less than or equal to $avg(WIN_{t+y})$. Therefore, a bid of $A_t + (A_{t+y} - A_t)(N/n)$ would be a winning bid at auction state $t + y$. \square

COROLLARY 5A. b_{t+y}^* (as defined in Equation (8)) $> A_{t+y}$, i.e., the guaranteed winning bid is greater than weighted average at time $t + y$.

PROOF. The proof directly follows from (8). Because $N > n$, we have that

$$\begin{aligned} b_{t+y}^* &= A_t + (A_{t+y} - A_t)(N/n) > A_t + (A_{t+y} - A_t) \\ &= A_{t+y}. \quad \square \end{aligned}$$

Figure 2 provides a visual explanation of Theorem 5 and Corollary 5A. Essentially, because the true average of the winning bids $avg(WIN_t)$ grows faster than weighted average A_t , then if we know that A_t is a winning bid at time $t + 1$ and increase it by the calculated growth of the winning bids between t and $t + y$, we are guaranteed to obtain a winning bid at time $t + y$.

Theorem 5 provides a guaranteed winning bid at any time during an auction as long as the bidder can observe a weighted average early enough in the auction (i.e., while $A_t < P$). However, if at time $t + y$

at least one winning bid is still smaller than the deal price, then the bidder would win by bidding only the weighted average (based on Corollary 3A), which is less than the guaranteed winning bid provided in Theorem 5.

We can use an estimate of lowest currently winning bid, α (as specified in Equation (4)), to create a better estimate of a lower winning bid at time $t + y$. Let $C = N(A_{t+y} - A_t)$, which also denotes the difference between the sums of winning bids at time $t + y$ and time t (as can be easily derived from Equation (1)). Then, a reasonable lower estimate for the winning bid would be $\alpha + C/n$ because C/n captures the average growth in a winning bid. Note that a bidder is not guaranteed to win at this level. Furthermore, theoretically it is possible to win below this estimate as well, depending on the distribution of bids that arrived between t and $t + y$. However, it is clear that the expression for guaranteed winning bid in (8) provides an upper bound. Therefore, a bidder has a range of bids to try to win the auction by heuristically bidding $b_h = \alpha + C/n + \varepsilon$, where ε is in the range $[0, A_t - \alpha]$. Clearly, the probability of winning increases as $\varepsilon \rightarrow A_t - \alpha$, i.e., as bid b_h approaches b_{t+y}^* . In the next section, we discuss this strategy by developing a simulation model and exploring the probability of winning as a function of ε using the data generated from the simulation model.

4. Simulation Experiment to Test Bidding Strategies

Here we first present the specification of the simulation experiment, where the strategic agent bids using estimates derived in §3 against a set of bidders who bid using a random strategy based on the weighted average. We then analyze the results of the experiment and develop a heuristic for improving the performance of the bidding agent. We also study the performance of the bidding agent in the case when bidding for multiple units is allowed.

4.1. Experiment Setup

Theoretical results in §§2 and 3 provide lower and upper estimates for the winning bid at a given point in the auction. Specifically, a bid of $\min(A_T, \alpha + C/n)$ gives a lower estimate at any given time T in the auction. Note that we advocate taking the minimum of

current weighted average A_T and the progression of lowest bid ($\alpha + C/n$) because toward the end of the auction all winning bids may be above deal price P . If so, A_T is not going to be a winning bid but may serve as a good lower bound. On the other hand, because C/n reflects the progression of winning bids, expression $\alpha + C/n$ grows faster than the weighted average; however, if $\alpha + C/n < A_T$, it provides a tighter lower bound because it reflects that the average of winning bids still may be below A_T .

As discussed in §3.3 (see Equation (8)), bidding at the extrapolated average of winning bids $A_t + C/n$ is guaranteed to put the bidder into the winning bid list. However, bidding this amount is guaranteed to result in loss of surplus for the bidder, i.e., he or she would be paying more than necessary to win the auction. Our suggested strategy is to find ways to use the upper and lower estimates and intelligently make bids that are smaller than the guaranteed winning bid, while providing an acceptable probability of winning.

Let $b_{agg} = A_t + C/n$ denote the aggressive bid that is guaranteed to win, and $b_{opp} = \min(A_T, \alpha + C/n)$ denote the opportunistic bid that attempts to bid just above the lowest winning bid. We can then use a mixed strategy based on these bids by bidding

$$b_{mix} = b_{opp} + \rho(b_{agg} - b_{opp}), \quad (9)$$

where ρ is the parameter of aggressiveness, $0 \leq \rho \leq 1$. When $\rho = 1$, the bidder wants a guaranteed win and will pursue the aggressive strategy; when $\rho = 0$, the bidder cares mostly about paying the lowest amount possible in case of winning, but can tolerate the possibility of not winning at all. Intermediate values of ρ represent a continuum of mixed-bidding strategies.

To analyze the behavior of the auction under different bidding strategies by varying the values of ρ , we developed a simulation model in which the strategic bidder uses different strategies trying to win the auction. Our implementation of simulation was largely influenced by the actual auctions observed at the Internet consumer auction site Dealspin (now defunct).¹⁰ Its business model appeared to be based

on buying certain consumer goods in large volumes, thus realizing some volume discounts from the manufacturers or the wholesalers. These items are then sold in online auctions in lots of five. This site also solicited losing bidders in an auction by offering the item at a fixed price afterwards. Note that Dealspin used the information revelation mechanism similar to the one described in this paper. In other words, rather than providing the information about whether a bid is winning, it reports average spin, which is defined similarly to our definition of the weighted average. The following text defined the average spin on the site:

The “average spin” is a weighted average that guides users in making winning spins. It includes the top 5 spins, the total number of spinners (DealGroup), and the aggregate discount of the DealGroup (DealPrice)... As a result winning spins may lie above or below the “Average Spin.”

While observing several auctions at the site, we noted some interesting properties:

- There were always five units for sale in an auction, and a bidder was allowed to update their bids up to five times.
- While the range of list prices for items sold was anywhere between \$40 for floor scales to \$4,000 for plasma TVs, the majority of items were in \$150–\$1,000 range.
- Auction duration is fixed at 24 hours. There is no extended bidding; therefore, the end point of the auction is uniquely defined. This allows bidders to use time-based strategies.
- Bidding activity is slow at the beginning and end of an auction, but during some time periods it is very intense. Usually this happens in the 7th–8th hour of an auction. Bidding is so fast that one can see that the value of the average spin changes from the time the user decided to bid to the confirmation of the bid. It is virtually impossible for humans to process information that fast; only an agent can perform active bidding in this situation.

To model the auction mechanism, we built an auction simulator using Java programming language. We are interested in modeling the strategic behavior of an agent that is trying to win the auction against everyone else. Thus, there are two conceptual parts to the simulator: the strategic agent and the environment. The strategic agent always observes the flow of weighted average A_t . On the basis of this information,

¹⁰ Some archived Web pages from Dealspin.com can be found at <http://www.archive.org>.

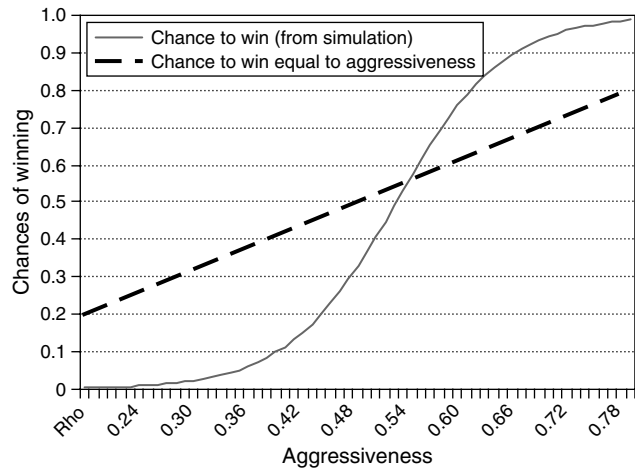
it places bids and performs inference as described in our theoretical results in §§2 and 3. If the strategic agent chooses not to bid in a particular round, then there will be a bid from a generic agent. Because we only track bids of the strategic agent, generic bids are represented by a random number drawn from a uniform distribution with the mean of current A_i and the range of $\pm 20\%$ around it. The first bids for all agents were initialized to the minimum required bid. All auctions consisted of 50 rounds of bidding.

Most of the complexity in the simulation resides in the behavior of the strategic agent. It places up to five bids depending on the current state of the auction. The first bid is placed at the beginning to start observing A_i , and it is the minimum required. The second and third bids are strategic and are used to estimate important auction parameters, as discussed in §3. The fourth bid is placed toward the end of the auction at the level of $\min(A_T, \alpha + C/n)$. If this bid does not win (as indicated by the unchanged weighted average¹¹), the fifth and final bid is placed using the mixed strategy defined in (9). In different replications of this auction, we change the value of aggressiveness parameter ρ to study the impact of ρ on the probability of winning. To control against the effect of “sniping” (winning the auction by bidding as close to the end as possible) and to focus on the pure effect of our suggested strategies, we have forced the strategic agent to place its final bid no later than when 80% of auction duration has passed.

Simulation allows us to study the relative performance of different bidding strategies. To do so, we considered 11 levels of ρ : 0.25; 0.30; 0.35...; 0.70; 0.75. To check for the possibility of nonstrategic influences, we also considered 59 price levels from \$100 to \$3,000 in increments of \$50. Thus, we had simulated 649 distinct auctions. To obtain reliable sampling estimates, we repeated each auction 30 times. This produced 19,470 data points for analysis.

¹¹ Clearly, this strategy is relevant when the number of bids that can be placed by an agent is restricted. If agents are allowed unlimited bids, another possible strategy is to use Theorem 3 and keep submitting increasingly higher bids until the reported weighted average changes. This is yet another task suitable for an intelligent agent.

Figure 3 Predicted Auction Win Dynamics



For every auction, we recorded the price level and aggressiveness combination as well as a binary variable indicating whether the strategic agent had been among the winners of the auction. Other parameters of simulation were held constant, i.e., the range of generic bids around the current weighted average (A_T) at $\pm 20\%$, starting bid at \$50, duration of auctions at 50 rounds, number of items for sale at 5, deal-group size at 50, and strategic increment at \$3.

4.2. Analysis of Simulation Data

When analyzing the data, we first explored choosing the appropriate level of strategic aggressiveness. Figure 3 depicts the winning probability line along with a straight line that depicts the proportional rate of growth for winning percentage as aggressiveness ρ increases. The winning probability curve is the outcome of logistic regression that is mapping agent aggressiveness to the agent’s chances of winning the auction. It demonstrates that bidders that do not want to bid more than at the aggressiveness level of 0.40 may not be able to win often. However, if bidders are willing to bid above $\rho = 0.40$, their chances of winning increase substantially.

As Figure 3 indicates, the area of fastest growth in winning percentage lies between aggressiveness levels of about 0.40 and 0.70. Using strategies that are too opportunistic is not beneficial because they will often lose; using strategies that are too aggressive is also not advisable because they will result in spending more money than is necessary to win. For example, if $\rho = 0.75$ is used, then the winning probability

is almost 0.99, whereas using $\rho = 0.6$ results in a winning probability of 0.70—a reasonable probability of winning at a substantial discount.

The mere fact of winning may not be an indication of a truly intelligent strategy—an agent may win by placing an extremely high bid, but the price it pays may be much higher than the value of the item to the agent’s owner. Therefore, we explored additional dimensions of strategy appropriateness. First, we considered the *position of the agent among the winners of the auction*. We have ordered the auction winners according to their final bid—the winner with the lowest bid being in position 1 and the winner with highest bid in position 5. Obviously, it is most efficient to win the auction in position 1 because other winners overpay relative to the efficient winner. Because we are not making any assumptions as to what the bidders’ true valuations are, the only comparisons that we can make are relative among bidders. Thus, we also computed two other value-related metrics: *overpayment w.r.t. the lowest winning bid* and *savings w.r.t. the average of winning bids*. Overpayment with respect to the lowest winning bid is the closest approximation to the surplus lost by all other winners, including the strategic agent; it is measured as a difference between the agent’s winning bid and the lowest winning bid expressed as a percentage of the lowest winning bid. Obviously, it should be zero in the efficient case (i.e., when the bid wins at position 1). Similarly, savings with respect to the average of winning bids are measured as a difference between the true average of all winning bids and the agent’s winning bid expressed as a percentage of the true average of all winning bids; it should be as large as possible in the efficient case.

We found that the agent’s aggressiveness level is a good predictor of all these additional metrics (quadratic regression models yielded p -values of 0.000 and R^2 between 20% and 50%). Table 1 presents the summary of our findings.

As can be seen in Table 1, increased aggressiveness of an agent increases its chances of winning the auction at the cost of being less efficient. While winning at the higher levels of aggressiveness, the agent ends up in the higher positions on the winning bid list and overpays a larger amount as compared to more efficient bidders. Eventually (i.e., at the highest

levels of aggressiveness), the agent starts losing edge even against the average of winning bids. Thus, it is important to devise a strategy that allows the agent to have sufficient chances of winning, while also being as efficient as possible. We have created a heuristic that serves this purpose, which will be discussed in the next subsection.

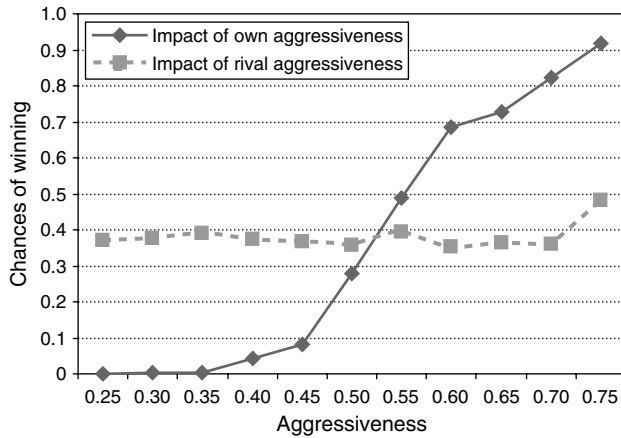
To complete the testing process, we have also studied the performance of the agent in the presence of other intelligent agents using the same bidding strategy. We considered cases where there were 1–3 rival agents with varying levels of aggressiveness, accounting up to 40% of all bids in the auction. This simulation produced very large data sets (e.g., agent with three rivals with 11 possible aggressiveness levels and 30 repeats per auction configuration has generated 439,230 data points), which causes almost all the lack-of-fit tests to come out as statistically significant. Thus, we cannot rely solely on statistics to judge the importance of these results, as described by Hosmer et al. (1997). From the practical viewpoint, our empirical findings were that:

- Presence and strategy of rivals has a much smaller effect on the outcome of an auction than the agent’s own strategy. Figure 4 illustrates this in the case of one strategic rival.
- Presence of strategic rivals had very little or no effect on agents pursuing very passive or very aggressive strategies.
- The nature of bidding agents in the auction has no impact on the auctioneer’s revenue.

Table 1 Key Metrics of Agent Performance

Aggressiveness	No. of wins (out of 30)	Mean win position	Overpayment vs. lowest bid, %	Savings against average, %
0.25	0	N/A	N/A	N/A
0.30	0	N/A	N/A	N/A
0.35	0	N/A	N/A	N/A
0.40	4	1.250	0.033	1.439
0.45	3	1.333	0.066	1.053
0.50	7	2.143	0.315	0.466
0.55	12	2.833	0.900	0.072
0.60	23	3.391	1.089	−0.206
0.65	26	3.577	1.212	−0.406
0.70	21	4.190	1.907	−0.962
0.75	27	4.222	1.650	−0.769

Note. N/A—Metric not applicable, as the agent did not win any auctions.

Figure 4 Impact of the Rival is Substantially Smaller Than the Agent's Own Strategy


Sample results in Table 2 illustrate the savings of the strategic agent against the average of winning bids in the presence of rivals.

4.3. Improved Estimation Heuristic

In this section, we present the approach to improve the agent's computation of the guaranteed winning bid. An ideal situation for the agent is to make sure that the upper bound on the guaranteed winning bid, $A_t + C/n$, at the time of the agent's bid is as close as possible to the true average of winning bids, $avg(WIN_{t+y})$, which is not directly observed. Theorem 1 implies that, if weighted average A_t is equal to deal price P , then it is also equal to $avg(WIN_t)$. Thus, if we observe A_t at the time when it is equal to (or very close to) the deal price, we will get the most accurate estimate of true average $avg(WIN_t)$. This

Table 2 Expected Savings (%) of the Strategic Agent w.r.t. the Winning Bid Average in the Presence of Rivals

Aggressiveness	0 rivals	1 rival	2 rivals	3 rivals
0.25	N/A	N/A	2.064	2.149
0.30	N/A	2.081	2.085	1.976
0.35	N/A	2.415	1.637	1.674
0.40	1.439	1.344	1.316	1.385
0.45	1.053	1.115	0.973	1.055
0.50	0.466	0.733	0.735	0.76
0.55	0.072	0.403	0.434	0.459
0.60	-0.206	0.103	0.123	0.165
0.65	-0.406	-0.221	-0.194	-0.154
0.70	-0.962	-0.639	-0.54	-0.475
0.75	-0.769	-0.885	-0.917	-0.832

Figure 5 Improved Upper Bound Estimate for a Guaranteed Winning Bid

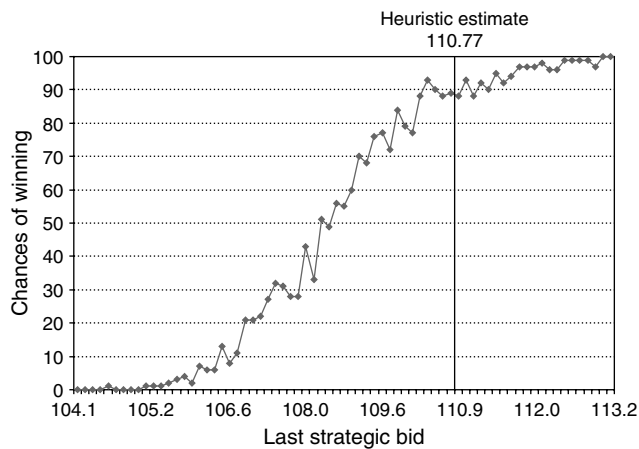
observation is represented in Figure 5—the upper bound for the winning bid based on computation at point t_2 is tighter than at t_1 (i.e., b_2^* is a tighter bound than b_1^*) because t_2 is closer to the point when A_t and $avg(WIN_t)$ are equal.

Therefore, we can leverage the agent's capability of being constantly present in the auction, i.e., it is possible to estimate when the critical point of $A_t = avg(WIN_t)$ is nearly achieved by continuously observing the changes in A_t .¹² Assuming that the auction participants base their strategies on the observed weighted average—the only available feedback from the auctioneer—they may be expected to bid above or below A_t with equal probability. When the bidding progression has not yet reached the deal price, then bids both above and below A_t have a chance to enter the winning list. However, the higher bidding goes above the deal price, the smaller are the chances for low bids to win (see Theorem 3). Therefore, the likelihood of bids entering the winning list decreases and, therefore, the growth of the weighted average will slow down.

Based on the realization that the growth rate of A_t slows down as the auction progresses, we implemented a detection heuristic in our agent. This heuristic is based on tracking the change of weighted average A_t over time and computing the moving

¹² The estimation procedure should be done while $A_t < P$ still holds. On the other hand, as illustrated in Figure 5, as the estimation is carried out closer in time to the crossing point, the resulting estimate becomes more accurate. The question of practical importance then becomes: How long to wait before starting the estimation? An agent, being omnipresent in an auction, can detect this moment more accurately than a human bidder; the ability to place bids in rapid succession also reduces the chance of missing the crossing point while the bid is being placed.

Figure 6 Heuristic Effectiveness in Winning the Auction



average of its change (i.e., growth) for the past three periods. Once the growth of the moving average gets below 1% (based on last estimated A_t), the agent starts its estimation strategy. To evaluate the heuristic performance, we computed the difference of the agent's estimate of guaranteed winning bid and the true lowest winning bid in the auction, as well as the true average of winning bids. For an agent with fixed-time bidding strategy (i.e., without the tracking heuristic), this difference was in the range of 4.23%–5.86% of the lowest winning bid, while for the heuristic-equipped agent, it improved to 1.53%–3.04%. The difference between the upper bound estimate and the true average of winning bids improved from 3.82% to 1.12%. Bidding at the heuristic level effectively puts the agent in a position to win more than 90% of the auctions. Figure 6 illustrates this for auctions where deal price $P = \$100$ and provides the probability of winning based on different levels of aggressiveness (as represented by the last strategic bid).

Because in real auctions it is impossible to predict the dynamics of bidding and place strategic bids at an exact, prescribed time, this heuristic leverages the omnipresent nature of the bidding agent to adapt to the bidding pattern of the auction. In other words, the proposed heuristic provides for an intelligent bidding strategy because it makes the bidding agent much more autonomous—the agent does not have to follow some predefined bidding scenario, but can time its bids according to the progression of a given auction.

4.4. Bidding Agent Performance in a Multiunit Scenario

To complete the analysis, we also discuss the performance of our proposed strategies in cases when the bidders are allowed to bid for more than one item. This discussion is important because it demonstrates possible similarities and differences between our weighted average auction and Spanish and other forms of multiunit or divisible-goods auctions. When it is desirable for agents to win more than one unit in an auction, theoretical analysis becomes much more complicated. There are two possible problems in this situation: the potential distortion of the estimation procedure and the need to consider the possibility of partial fulfillment in the outcome of the auction.

The problem of distorted estimates because of bid interference was explored in §3.2, where we have shown that interference of other bids during the strategic bidding procedure results in inflated estimates of bid growth. Recall that in the strategic bidding phase, the agent bids on the current weighted average, and, if it is early enough in the auction, such a bid is guaranteed to displace *one* bid from previous winners. In the case when an agent bids for multiple units during the parameter estimation stage (i.e., the early part of the auction), it risks introducing distortion into calculation. For example, if the agent bids the current average for three units, it may get one, two, or three one-item current winning bids at the same time. Unfortunately, we cannot know exactly how many one-item bids will enter the winning list as a result of such a strategy. Thus, for strategic purposes, an agent should not bid on more than one item during the parameter estimation stage, regardless of the number of items it is intending to win at the end.

Similarly, toward the end of an auction, the agent knows with some degree of confidence that placing a strategically computed bid will earn one spot on the list of winners. It is impossible to predict if it will be a high spot or a low spot (although more aggressive bidding will result in a place among higher winning bids, as was shown in our simulation results in §4.2). This may create a problem when the agent's goal is to win more than one item, e.g., if its final bid ends up being the second-lowest winning bid and the agent needs three items, it will have to settle for only two.

Table 3 Performance of the Bidding Agent in the Multiunit Case

Aggressiveness level	Auctions lost	Auctions won		Winning %	Partial fulfillment %
		Partial fulfillment	Total fulfillment		
0.50	27	3	0	10.00	100.00
0.55	21	9	0	30.00	100.00
0.60	16	11	3	46.67	78.57
0.65	14	13	3	53.33	81.25
0.70	14	9	7	53.33	56.25
0.75	12	10	8	60.00	55.56
0.80	10	8	12	66.67	40.00
0.85	3	13	14	90.00	48.15
0.90	4	6	20	86.67	23.08
0.95	2	9	19	93.33	32.14
1.00	3	5	22	90.00	18.52

We have conducted additional simulation analysis to see how the agent following our proposed strategies performs in multiunit settings. We have considered the setup when the agent's goal is to win four items at the end of a five-item auction. Because the number of items that a bidder can win depends on the position of her highest bid relative to other winning bids, the goals of winning two or three items are automatically included in our analysis.¹³ We also used the previously discussed estimation heuristic, which is aiming to win a single item in 90% of auctions. We ran the experiments for agent aggressiveness levels from 0.50 to 1.00 in increments of 0.05, keeping all other parameters as before (30 simulations at every level of aggressiveness, the range of generic bids around current weighted average A_T at $\pm 20\%$, starting bid at \$50, duration of auctions at 50 rounds, number of items for sale at 5, deal-group size at 50, and strategic increment at \$3). Our findings indicate that the agent is performing well even in the multiunit case. (See Table 3.)

We also explored the agent's ability to win the necessary number of units versus settling for partial fulfillment. The results indicate that with the increase of aggressiveness, the agent: (a) wins more

¹³ For example, if the agent's highest bid is the second-highest among five winners, it is in position to win four items, which also includes subsets of two and three items. However, if the agent's winning bid is the third-highest, it will win three items, but not four. Thus, the partial fulfillment problem is most severe with a higher number of items desired, and the results for fewer items will be more optimistic.

auctions in general, consistent with a single-unit scenario, and (b) gets all required items more often, i.e., it does not have to settle for partial allocation as frequently. Therefore, with the use of timely estimation and a strategic bidding procedure, the bidding agent is capable of performing its task in single-unit as well as multiunit situations.

5. Conclusions

In this paper, we argue that analysis of dynamic information provided by an online auction mechanism can be used to derive information that can help in formulating intelligent bids. We chose an auction mechanism that by design is meant to conceal information from the bidders. Our model provides implications for a class of auctions based on weighted average reporting, which are used in multiple domains, e.g., Treasury Bill Auctions, B2C auctions, and greenhouse emission allowance auctions. We derive analytical results based on strategic bid placements to get estimates of some key auction parameters. We argue that, in reality, smart agents should be designed to perform the task of placing strategic bids and extracting information because of their inherent advantage in quick placement of bids and information retrieval.

We then design strategies for intelligent bidding based on theoretical results about the guaranteed winning bid at any point in an auction and some reasonable lower bounds. We also design a simulation to test the heuristics for retaining larger surplus for bidders, while keeping a reasonably high probability of winning.

Although we consider the case where bidders are restricted to a certain number of bids, this restriction has very little effect on a bidder's ability to place strategic bids. While a large number of bids allow a bidder to compute parameters such as deal-group size with more certainty, the impact of potential estimation errors is likely to be small. However, from the seller's perspective, the number of allowable bids is more important, especially if the seller wants to restrict strategic bidding. For the auction mechanism explored in this paper, if bidders are allowed to place more than three bids, then bidders can use strategic bidding. In other words, with the current set of rules of the auction (as defined in §2), if the total number of bids allowed per person is three or fewer, then there

is no opportunity for strategic bidding by any single bidder. However, such restrictions may give rise to third-party strategic agents that may sell the information about estimates of key parameters, and thus reduce incentives for bidding higher. Another important realization for sellers is that if the deal-group size is advertised, there is no need for strategic bidding and bidders can always estimate the guaranteed winning bids at any point in time.

The limitations of this work are that we have not considered the endogenous impact of the choice of parameters made by the auctioneer. For example, a different set of parameters may result in significant changes in participation incentives of the bidders, thereby affecting the revenue generated from an auction. In future research, we will explore the issue of changes in auction parameters and its impact on auction participation, as well as agent strategies for participating in multiple auctions over time, e.g., whether the deal-group size can be estimated by monitoring multiple auctions for the same item.

Acknowledgments

Author names are in alphabetical order. The second author's research was supported in part by National Science Foundation (NSF) Career Grant IIS-0301239, but does not necessarily reflect the view of the NSF.

References

- Abbink, K., J. Brandts, P. Pezanis-Christou. 2006. Auctions for government securities: A laboratory comparison of uniform, discriminatory and Spanish designs. *J. Econom. Behavior Organ.* **61**(2) 284–303.
- Adomavicius, G., A. Gupta. 2005. Towards comprehensive real-time bidder support in iterative combinatorial auctions. *Inform. Systems Res.* **16**(2) 169–185.
- Alvarez, F., C. Mazon, E. Cerda. 2003. Treasury auctions in Spain: A linear approach. *Spanish Econom. Rev.* **5** 25–48.
- Ba, S., J. Stallaert, A. B. Whinston. 2001. Introducing a third dimension in information systems design: The case for incentive alignment. *Inform. Systems Res.* **12**(3) 225–239.
- Banks, J., M. Olson, D. Porter, S. Rassenti, V. Smith. 2003. Theory, experiment and the Federal Communications Commission spectrum auctions. *J. Econom. Behav. Organ.* **51** 303–350.
- Bapna, R., P. Goes, A. Gupta. 2003a. Analysis and design of business-to-consumer online auctions. *Management Sci.* **49**(1) 85–101.
- Bapna, R., P. Goes, A. Gupta. 2003b. Replicating online Yankee auctions to analyze auctioneers' and bidders' strategies. *Inform. Systems Res.* **14**(3) 244–268.
- 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. 2002. Optimal design of the online auction channel: Analytical, empirical and computational insights. *Decision Sci.* **33**(4) 557–577.
- Bicharra Garcia, A., A. Lopes, C. Bentes. 2001. Electronic auction with autonomous intelligent agents: Finding opportunities by being there. *Inteligencia Artificial, Revista Iberoamericana de Inteligencia Artificial* **13** 45–52.
- Chang, K. 2002. New age bidding—Are bots smarter? Pricingsociety.com. Available via archive.org at http://www.pricingsociety.com/Art_New_Age_Bidding_Are_Bots_Smarter_FF.htm.
- Fabra, N., N.-H. von der Fehr, D. Harbord. 2002. Modeling electricity auctions. *Electricity J.* **15**(7) 72–81.
- Greenwald, A., N. Jennings, P. Stone. 2003. Agents and markets. *IEEE Intelligent Systems* **18**(6) 12–14.
- Hosmer, D., T. Hosmer, S. Le Cessie, S. Lemeshow. 1997. A comparison of goodness-of-fit tests for the logistic regression model. *Statist. Medicine* **16** 965–980.
- Karuga, G., S. Maganti. 2004. Bidding agents in online auctions: What are they doing for the principal? *Proc. of the Tenth Americas Conf. on Inform. Systems*. New York, 1786–1793.
- Kauffman, R., C. Wood. 2006. Doing their bidding: An empirical examination of factors that affect a buyer's utility in Internet auctions. *Info. Tech. Management* **7**(3) 171–190.
- Ketter, W., J. Collins, M. Gini, A. Gupta, P. Schrater. 2006. Identifying and forecasting economic regimes in TAC SCM. H. La Poutre, N. Sadeh, S. Janson, eds. *Agent-Mediated Electronic Commerce: Designing Trading Agents and Mechanisms. Lecture Notes in Artificial Intelligence*, Vol. 3937. Springer, Berlin, 113–125.
- Klemperer, P. 2002. What really matters in auction design. *J. Econom. Perspect.* **16**(1) 169–189.
- Klusch, M. 2000. Agent-mediated trading: Intelligent agents and e-business. A. L. G. Hayzelden, R. Bourne, eds. *Agent Technology Applied to Networked Systems*. John Wiley, Chichester, UK.
- Kwasnica, A., J. Ledyard, D. Porter, C. DeMartini. 2005. A new and improved design for multiobject iterative auctions. *Management Sci.* **51**(3) 419–434.
- March, S., A. Hevner, S. Ram. 2000. Research commentary: An agenda for information technology research in heterogeneous and distributed environments. *Info. Systems Res.* **11**(4) 327–341.
- Mehta, K., S. Bhattacharya. 2006. Design, development and validation of an agent-based model of electronic auction. *Inform. Tech. Management* **7**(3) 191–212.
- Milgrom, P. 1989. Auctions and bidding: A primer. *J. Econom. Perspectives* **3**(3) 3–22.
- Milgrom, P. R., R. J. Weber. 1982. A theory of auctions and competitive bidding. *Econometrica* **50**(5) 1089–1122.
- Nautz, D., E. Wolfstetter. 1997. Bid sharing and risk aversion in multi-unit auctions with many bidders. *Econom. Lett.* **56** 195–200.
- Parkes, D. 2000. Optimal auction design for agents with hard valuation problems. A. Moukas, C. Sierra, F. Ygge, eds. *Agent Mediated Electronic Commerce II: Towards Next Generation Agent-Based Electronic Commerce System*. Springer, Berlin, 206–219.
- Roth, A., 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., S. Park. 2001. An elementary introduction to auctions. *Interfaces* **31**(6) 83–97.
- Sandholm, T. 2002. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence* **135**(1) 1–54.
- Stone, P., A. Greenwald. 2005. The first international trading agent competition: Autonomous bidding agents. *Electronic Commerce Res.* **5**(2) 229–265.
- Thomas, R. J., T. Mount, R. Zimmerman, W. Schulze, R. Schuler, L. Chapman. 2002. Testing the effects of price responsive demand on uniform price responsive demand on pay-as-offered and soft-cap electricity auctions. *35th Annual Hawaii Internat. Conf. Systems (HICCS)*-35.
- Tu, Z. 2005. Why do we use the Dutch auction to sell flowers?—Information disclosure in sequential auctions. <http://cedr.whu.edu.cn/cedrpaper/2007101516252.pdf>.
- Wurman, P., M. Wellman, W. Walsh. 2001. A parameterization of the auction design space. *Games Econom. Behav.* **35**(1–2) 304–338.