Optimal Bidding in On-line Auctions

Dimitris Bertsimas * Jeffrey Hawkins† Georgia Perakis ‡

December 2001, Revised December 2002

Abstract

On-line auctions are arguably one of the most important and distinctly new applications of the Internet. The predominant player in on-line auctions, eBay, has over 42 million users, and it was the host of over $9.3 billion worth of goods sold in the year 2001. Using methods from approximate dynamic programming and integer programming, we design algorithms for optimally bidding for a single item in an on-line auction, and in simultaneous or overlapping multiple on-line auctions. We report computational evidence using data from eBay's web site from 1772 completed auctions for personal digital assistants and from 4208 completed auctions for stamp collections that shows that (a) the optimal dynamic policy outperforms simple but widely used static heuristic rules for a single auction, and (b) a new approach for the multiple auctions problem that uses the value functions of single auctions found by dynamic programming in an integer programming framework produces high quality solutions fast and reliably.

*Boeing Professor of Operations Research, Sloan School of Management and Operations Research Center, Massachusetts Institute of Technology, E53-363, Cambridge, MA 02139, dbertsim@mit.edu.

†Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02139, jhawkins@mit.edu.

‡Sloan Career Development Associate Professor of Operations Research, Sloan School of Management and Operations Research Center, Massachusetts Institute of Technology, E53-359, Cambridge, MA 02139, georgiap@mit.edu.

1 Introduction

On-line auctions have become established as a convenient, efficient, and effective method of buying and selling merchandise. The largest of the consumer-to-consumer on-line auction web sites is eBay which has over 42 million registered users and was the host of over $9.3 billion worth of goods sold[1] in over 18 000 categories, ranging from consumer electronics and collectibles to real estate and cars. Because of the ease of use, the excitement of participating in an auction, and the chance of winning the desired item at a low price, the auctions hosted by eBay attract a wide variety of bidders in terms of experience and knowledge concerning the item for auction. Indeed, even for standard items like personal digital assistants we have observed a large variance in the selling price, which illustrates the uncertainty one faces when bidding.

eBay auctions have a finite duration (3, 5, 7, or 10 days). The data available to bidders during the duration of the auction include: the items description, the number of bids, the ID of all the bidders and the time of their bid, but not the amount of their bid (this becomes available after the auction has ended), the ID of the current highest bidder, the time remaining until the end of the auction, whether or not the reserve price has been met, the starting price of the auction, and the second highest price of the item, referred to as the listed price. The auction ends when time has expired, and the item goes to the highest bidder at a price equal to a small increment above the second highest bid.

eBay publishes on the web the bidding history of all of the auctions completed through its web site from the past thirty days. The bidding history includes the starting and ending time of the auction, the amount of the minimum opening bid set by the seller, the price for which the item was sold and, apart from the winning bid of the auction, the amount of every bid, and when and by whom it was submitted. For the winning bid of the auction only the identity of the bidder and submission date are revealed. In addition, if the auction was a reserve auction, then an indication of whether or not the reserve price was met. However, eBay does not publish the reserve price set by the sellers, and without this information we felt we could not properly model reserve price auctions. As a result we only consider auctions without a reserve price.

[1]http://www.shareholder.com/ebay/annual/2001_annual_10K.pdf

Literature Review

The literature for traditional auctions is extensive. For a survey of auction theory see Klemperer [10], Milgrom [12], and McAfee and McMillan [11]. The mechanism for determining a winner in an eBay auction is similar to that of a second-price sealed bid auction, also known as a Vickrey auction, see Vickrey [20]. In such auctions the optimal action, regardless of what the opponents are doing, is at some point to submit a bid equal to one's valuation of the item. The primary difference between a Vickrey auction and an eBay auction is that eBay reveals the identity of bidders and the value of the highest bid to date. In addition, the end of an auction on eBay is fixed in advance (i.e., there is a hard stop time). This makes it possible for bidders to submit bids close enough to the ending time of the auction and as a result, to not allow for competitors to respond. Such a strategy, known as sniping, has become so popular that a number of web sites exist to assist bidders in sniping (for example, see www.esnipe.com). In fact, we have found that for a personal digital assistant, model Palm Pilot III, the bids received per second in the final ten seconds is over 100 times greater than those received in the final day.

On-line auction allow bidders to participate in many auctions at once, or perhaps in many auctions in a short time span. The popularity of on-line auctions has motivated both theoretical and empirical investigations of bidding strategies. Taking into account network congestion, response time, and potentially other factors, Roth and Ockenfels [18] (see also Ockenfels and Roth [16]) provide evidence that there is a small but significant probability that a bid placed at the last seconds of an auction will not register on eBay's web site. This is an effect that our proposed algorithm explicitly accounts for. Roth and Ockenfels [18] show that if one is not certain that a submitted bid will be accepted, then there is no dominate bidding strategy. Furthermore, they argue that it is an undominated strategy to submit multiple bids. Bajari and Hortacsu [2] show that in a common value environment, sniping is an equilibrium behavior. Late bidding in on-line auctions has attracted a lot of interest from both practitioners and academics. Landsburg [13], suggests submitting bids late and bidding multiple times in order to keep others from learning and out-bidding him. Hahn [8] provides evidence that late bidding makes up 45% of all bids, but also that there is also a substantial amount of early bidding. Nonetheless, Hasker et al. [9] statistically reject that bidders commonly use a "Jump-call" strategy (a derivative of "Jump-bidding" from

Avery [1] for English auctions), but also a "Snipe-or-war" strategy. Mizuta and Steiglitz [15] simulate a bidding a bidding environment with early bidders and snipers and find out that early bidders win at a lower price but win fewer times on average. There is also evidence that bidders react to the ratings of sellers (see Lucking-Reiley et al. [14] and Dewan and Hsu [7]). In this paper we ignore this effect.

There has also been work done on bidding in multiple auctions. Oren and Rothkopf [17] consider the effects of bidding in sequential auctions against intelligent competitors and derive an infinite horizon optimal bidding strategy. Boutilier et al. [6] develop a piece-wise linear dynamic programming approximation scheme for bidding in multiple sequential auctions with complementaries and substitutability. Zheng [21] finds empirical evidence from eBay that bidders bid across multiple auctions simultaneously and that they tend to bid for the item with the lowest listed price. They also show that such a bidding strategy is a Nash equilibrium and results in lower payments for winners. Stone and Greenwald [19] consider a number of automated trading agents programmed to bid in multiple simultaneous auctions for complementary and substitutable goods. Bapna et al. [4] provide an empirical and theoretical study of observed bidding strategies in on-line auctions with multiple identical items.

Philosophy and contributions

Our objective in this paper is to construct algorithms that determine the optimal bidding policy for a given utility function for a single item in an on-line auction, as well as multiple items in multiple simultaneous or overlapping on-line auctions. In order to explain our modeling choices (see Section 2 for more details), we require that the model we build for optimal bidding for a potential buyer, called *the agent* throughout the paper, satisfies the following requirements:

(a) It captures the essential characteristics of on-line auctions.

(b) It leads to a computationally feasible algorithm that is directly usable by bidders.

(c) The parameters for the model can be estimated from publicly available data.

To achieve our goals we have taken an optimization, as opposed to a game theoretic approach. The major reason for this is the requirement of having a computationally feasible algorithm that is based on data and is directly applicable by bidders. Furthermore, our goal is to impose as few behavioral assumptions as possible and yet come up with bidding strategies that work well in practice (see also Sections 2.3, 2.4, 3.6 for some empirical evidence). Given that auctions evolve dynamically, in this paper we adopt a dynamic programming framework. We model the rest of the bidders as generating bids from a probability distribution which is dependent on the time remaining in the auction and the listed price, and can be directly estimated using publicly available data. Furthermore, we have tested our approach in a setting where there is a population and an additional competitor. We intend to show that by incorporating other strategies into the population bidding distribution (i.e., the agent is aware that the population may be also using "smarter" strategies) the approach suggested in this paper performs better when competing against other strategies. Finally, the first author has applied the algorithms in this paper many times in a real world setting to buy stamps and collections of stamps. The author's findings are that the algorithm is highly effective in that it both increases the chances of winning and decreases the amount paid per win. As a result, we feel that a dynamic programming approach gives rise to practical, realistic and directly applicable bidding strategies.

We feel that this paper makes the following contributions:

1. We propose a model for on-line auctions that satisfies requirements (a)-(c), mentioned above. The model gives rise to an exact optimal algorithm for a single auction based on dynamic programming.

2. We show in simulation using real data from 1772 completed auctions for personal digital assistants and 4208 completed auctions for stamp collections that the proposed algorithm outperforms simple, but widely used static heuristic rules.

3. We extend our methods to multiple simultaneous or overlapping on-line auctions. We provide five approximate algorithms, based on approximate dynamic programming and integer programming. The strongest of these methods is based on combining the value functions of single auctions found by dynamic programming using an integer programming framework. We provide computational evidence that the method produces high quality solutions fast and reliably. To the best of our knowledge, this method is new and may have wider applicability to high dimensional dynamic programming problems.

4. We test our algorithm in a multi-bidder environment against widely used bidding heuristics for both single and multiple simultaneous auctions. We show how our algorithms can be improved by incorporating different bidding strategies into the probability distribution of the competing bids.

Structure of the paper

The paper is structured as follows. In Section 2, we present our formulation and algorithm for a single item on-line auction. In Section 3, we present several algorithms based on approximate dynamic programming and integer programming for the problem of optimally bidding on multiple simultaneous auctions, and in Section 4, we consider multiple overlapping on-line auctions. The final section summarizes our contributions.

2 Single item auction

In this section, we outline the model in Section 2.1, the process we used to estimate the parameters of the model in Section 2.2, and the empirical results from the application of the proposed algorithm in Section 2.3.

2.1 The model

The length of the auction is discretized into T periods during which bids are submitted and where the winner, the highest bidder, is declared in period $T+1$. As the majority of the activity in an eBay auction occurs near the end of the auction, see Section 2.2 and [16], we have used the following $T = 13$ periods to indicate the time remaining in the auction: 5 days, 4 days, 3 days, 2 days, 1 day, 12 hours, 6 hours, 1 hour, 10 minutes, 2 minutes, 1 minute, 30 seconds, and 10 seconds remaining in the auction. These periods are indexed by $t = 1, \ldots, 13$ respectively. These time intervals were selected for two reasons: First, they were chosen in decreasing size in order to match the increasing intensity of bids as the auction draws to a close. Second, the periods were chosen at times which are naturally convenient for bidders to follow.

State

A key modeling decision is the description of the state. We define the state to be (x_t, h_t) for $t = 1, \ldots, T+1$, where

$$\begin{aligned} x_t &= \text{listed price at time } t, \\ h_t &= \begin{cases} \text{the agent's proxy bid if the highest bidder at time } t, \\ 0, \text{ otherwise.} \end{cases} \end{aligned}$$

We will often use the indicator $w_t = 1$ if $h_t > 0$, and zero, otherwise, to indicate if the agent is the highest bidder or not.

Control

The control at time t is the amount u_t the agent bids. We assume that the agent has a maximum price A up to which he is willing to bid for. Clearly, $u_t \in F_t = \{0\} \cup \{u_t | \ x_t \leq u_t \leq A\}$ if $w_t = 0$, and $u_t \in F_t = \{u_t | \ h_t \leq u_t \leq A\}$ if $w_t = 1$,

Randomness

There are three elements of randomness in the model:

(a) How the other bidders (*the population*) will react. In order to model the population's behavior, we let q_t be the population's bid. Note that $q_t = 0$ means that the population does not submit a bid at time t. We assume that $P(q_t = j | x_t, h_t)$ is known and estimated from available data, as described in Section 2.2.

(b) The proxy bid $\overline{h}_t$ at time t which is the highest bid to date if $w_t = 0$ (the agent is not the highest bidder). If, however, $w_t = 1$, then $\overline{h}_t$ is defined to be zero. The reason for this is that in this case, the proxy bid is known to the agent and is part of the state (denoted at h_t). In an eBay auction bidders know the listed price, but not the value of the proxy bid, unless of course they are the highest bidder. If a submitted bid is higher than the proxy bid, then the new listed price becomes equal to the old proxy bid plus a small increment. The exception to this is if a bidder out-bids his own proxy bid, in which case the listed price remains unchanged. For a given listed price, the minimum allowable bid is a small increment above the current listed price. We assume that if the agent is not the highest bidder, then the distribution of the proxy bid $P(\overline{h}_t = j | x_t, h_t = 0)$ is known and estimated from available data, as described in Section 2.2.

(c) Whether or not the bid will be accepted. As we have mentioned, near the last seconds in the auction, that is for $t = T$, there is evidence (see [16]) that a bid will be accepted with probability $p < 1$. This models increased congestion due to increased activity, low speed connections, network failures, etc. In all other times $t = 1, \dots, T-1$ the bid will be accepted. We use the random variable v_t, which is equal to one if the bid is accepted, and zero, otherwise. From the previous discussion, $P(v_t = 1) = 1$, for $t = 1, \dots, T-1$, and $P(v_T = 1) = p$.

Dynamics

The dynamics of the model are of the type

$$
\begin{aligned}
x_{t+1} &= f(x_t, h_t, u_t, v_t, q_t, \overline{h}_t) \\
h_{t+1} &= g(h_t, u_t, v_t, q_t, \overline{h}_t),
\end{aligned}
$$

where the functions $f(\cdot)$, $g(\cdot)$ are as follows:

$$w_t = 0,\ q_t \geq u_t \geq \overline{h}_t,\ v_t = 1 \ \Rightarrow\ x_{t+1} = u_t,\ h_{t+1} = 0, \tag{1}$$

$$w_t = 0,\ q_t \geq \overline{h}_t \geq u_t,\ v_t = 1 \ \Rightarrow\ x_{t+1} = \overline{h}_t,\ h_{t+1} = 0, \tag{2}$$

$$w_t = 0,\ \overline{h}_t \geq q_t \geq u_t,\ v_t = 1 \ \Rightarrow\ x_{t+1} = \max(q_t, x_t),\ h_{t+1} = 0, \tag{3}$$

$$w_t = 0,\ u_t > q_t \geq \overline{h}_t,\ v_t = 1 \ \Rightarrow\ x_{t+1} = q_t,\ h_{t+1} = u_t \tag{4}$$

$$w_t = 0,\ u_t > \overline{h}_t \geq q_t,\ v_t = 1 \ \Rightarrow\ x_{t+1} = \overline{h}_t,\ h_{t+1} = u_t \tag{5}$$

$$w_t = 0,\ \overline{h}_t \geq u_t \geq q_t,\ v_t = 1 \ \Rightarrow\ x_{t+1} = \max(u_t, x_t),\ h_{t+1} = 0, \tag{6}$$

$$w_t = 0,\ q_t \geq \overline{h}_t,\ v_t = 0 \ \Rightarrow\ x_{t+1} = \overline{h}_t,\ h_{t+1} = 0, \tag{7}$$

$$w_t = 0,\ \overline{h}_t \geq q_t,\ v_t = 0 \ \Rightarrow\ x_{t+1} = \max(q_t, x_t),\ h_{t+1} = 0, \tag{8}$$

$$w_t = 1,\ q_t \geq u_t \geq h_t,\ v_t = 1 \ \Rightarrow\ x_{t+1} = u_t,\ h_{t+1} = 0, \tag{9}$$

$$w_t = 1,\ u_t > q_t \geq h_t,\ v_t = 1 \ \Rightarrow\ x_{t+1} = q_t,\ h_{t+1} = u_t \tag{10}$$

$$w_t = 1,\ u_t \geq h_t > q_t,\ v_t = 1 \ \Rightarrow\ x_{t+1} = \max(q_t, x_t),\ h_{t+1} = u_t \tag{11}$$

$$w_t = 1,\ u_t > h_t \geq q_t,\ v_t = 1 \ \Rightarrow\ x_{t+1} = \max(q_t, x_t),\ h_{t+1} = u_t \tag{12}$$

$$w_t = 1,\ q_t \geq h_t,\ v_t = 0 \ \Rightarrow\ x_{t+1} = h_t,\ h_{t+1} = 0, \tag{13}$$

$$w_t = 1,\ h_t > q_t,\ v_t = 0 \ \Rightarrow\ x_{t+1} = \max(q_t, x_t),\ h_{t+1} = h_t \tag{14}$$

Eqs. (1)-(8) are for the case when $w_t = 0$. Eqs. (1)-(3) address the case that the population's bid is higher than the agent's bid, and the agent's bid is accepted. In Eq. (1), both the population and the agent bid above the proxy bid at time t, and thus the next listed price is u_t, and the agent is not the highest bidder. In Eq. (2) the highest price at time t is between the population's and the agent's bid, and thus the next listed price will be h_t, and the agent is not the highest bidder. In Eq. (3) both the population and the agent bid lower than the proxy bid at time t, and thus the next listed price is q_t, and the agent is not the highest bidder. Note that the max operator in Eqs. (3) and (6) cover the case that neither the population nor the agent bids ($q_t = u_t = 0$).

Eqs. (4)-(6) address the case that the population's bid is lower than the agent's bid, and the agent's bid is accepted, analogously to Eqs. (1)-(3). Finally, Eqs. (7), (8) cover the case that the agent's bid is not accepted. Note that the max operator in Eq. (8) covers the case that the population does not bid ($q_t = 0$).

Eqs. (9)-(14) address the case that $w_t = 1$, that is the agent is the highest bidder and hence has a proxy bid. In Eq. (9) both the population and the agent bid above the proxy bid at time t and the population bids higher, and thus the next listed price is u_t, and the agent is not the highest bidder. In Eq. (10) the agent bids higher than the population, and thus the next listed price is q_t, and the agent is the highest bidder. In Eqs. (11) and (12), the agent bids higher than the proxy bid, and thus the proxy bid is equal to u_t, while the listed price is updated to $\max(q_t, x_t)$. Note that we use strict inequalities to ensure that the agent is the highest bidder, since in the case of ties, the population is the highest bidder. Finally in Eqs. (13)-(14) $v_t = 0$ and so the populations bid is competing against the agent's proxy bid.

Objective

We assume that the agent wants to maximize the expected utility

$$\text{maximize } E[\mathcal{U}(x_{T+1}, h_{T+1})].$$

We will focus on the utility function

$$\mathcal{U}(x_{T+1}, h_{T+1}) = w_{T+1}(A - x_{T+1}). \tag{15}$$

The utility (15) implies that the agent will not bid for an item beyond his budget A, he wants to win the auction at the lowest possible price, and he is indifferent between not winning the auction and winning it at the budget A.

The choice of this particular model is guided by the requirements (a)-(c) outlined in the Introduction. We could include a more intricate state; for example we could include the number of bids at time t as an indicator of the auction's interest; however, the tractability of the model would decrease, but most importantly the estimation of the relevant probability distributions would become substantially more difficult given the sparsity of the data.

Bellman equation

The problem of maximizing the expected utility in a single item auction can be solved using the Bellman equation:

$$
\begin{aligned}
J_{T+1}(x_{T+1}, h_{T+1}) &= \mathcal{U}(x_{T+1}, h_{T+1}) \\
&\text{If } w_t = 0 \text{ then} \\
J_t(x_t, h_t) &= \max_{u_t \in F_t(x_t,h_t)} E_{q_t, v_t, \overline{h}_t}[J_{t+1}(x_{t+1}, h_{t+1})], \quad t = 1, \ldots, T, \\
&= \max_{u_t \in F_t(x_t,h_t)} \sum_{q=0}^{A} \sum_{v=0}^{1} \sum_{\overline{h}=x_t}^{A} J_{t+1}(f(x_t, h_t, u_t, q, v, \overline{h}), g(h_t, u_t, q, v, \overline{h})) \\
&\quad \cdot P(q_t = q, \overline{h}_t = \overline{h} | x_t) P(v_t = v) \qquad (16)
\end{aligned}
$$

$$
\begin{aligned}
&\text{If } w_t = 1 \text{ then} \\
J_t(x_t, h_t) &= \max_{u_t \in F_t(x_t,h_t)} E_{q_t, v_t}[J_{t+1}(x_{t+1}, h_{t+1})], \quad t = 1, \ldots, T, \\
&= \max_{u_t \in F_t(x_t,h_t)} \sum_{q=0}^{A} \sum_{v=0}^{1} J_{t+1}(f(x_t, h_t, u_t, q, v, 0), g(h_t, u_t, q, v, 0)) \\
&\quad \cdot P(q_t = q | x_t) P(v_t = v). \qquad (17)
\end{aligned}
$$

Note that in Eq. (16), when the agent does not have a proxy bid, the expectation is taken over q_t, $\overline{h}_t$ and v_t, whereas in Eq. (17), when the agent has a proxy bid, the expectation is taken over only q_t and v_t ($\overline{h}_t = 0$). We set $P(q_t = A, \overline{h}_t = \overline{h}|x_t)$, $P(q_t = q, \overline{h}_t = A|x_t)$ equal to $P(q_t \geq A, \overline{h}_t = \overline{h}|x_t)$ and $P(q_t = q, \overline{h}_t \geq A|x_t)$ respectively, since if a bid from the population is ever greater than or equal to A then the agent cannot win. If the agent has a proxy bid, then we set $P(q_t = A|x_t) = P(q_t \geq A|x_t)$.

2.2 Estimation of parameters

As we have mentioned, perhaps the most important guiding principle for the current model, is that the model's parameters should be estimated from the data that is publicly available from eBay. eBay publishes the history of auctions, and thus the prices h_t are readily available, with the exception of h_{T+1}, which is not publicized. Given this information, and the time of bids and identity of bidders, we calculate the listed price reported to the bidder when their bid was submitted. We can thus find the empirical distribution for $P(q_t = j|x_t, w_t)$ and $P(h_t = j|x_t, w_t)$. We have found no dependence on w_t, and thus we calculated $P(q_t = j|x_t)$ and $P(h_t = j|x_t)$. To reduce the size of the estimation problem, and to eliminate having to deal with extremely sparse distribution matrices, we round up bids u_t and listed prices x_t to $\lceil u_t/10 \rceil$ and $\lceil x_t/10 \rceil$. For example, an observed listed price of \$45 at time t is counted as $x_t = 5$.

Since we are modeling only a single competing bid from the population and not the many that could arrive during a given time period, we calculate the distribution of the maximum bid to occur for a given (x_t, t). Let $\hat{q}_s$ be an actual bid at a real time s, and similarly for $\hat{x}_s$, and let $\hat{s}_t$ be the actual time, in seconds, at which period t begins. Thus, we calculate $\max_{\hat{s}_t \leq s < \hat{s}_{t+1}} \hat{q}_s$. Then,

$$P(q_t = q | x_t) = P\left(q_t = \left\lceil \frac{\max_{\hat{s}_t \leq s < \hat{s}_{t+1}} \hat{q}_s}{10} \right\rceil \;\middle|\; x_t = \left\lceil \frac{\hat{x}_{\hat{s}_t}}{10} \right\rceil \right),$$

where the right hand side is calculated empirically.

We have calculated the empirical bidding distribution, adjusted as described above, for personal digital assistants (PDAs) and stamp collections, in an attempt to capture the effect of private and common value auctions, respectively. For PDAs, we looked at the Palm Pilot III model, whose final selling price was between \$70 and \$200. In total, there were 22478 bids in 1772 auctions over a two week period, with the mean auction lasting 5 days and receiving bids from just over 4 unique bidders on average. As an example, Figure 1 presents the empirical distribution of bids submitted between 6 and 12 hours from the end of the auction. Note that for a given listed price, bids are either zero (no bid), or they are distributed at values above the listed price (the 22478 recorded bids do not include 'zero bids'). Similarly, we have also calculated the bidding distribution of the population for stamp collections with final selling prices ranging from \$100 to \$1000. The data was taken from 4208 completed auctions with 50766 total bids during the same period, with the mean auction lasting 7.5 days and receiving bids from 3 different bidders on average. For this set of data, bid increments of \$50 were used. The empirical distribution of $P(h_t = j | x_t)$ has been calculated similarly.

As noted earlier, Roth and Ockenfels [18] observed that the number of bids increases as auctions near their end, and that the distribution of the arrival time of bids in the final seconds obeys a power law. In order to capture this phenomenon for the data we used, we consider different time horizons, denoted by S, before the end of the auction: 3 days, 6 hours, 10 minutes and 1 minute. For each separate S, we partition the time interval $[0, S]$ into ten subintervals $a_1 = [0, 0.1S]$, $a_2 = [0.1S, 0.2S], \ldots, a_{10} = [0.9S, S]$, so that a_1 represents the final tenth of an interval of total length S. For example, when $S = 10$ minutes, a_1 is the final minute of the auction. For each

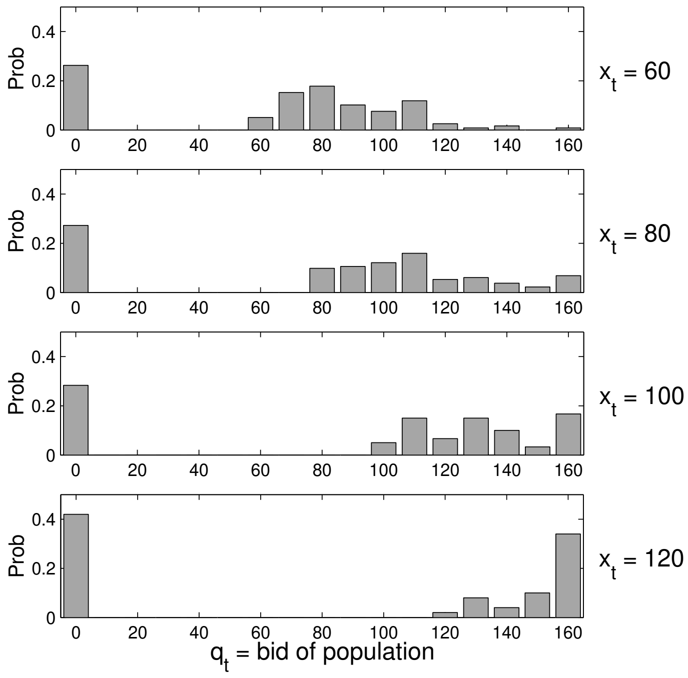

Figure 1: The empirical distribution of the population's bid q_t for Palm Pilot IIIs for a given listed price x_t. Here t represents the time period of 6 to 12 hours remaining in the auction. Note that since the agent's budget is $A = \$150$, bids by the population above \$160 are counted as \$160.

interval a_i, $i = 1, \ldots, 10$ we record the fraction of all the bids in $[0, S]$ that arrived within this period. Figure 2 shows the fraction of bids in each interval a_i as a function of the percentage of the respective time scale, that is $0.1 \times i$, for all the four values of S for the data for Palm Pilots III and stamp collections. Figure 2 suggests that the distribution of the timing of these bids is identical for the times S equal to 3 days, 6 hours and 10 minutes. For S equals 1 minute, it is still the same for all but the first interval a_1, that is, within 6 seconds, before the end of the auction. An explanation of this phenomenon is to assume that due to network congestion and other phenomena, there is a probability p of a bid being accepted during the last seconds of an auction. An approximate estimate of p is then given as follows.

We first make a distinction between submitted bids, and accepted bids. The former are bids intended to be submitted, and the latter is what registers on the web site. For all subintervals except that of a_1 for interval S=1 minute, submitted bids are accepted bids. However, for the final subinterval for S=1 minute, submitted bids are accepted bids with probability p. We assume that the distribution of submitted bids over a particular interval S is the same for all intervals. For $S = 1$ minute, the observed fraction of bids arriving in interval a_1 is

$$\begin{aligned} P(\text{accepted in } a_1|\text{accepted}) &= \frac{P(\text{accepted in } a_1)}{P(\text{accepted})} \\ &= \frac{P(\text{accepted in } a_1|\text{submitted in } a_1)P(\text{submitted in } a_1)}{\sum_{i=1}^{10} P(\text{accepted in } a_i|\text{submitted in } a_i)P(\text{submitted in } a_i)}. \end{aligned} \quad (18)$$

$P(\text{accepted in } a_i|\text{submitted in } a_i) = p$ for $i = 1$ and equals 1, otherwise. Figure 2 suggests that $0.41 = (0.45 + 0.42 + 0.36)/3 = P(\text{submitted in } a_1)$. Figure 2 also suggests that for S=1 minute $P(\text{accepted in } a_1|\text{submitted in } a_1) = 0.16$. Then, from Eq. (18) we have $0.41p/(0.41p+(1-0.41)) = 0.16$, leading to an estimate of $p \approx 0.27$. We have tested DP over a broad range of p and found that for the different values there is no qualitative difference in the results. In our experiments, we use a p-value of 0.8, which is based roughly on experience.

2.3 Empirical results

Having estimated its parameters, we have applied the model as follows:

(a) For bidding for a Palm Pilot III, we used a utility of the form (15) with a budget of \$150. Since we have clustered the data into \$10 increments the utility function becomes $\mathcal{U}(x_{T+1}, h_{T+1}) =$

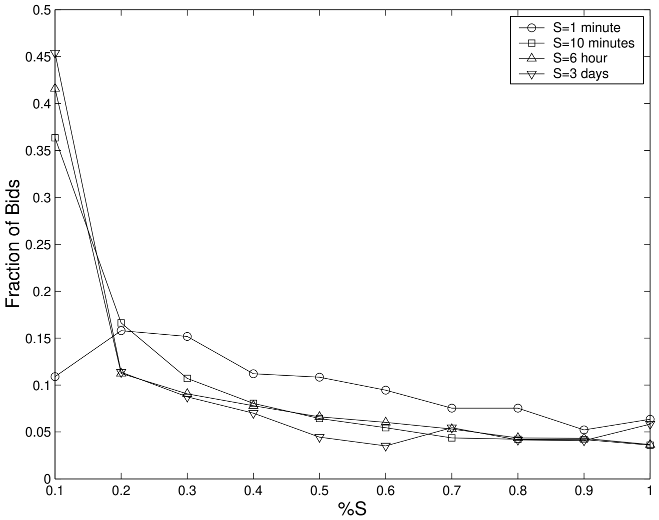

Figure 2: The fraction of bids in each interval a_i as a function of the percentage of the respective time scale, that is $0.1 \times i$, for all the four values of S for the data for Palm Pilots III and stamp collections.

$10(A - x_{T+1})w_{T+1}$ with $A = 15$ and where x_t measures the listed price in tens of dollars. We set $T = 13$ using the time steps described earlier.

(b) For bidding for stamp collections, we used a budget of \$500, \$50 increments, and a utility function $\mathcal{U}(x_{T+1}, h_{T+1}) = 50(A - x_{T+1})w_{T+1}$ with $A = 10$, to represent the budget of \$500.

To test the performance of the algorithm in simulation we first compute the optimal cost to go and optimal decision for every state (x_t, h_t), for $t = 1, \ldots, T$ using Eq. (16). For the purposes of the simulation experiment, bids are drawn from the same distribution for which the algorithm was constructed and upon arriving in a new state of the auction, the optimal bid is determined following the dynamic programming algorithm. The next states are computed using update rules (1)-(8) and the auction proceeds. At the end of the auction, period $T+1$, the winner is declared and the appropriate utility is received. The following reported results are based on 10,000 simulations.

The optimal bidding policy depends on the estimated data. In Tables 1 and 2, we report the empirically observed optimal bidding policy for the estimated data for $p = 0.27$ and $p = 0.8$ respectively. Bidding in the early stages of the auction is not optimal since it can only lead to higher listed prices later on in the auction. However, because bids submitted in period T are not

State	Time Period			
	1-11	12		13
	$w_t = 0,\ w_t = 1$	$w_t = 0$	$w_t = 1$	$w_t = 0,\ w_t = 1$
x_t, h_t	no bid	$\min(x_t + 10, A)$	$\min(\max(x_t + 11, h_t), A)$	A

Table 1: Approximation of the optimal bidding policy for Palm Pilot IIIs with $p = 0.27$.

State	Time Period			
	1-11	12		13
	$w_t = 0,\ w_t = 1$	$w_t = 0$	$w_t = 1$	$w_t = 0,\ w_t = 1$
x_t, h_t	no bid	$\min(x_t + 7, A)$	$\min(\max(x_t + 9, h_t), A)$	A

Table 2: Approximation of the optimal bidding policy for Palm Pilot IIIs with $p = 0.8$.

guaranteed to be accepted, it is optimal to submit a bid in $T-1$. As bids in period $T-1$ also lead to a higher listed price in T, a trade off emerges between having a proxy bid and causing the listed price to be too large. As expected, the tables show that it is optimal to bid more when p is smaller in time period $T-1$. Note that the algorithm suggests bidding more when the agent is the highest bidder in $T-1$. This is due to the reduction of uncertainty one faces when he is the highest bidder. These two tables show that qualitatively, there is a difference in bidding strategy resulting from the two values of p tested, but it is very small.

Tables 3 and 4 show the effect p has on the performance of *DP* for Palm Pilots and stamp collections, respectively. In both cases the effects are small.

p-value	Win %	Avg. Utility	Avg. Spent per Win
0.27	69.0	30.5	105.8
0.7	69.5	30.7	105.8
0.8	69.6	30.7	105.8
0.9	69.4	31.0	105.4
1.0	70.2	32.0	104.4

Table 3: Performance of DP Policy for Palm Pilot IIIs for a range of p-values.

p-value	Win %	Avg. Utility	Avg. Spent per Win
0.27	98.9	332.7	163.6
0.7	98.9	332.7	163.6
0.8	98.9	332.7	163.6
0.9	98.9	332.8	163.5
1.0	99.0	333.7	163.0

Table 4: Performance of DP Policy for Stamps for a range of p-values.

Table 4 shows the varying effect p has on the average utility and winning percentage when bidding for stamps. The high budget of the seller relative to what the market is willing to pay means that the effect of p is small. The conclusion drawn from Tables 1 - 4 is that the effect of the value of p on the performance of DP is small. For the remainder of this paper we will use $p = 0.8$.

Table 5 shows the results of the algorithm after 10,000 simulations with $A = 15$, for four different bidding strategies for stamp collections: (a) The dynamic programming policy; Bidding the budget A (b) at time $t = 0$ (the beginning of the auction); (c) at time $t = T - 1$; (d) at time $t = T$. The dynamic programming based policy was clearly the best. Although it didn't lead to wins as often as bidding A at $t = 0$ or $t = T - 1$, the average utility was far larger. Note that the average utility is equal to the probability of winning times 150 minus the average spent per win.

The reason for dynamic programming's success is that it is not restricted to making bids at specified times, but can instead manipulate the auction and bid when required. On average, the agent spent \$105.8 per win using the dynamic programming based policy. We implemented this algorithm using similar data to bid for a Palm Pilot III in an on-line auction and the item was won for \$92.

Table 6 shows the results of the algorithm after 10,000 simulations with $A = 10$, for different bidding strategies for stamp collections. In this case, the listed price and all bids were rounded to \$50 increments. Again the optimal policy is the clear winner. Not only does it win 99% of the time, it spends \$163.6 per win, versus 98% winning percentage and \$168.4 per win for the next closest policy. We have used this algorithm to win over one thousand stamp collections and individual stamps in eBay.

Policy	Win %	Avg. Utility	Avg. Spent per Win
DP	69.6	30.7	105.8
Bid A at $t = 0$	51.8	23.1	105.4
Bid A at $t = T - 1$	55.9	24.3	106.5
Bid A at $t = T$	46.5	20.8	105.3

Table 5: Performance of bidding strategies for Palm Pilot IIIs.

Policy	Win %	Avg. Utility	Avg. Spent per Win
DP	98.9	332.7	163.6
Bid A at $t = 0$	93.9	285.0	196.6
Bid A at $t = T - 1$	98.2	325.9	168.4
Bid A at $t = T$	78.6	260.0	169.3

Table 6: Performance of bidding strategies for stamp collections.

2.4 Bidding against multiple competitors

In this section we consider an agent bidding against both the population's bid and an additional competitor's bid. The purpose of this analysis is to illustrate the robustness of the DP as well as show how information about the competitor's strategy can be used to improve the performance of the DP.

Tables 7 and 8 show the results of bidding against a competitor of different budgets for Palm Pilots when the agent's budget is 150. In Table 7, the strategy of the competitor is to bid his budget at time $T - 1$. Table 8 applies to the case that the competitor's strategy is also a DP strategy solved for the particular budget. In both cases the competitor's utility is equal to his budget minus the price paid if he wins the object, and zero otherwise. Both tables show that as the competitor's budget increases the agent's expected utility decreases only slightly, except for when both parties have the same budget. Note that because the DP's objective is to maximize expected utility, and not the probability of winning, the strategy employed by the agent allows the competitor to sometimes win even with a smaller budget.

In the case of a Palm Pilot, Table 9 shows the results of bidding against a competitor, when

Competitor's Budget	*DP* Win %	*DP* Avg. Utility	*DP* Avg. Spent per Win	Competitor Win %	Competitor Avg. Utility	Competitor Avg. Spent per Win
100	67.8	25.7	112.0	1.0	0.1	89.7
110	68.5	22.7	116.9	2.0	0.2	98.6
120	65.5	17.5	123.3	3.4	0.5	105.0
130	65.7	12.6	130.8	4.4	0.9	109.3
140	58.5	5.8	140.0	7.4	1.4	120.8
150	6.5	0.0	150.0	58.0	2.5	145.7

Table 7: Performance of bidding against Policy 'Bid A at T-1' for Palm Pilot IIIs, the agent's budget is 150.

Competitor's Budget	*DP* Win %	*DP* Avg. Utility	*DP* Avg. Spent per Win	Competitor Win %	Competitor Avg. Utility	Competitor Avg. Spent per Win
100	60.1	20.2	116.3	0.0	0.0	95.2
110	60.5	18.2	119.9	0.1	0.0	90.0
120	59.5	15.4	124.1	0.3	0.0	107.8
130	59.4	11.6	130.5	0.6	0.1	115.4
140	45.7	4.9	139.3	3.8	0.1	136.6
150	22.0	0.8	146.5	22.0	0.8	146.5

Table 8: Performance of bidding against DP Policy. for Palm Pilot IIIs, the agent's budget is 150.

Probability of Competitor's Entrance	*DP*			Competitor		
	Win %	Avg. Utility	Avg. Spent per Win	Win %	Avg. Utility	Avg. Spent per Win
0.0	58.5	5.8	140.0	7.4	1.4	120.8
0.25	57.5	5.8	140.0	10.1	3.3	107.4
0.5	56.0	5.6	140.0	11.7	3.9	106.3
0.75	56.8	5.7	140.0	11.7	4.1	105.4
1.0	62.5	6.2	140.0	0.0	0.0	-

Table 9: Performance of bidding against Policy 'Bid A at T-1' with competitor's budget of 140 for Palm Pilot IIIs, for different anticipated entrance probabilities.

the agent's budget is 150 (this involves first solving the Bellman Equations (16) and (17) with the population bids q_t and $\bar{h}_t$ and the competitor's bid). However, the competitor's bid was present only with a given probability to reflect the agent's uncertainty as to whether the competitor would be present or not later in the auction. Nevertheless, in the simulations the competitor did bid in every auction. The agent's anticipated probability of the competitor bidding is reflected in the column 'Probability of Competitor's Entrance'. The simulations use a competitor with a budget of 140. The poor performance of the DP for entrance probabilities of less than one, is a result of the DP attempting to win at low prices without anticipating the scenario in which there is a competitor bidding.

3 Multiple auctions

We consider an agent interested in participating in N simultaneous auctions all ending at the same time. In each auction $i = 1, \ldots, N$, the agent is willing to bid no more than A_i, and no more than A over all auctions.

For $t = 1, \ldots, T+1$, and $i = 1, \ldots, N$ the state of each auction is (x_t^i, h_t^i); the control is u_t^i; randomness is given by the vector $(q_t^i, v_t^i, \overline{h}_t^i)$. We denote the corresponding vectors by $(\boldsymbol{x_t}, \boldsymbol{h_t})$, $\boldsymbol{u_t}$ and $(\boldsymbol{q_t}, \boldsymbol{v_t}, \overline{\boldsymbol{h}}_{\boldsymbol{t}})$. We use $w_t^i = 1$ if $h_t^i > 0$, and zero, otherwise; $\boldsymbol{w_t}$ denotes the vector of w_t^i. The set of feasible controls is given by:

$$F_t(\boldsymbol{x_t}, \boldsymbol{h_t}) = \left\{ \boldsymbol{u_t} \,\middle|\, u_t^i \in F_t(x_t^i, h_t^i),\ i = 1, \ldots, N,\ \sum_{i=1}^{N} u_t^i \leq A \right\}.$$

The utility is given by

$$\mathcal{U}(\boldsymbol{x_{T+1}}, \boldsymbol{h_{T+1}}) = \sum_{i=1}^{N} (A_i - x_{T+1}^i) w_{T+1}^i,$$

and the dynamics are given analogously to Eqs. (1)-(8). We denote by $\boldsymbol{f_t}(\boldsymbol{x_t}, \boldsymbol{h_t}, \boldsymbol{u_t}, \boldsymbol{q_t}, \boldsymbol{v_t}, \overline{\boldsymbol{h}}_{\boldsymbol{t}}) = (f_t(x_t^1, h_t^1, u_t^1, q_t^1, v_t^1, \overline{h}_t^1), \ldots, f_t(x_t^N, h_t^N, u_t^N, q_t^N, v_t^N, \overline{h}_t^N))$, and likewise for $\boldsymbol{g_t}(\cdot)$. Note that with the utility function as described, the agent's goal is to win each of the N auctions at the lowest possible price. If the agent's goal is to win fewer than $M < N$ items, then the same utility function is used, but the agent must constrain his bidding so that he is never the leading bidder in more than M auctions at a time.

We assume that $P(q_t^i = j, \overline{h}_t^i = j | \boldsymbol{x_t})$ is known, in other words the bids of the population and the proxy bids depend on the listed prices of all auctions. To simplify notation, we use

$$\sum_{\boldsymbol{q}=\boldsymbol{0}}^{\boldsymbol{A}} = \sum_{q^1=0}^{A_1} \cdots \sum_{q^N=0}^{A_N}, \qquad \sum_{\boldsymbol{v}=\boldsymbol{0}}^{\boldsymbol{1}} = \sum_{v^1=0}^{1} \cdots \sum_{v^N=0}^{1}, \qquad \sum_{\overline{\boldsymbol{h}}=\boldsymbol{x}}^{\boldsymbol{A}} = \sum_{\overline{h}^1=x^1}^{A_1} \cdots \sum_{\overline{h}^N=x^N}^{A_N}.$$

Bellman's equation is thus given by:

$$\begin{aligned}
J_{T+1}(\boldsymbol{x_{T+1}}, \boldsymbol{h_{T+1}}) &= \mathcal{U}(\boldsymbol{x_{T+1}}, \boldsymbol{h_{T+1}}) \\
J_t(\boldsymbol{x_t}, \boldsymbol{h_t}) &= \max_{\boldsymbol{u_t} \in F_t(\boldsymbol{x_t}, \boldsymbol{w_t})} E_{\boldsymbol{q_t}, \boldsymbol{v_t}, \overline{\boldsymbol{h}}_{\boldsymbol{t}}} [J_{t+1}(\boldsymbol{x_{t+1}}, \boldsymbol{w_{t+1}})], \quad t = 1, \ldots, T, \\
&= \max_{\boldsymbol{u_t} \in F_t(\boldsymbol{x_t}, \boldsymbol{w_t})} \sum_{\boldsymbol{q}=\boldsymbol{0}}^{\boldsymbol{A}} \sum_{\boldsymbol{v}=\boldsymbol{0}}^{\boldsymbol{1}} \sum_{\overline{\boldsymbol{h}}=\boldsymbol{x}}^{\boldsymbol{A}} J_{t+1}(\boldsymbol{f_t}(\boldsymbol{x_t}, \boldsymbol{h_t}, \boldsymbol{u_t}, \boldsymbol{q}, \boldsymbol{v}, \overline{\boldsymbol{h}}), \boldsymbol{g_t}(\boldsymbol{h_t}, \boldsymbol{u_t}, \boldsymbol{q}, \boldsymbol{v}, \overline{\boldsymbol{h}})) \\
&\quad \cdot \prod_{i:\ w_t^i=0} P(q_t^i = q^i, \overline{h}_t^i = \overline{h}^i | \boldsymbol{x_t}) \cdot \prod_{i=1}^{N} P(v_t^i = v^i). \qquad (19)
\end{aligned}$$

Note that $\overline{h}_t^i = 0$, when $w_t^i = 1$, and thus we only take expectations in Eq. (19) over only those $\overline{h}_t^i$ for which $w_t^i = 0$. In practice of course, the computation from Eqs. (19) is barely feasible even for two auctions. Moreover, it is infeasible for three simultaneous auctions given the high dimension of Bellman's equation. For this reason, we propose in the next subsections several approximate dynamic programming methods.

3.1 Approximate dynamic programming method 1

The method we consider in this and the next section belongs in the class of methods of approximate dynamic programming (see Bertsekas and Tsitsiklis [3]). Under this method, abbreviated as $ADP1$, for each of the 2^N binary vectors $\boldsymbol{w_t} \in \{0,1\}^N$ we approximate the cost-to-go function $J_t(\boldsymbol{x_t}, \boldsymbol{h_t})$ as follows:

$$\hat{J}_t(\boldsymbol{x_t}, \boldsymbol{h_t}) = r_0(\boldsymbol{w_t}, t) + \sum_{i=1}^{N} r_i(\boldsymbol{w_t}, t) x_t^i,$$

where each of the coefficients $r_i(\boldsymbol{w_t}, t)$, $i = 0, 1, \ldots, N$ are defined for each of the 2^N vectors $\boldsymbol{w_t}$.

By its nature, this approach works only for up to $N = 5$ auctions. We use simulation to generate feasible states $(\boldsymbol{x_t}, \boldsymbol{h_t})$. The overall algorithm is as follows.

Algorithm $ADP1$:

1. For time period $t = T, \ldots, 1$ and each $\boldsymbol{w} \in \{0,1\}^N$ select by simulation a set $X_t(\boldsymbol{w})$ of states $(\boldsymbol{x_t}(k), \boldsymbol{h_t}(k))$ indexed by k.

2. For each $(\boldsymbol{x_t}(k), \boldsymbol{h_t}(k)) \in X_t(\boldsymbol{w})$ compute

$$\tilde{J}_t(\boldsymbol{x_t}(k), \boldsymbol{h_t}(k)) = \max_{\boldsymbol{u_t} \in F_t(\boldsymbol{x_t}(k), \boldsymbol{h_t}(k))} E[\hat{J}_{t+1}(\boldsymbol{x_{t+1}}, \boldsymbol{h_{t+1}})], \tag{20}$$

where

$$\hat{J}_t(\boldsymbol{x}, \boldsymbol{h}) = r_0(\boldsymbol{w}, t) + \sum_{i=1}^{N} r_i(\boldsymbol{w}, t) x^i.$$

3. For each $\boldsymbol{w} \in \{0,1\}^N$, find parameters $r(\boldsymbol{w}, t)$ by regression, i.e., solving the least squares problem:

$$\sum_{(\boldsymbol{x_t}(k), \boldsymbol{w}) \in X_t(\boldsymbol{w})} \left(\tilde{J}_t(\boldsymbol{x_t}(k), \boldsymbol{h_t}(k)) - r_0(\boldsymbol{w}, t) - \sum_{i=1}^{N} r_i(\boldsymbol{w}, t) x_t^i(k) \right)^2. \tag{21}$$

Notice that the algorithm is still exponential in N as the cost-to-go function for each time t is approximated by 2^N linear functions, each corresponding to a distinct vector $\boldsymbol{w}$.

3.2 Approximate dynamic programming method 2

This method, abbreviated as $ADP2$, is similar to the previous method, but instead of using 2^N linear (in $\boldsymbol{x_t}$) functions to approximate $J_t(\cdot)$ it uses $N + 1$ linear functions. In this method, the cost-to-go-function only depends on $a = \sum_{i=1}^{N} w_t^i$, that is, the number of auctions the agent is the highest bidder at time t. In this method, we only need to evaluate $N+1$ vectors $\boldsymbol{r}(a,t)$, $a = 0,\ldots,N$ and $t = 1,\ldots,T$. Although this uses a coarser approximation than method A, it is capable to solve problems with a larger number of auctions.

3.3 Integer programming approximation

Under this method, abbreviated as IPA, we let $d_t^i(x_t^i, h_t^i, j)$ denote the expected utility of bidding j in auction i given state (x_t^i, h_t^i) and optimally bidding in this single auction thereafter. This is calculated as

$$d_t^i(x_t^i, h_t^i, j) = E_{q_t^i, v_t^i, \overline{h}_t^i}[J_{t+1}^i(f(x_t^i, h_t^i, j, q_t^i, v_t^i, \overline{h}_t^i), g(h_t^i, j, q_t^i, v_t^i, \overline{h}_t^i)], \tag{22}$$

with

$$J_t^i(x_t^i, h_t^i) = \max_j d_t^i(x_t^i, h_t^i, j). \tag{23}$$

Starting with $J_{T+1}^i(x_{T+1}^i, h_{T+1}^i) = \mathfrak{U}(x_{T+1}^i, h_{T+1}^i) = (A_i - x_{T+1}^i)w_{T+1}^i$, we use Eqs. (22) and (23) to find $d_t^i(x_t^i, h_t^i, j)$.

For a fixed time t we define the following decision variables $u_i(j,t)$ as

$$u_i(j,t) = \begin{cases} 1, & \text{if the agent bids at least } j \text{ in auction } i \text{ at time } t, \\ 0, & \text{otherwise.} \end{cases}$$

Given the state $(\boldsymbol{x_t}, \boldsymbol{h_t})$, and constants A, A_i, the agent solves the following discrete optimization problem:

$$\begin{aligned}
\text{maximize} \quad & \sum_{i=1}^{N}\sum_{j=0}^{A_i} u_i(j,t)(d_t^i(x_t^i, h_t^i, j) - d_t^i(x_t^i, h_t^i, j-1)) && && (24)\\
\text{subject to} \quad & u_i(j,t) \le u_i(j-1,t) && \forall i,j && (25)\\
& \sum_{i=1}^{N}\sum_{j=1}^{A_i} u_i(j,t) \le A && && (26)\\
& \sum_{j=1}^{A_i} u_i(j,t) \ge h_t^i && \forall i && (27)\\
& u_i(j,t) \in \{0,1\} && \forall i,j, &&
\end{aligned}$$

where $d_t^i(x_t^i, h_t^i, -1) = 0 \; \forall i, j$. The cost coefficients in (24) represent the marginal increase in utility for bidding one unit higher in a given auction. Note that if the agent bids j_0 in auction i, that is $u_i(j,t) = 1$ for $j \leq j_0$, and $u_i(j,t) = 0$ for $j \geq j_0 + 1$, then the contribution to the objective function (24) is correctly $d_t^i(x_t^i, h_t^i, j_0)$. Constraint (25) ensures that if we bid at least j in auction i, then we had to have bid at least $j-1$ in auction i. Constraint (26) is the way auctions interact, that is through a global budget. Constraint (27) ensures that if the agent is the highest bidder in auction i at time t, that is $h_t^i > 0$, then his bid at time t should be larger than his proxy bid at time $t-1$.

Note that the solution to Problem (24) only provides an approximate solution method as it ignores the budget constraint in future periods. It also does not take into account the possibility that the bids of the population in different auctions might be correlated.

3.4 Pairwise integer programming approximation method 1

In this section, we propose a more elaborate approximation method based on integer programming. Under this method, abbreviated as $PIPA1$, we optimally solve all pairs of auctions using the exact dynamic programming method, and then at each time stage, for a given state of the auctions, find the bid that maximizes the sum of the expected cost-to-go over all pairs of auctions.

Let $M = \{(i,k) | \; i,k = 1,\ldots,N, \;\; i < k\}$ be the set of all $\binom{N}{2}$ pairs of auctions. As before we solve the two auction problem optimally by dynamic programming. This enables us to compute for all pairings (i,k) the quantity $d_t^{(i,k)}(r,s)$, the expected cost to go after bidding r in auction i and s in auction k at time t. Given the optimal cost to go function $J_t(\boldsymbol{x_t}, \boldsymbol{h_t})$ calculated from Eq. (19) for a two auction problem, the quantities $d_t^{(i,k)}(r,s)$ are given by:

$$d_t^{(i,k)}(r,s) = E[J_{t+1}(\boldsymbol{f_t}(\boldsymbol{x_t}, \boldsymbol{h_t}, (r,s), \boldsymbol{q_t}, \boldsymbol{v_t}, \overline{\boldsymbol{h_t}}), \boldsymbol{g_t}(\boldsymbol{h_t}, (r,s), \boldsymbol{q_t}, \boldsymbol{v_t}, \overline{\boldsymbol{h_t}})] \tag{28}$$

We define the decision variable $u_{(i,k)}(r,s,t)$, which is equal to one if the agent bids at least r in auction i and at least s in auction k at time t, and is 0, otherwise. At time t, for a given state $(\boldsymbol{x_t}, \boldsymbol{h_t})$ the agent solves the following discrete optimization problem:

$$
\begin{aligned}
\max \quad & \sum_{(i,k)\in M} \sum_{r=0}^{A_i} \sum_{s=0}^{A_k} u_{(i,k)}(r,s,t)(d_t^{(i,k)}(r,s,t) - d_t^{(i,k)}(r-1,s,t) \\
& \qquad -d_t^{(i,k)}(r,s-1) + d_t^{(i,k)}(r-1,s-1)) && (29)\\
\text{s.t.} \quad & u_{(i,k)}(r,s,t) \le u_{(i,k)}(r-1,s,t) && (30)\\
& u_{i,k}(r,s,t) \le u_m(r,s-1,t) && (31)\\
& u_{(i,k)}(r,s,t) - u_{(i,k)}(r-1,s,t) && \\
& -u_{(i,k)}(r,s-1,t) + u_{(i,k)}(r-1,s-1,t) \ge 0 \qquad \forall (i,k)\in M, \forall r,s && (32)\\
& u_{(i,k)}(r,0,t) - u_{(i,l)}(r,0,t) = 0 \qquad \forall i,k,l,r && (33)\\
& u_{(i,k)}(r,0,t) - u_{(l,i)}(0,r,t) = 0 \qquad \forall i,k,l,r && (34)\\
& u_{(k,i)}(0,r,t) - u_{(l,i)}(0,r,t) = 0 \qquad \forall i,k,l,r && (35)\\
& \sum_{r=1}^{A} u_{(1,2)}(r,0,t) + \sum_{n_2=2}^{N} \sum_{r=1}^{A} u_{(1,n_2)}(0,r,t) \le A, && (36)\\
& \sum_{r=1}^{A_1} u_{(1,2)}(r,0,t) \ge h_t^1 && (37)\\
& \sum_{r=1}^{A_k} u_{(1,k)}(0,r,t) \ge h_t^k && (38)\\
& u_{(i,k)}(r,s,t) \in \{0,1\},
\end{aligned}
$$

with $d_t^{(i,k)}(r,s,t) = 0$ if r or $s = -1$. The optimal bidding vector is

$$
\left(\sum_{r=1}^{A} u_{(1,2)}(r,0,t), \sum_{r=1}^{A} u_{(1,2)}(0,r,t), \ldots, \sum_{r=1}^{A} u_{(1,N)}(0,r,t) \right).
$$

The cost coefficients in (29) represent the marginal increase in utility for bidding one unit higher in both auctions of a given pair. Constraint (30) enforces that if the agent bids at least r in auction i, then he has to bid at least $r-1$. Likewise for Constraint (31). Constraint (32) enforces that if the agent bids at least r in auction i, at least $s-1$ in auction k, and at least $r-1$ in auction i and at least s in auction k, then he has to bid at least r in auction i and at least s in auction k. Constraints (33)-(35) enforce consistent decisions in each auction pairing. Constraint (36) is the global budget constraint. Finally, Constraints (37), (38) ensure that if the agent is the highest bidder in auction k at time t, that is $h_t^k > 0$, then his bid at time t should be larger than his proxy bid at time $t-1$.

3.5 Pairwise integer programming approximation method 2

The computational burden of the pairwise integer programming approximation is considerable as we need to solve $\binom{N}{2}$ pairs of auctions exactly. Alternatively, we can solve $N/2$ disjoint pairs of auctions and combine the cost to go functions in an integer programming problem. We omit the details as they are very similar to what we have already presented. We abbreviate the method as $PIPA2$.

3.6 Empirical results

We consider an agent bidding for an identical item in N multiple auctions for $N = 2, 3, 6$, where the item is valued at A. In this case $A_i = A$. The utility received at the end of the auction is

$$\mathfrak{U}(x_{T+1}, h_{T+1}) = C \sum_{i=1}^{N} (A_i - x_{T+1}^i) w_{T+1}^i. \tag{39}$$

We set $A = A_i = 15$ and $C = 10$ for Palm Pilots III, and $A = A_i = 10$ and $C = 50$ for stamp collections. We use $T = 13$ and $p = 0.8$ and the competing bidding distributions are calculated as in Section 2.

We have implemented all the methods proposed: the exact dynamic programming method for $N = 2$ abbreviated as DP; the approximate dynamic programming methods of Sections 3.1 and 3.2 abbreviated as $ADP1$ and $ADP2$ respectively; the integer programming based methods of Sections 3.3, 3.4 and 3.5 abbreviated as IPA, $PIPA1$ and $PIPA2$ respectively.

Tables 10-12 and 13-15 report simulation results averaged over 10,000 simulations of $N = 2, 3, 6$ simultaneous auctions using eBay data for Palm Pilots III, and stamp collections respectively.

In Table 10 we compare the performance of DP, $ADP1$, $ADP2$ and IPA for $N = 2$ auctions with the goal of giving insight on the degree of suboptimality of the approximate methods compared to the optimal one. Notice that for $N = 3$, solving the exact dynamic programming problem is computationally infeasible. In Table 11 in addition to $ADP1$, $ADP2$ and IPA, we include $PIPA1$ in the comparison. In Table 12, we compare IPA and $PIPA2$ for $N = 6$ auctions. The Column labeled "% Won" is the percentage of auctions that were won, the labeled "% at least one win" is the fraction of rounds (one round is one set of N simultaneous auctions) in which at least one auction was won, and the Column "Avg. Spent per Win" is the amount spent in dollars per auction

Method	% Won	Avg. Utility	Avg. Spent per Win
DP	41.4	39.3	102.5
$ADP1$	39.4	34.8	105.8
$ADP2$	38.4	34.7	104.8
IPA	42.0	39.0	103.6

Table 10: Comparison of DP, $ADP1$, $ADP2$ and IPA for $N = 2$ auctions, $A = 15$, $C = 10$ and data from Palm Pilots III.

Method	% Won	Avg. Utility	Avg. Spent per Win
$ADP1$	29.3	17.6	130
$ADP2$	31.5	9.5	140
IPA	30.3	45.8	99.7
$PIPA1$	29.9	46.6	98.1

Table 11: Comparison of $ADP1$, $ADP2$, IPA, and $PIPA1$ for $N = 3$ auctions, $A = 15$, $C = 10$ and data from Palm Pilots III.

won. If we set our budget $A = A_1 = \cdots = A_6$, then "% at least one win" shows how much more often we win by competing in more auctions than just one. Tables 13-15 have the same comparisons but for stamp collections.

The results in Tables 10-12 and 13-15 suggest the following insights:

(a) The integer programming based methods (IPA, $PIPA1$) clearly outperform the approximate dynamic programming methods ($ADP1$, $ADP2$) (see Tables 10, 11, 13, 14).

(b) When it is computationally feasible to find the optimal policy ($N = 2$), IPA is almost optimal

Method	% Won	% at least one win	Avg. Utility	Avg. Spent per Win
IPA	15.5	93.1	54.9	92.1
$PIPA2$	15.9	94.7	55.4	91.7

Table 12: Comparison of IPA and $PIPA2$ for $N = 6$ auctions, $A = 15$, $C = 10$ and data from Palm Pilots III.

Method	% Won	Avg. Utility	Avg. Spent per Win
DP	90.5	626.4	153.9
$ADP1$	64.0	413.2	177.0
$ADP2$	57.4	364.0	183.0
IPA	90.3	624.3	154.5

Table 13: Comparison of DP, $ADP1$, $ADP2$ and IPA for $N = 2$ auctions, $A = 10$, $C = 50$ and data from stamp collections.

Method	% Won	Avg. Utility	Avg. Spent per Win
$ADP1$	51.9	487.2	187.1
$ADP2$	47.2	429.2	196.9
IPA	64.0	677.0	147.3
$PIPA1$	64.3	686.6	144.3

Table 14: Comparison of $ADP1$, $ADP2$, IPA, and $PIPA1$ for $N = 3$ auctions, $A = 10$, $C = 50$ and data from stamp collections.

Method	% Won	% at least one win	Avg. Utility	Avg. Spent per Win
IPA	34.2	99.3	759.6	129.4
$PIPA2$	34.0	99.7	762.8	126.4

Table 15: Comparison of IPA and $PIPA2$ for $N = 6$ auctions, $A = 10$, $C = 50$ and data from stamp collections.

Method	% Won	% Single Win	% Double Win	% Triple Win	Avg. Utility	Avg. Spent per Win
IPA	54.4	28.1	66.4	0.8	71.6	106.1
$PIPA1$	54.4	28.1	66.4	0.8	71.6	106.1

Table 16: Comparison of IPA and $PIPA1$ for $N = 3$ auctions, $A_1 = A_2 = A_3 = A/2$, $A = 30$, $C = 10$, and Palm Pilots III data.

(see Tables 10, 13). The exact dynamic programming policy leads to slightly higher utility.

(c) The more sophisticated $PIPA1$ (for $N = 3$) leads to slightly better solutions compared to IPA for Palm Pilots III data (see Table 11) and the same solutions for stamp collections data (see Table 14), but at the expense of much higher computational effort.

(d) IPA is outperformed only slightly by $PIPA2$ (see Tables 12, 15). For all its computational effort, $PIPA2$ has slightly greater average utility than IPA.

The emerging insight from all the computational results is that IPA seems an attractive method relative to the other methods. It is certainly significantly faster than all other methods, and its performance is very close to the more sophisticated $PIPA1$.

We next examine the robustness of this conclusion relative to the budget A. In Tables 16 and 17, we consider the case of bidding in $N = 3$ auctions with $A_1 = A_2 = A_3 = A/2$. For Palm Pilots III data we set $A = 30$, $C = 10$ and for stamp collections $A = 20$, $C = 50$. The columns labeled "% Single Win", "% Double Win" and "% Triple Win" are the percentage f_1, f_2, f_3 of simulations in which 1 out of 3, 2 out of 3, and all 3 out of 3 auctions were won, respectively. The column labeled "% Won" is the fraction f of auctions won, i.e., $f = (f_1 + 2f_2 + 3f_3)/3$. Note that the expected utility is equal to the fraction of wins f times N times the difference of $A/2$ and the average spent per win. The results in Tables 16 and 17 show that the performances of IPA and $PIPA1$ are identical. Thus, given that computationally IPA is faster and simpler, IPA is our proposed approach for the problem of multiple simultaneous auctions.

Method	% Won	% Single Win	% Double Win	% Triple Win	Avg. Utility	Avg. Spent per Win
IPA	97.4	0.2	7.3	92.5	983.8	163.4
$PIPA1$	97.4	0.2	7.3	92.5	983.8	163.4

Table 17: Comparison of IPA and $PIPA1$ for $N = 3$ auctions, $A_1 = A_2 = A_3 = A/2$, $A = 20$, $C = 50$, and stamp collections data.

3.7 Bidding against a sophisticated competitor in multiple auctions

With the tremendous volume of trade occurring on eBay, it comes as no surprise that many similar goods are being auctioned off concurrently. As Zheng reports in [21], it is interesting to observe that bidders have taken advantage of this trend by employing the simple heuristic of bidding in the auction with the lowest listed price for a particular item. In this section we examine how IPA performs in a multi-bidder environment while competing in three simultaneous auctions. In addition to competing against bids from the population, we now consider a setting with an additional agent bidding in the same three auctions who has budget A_2 and employs the following strategy: Bid A_2 at time $T-1$ in the auction with the lowest listed price. If outbid then bid A_2 at time T in the auction with the lowest listed price, otherwise do not bid. In this three bidder environment ties between IPA and the competitor are randomly decided, while any tie with the population is won by the population.

Table 18 shows the results of simultaneously bidding for Palm Pilots in three auctions against a population bid and a competitor. Here, 'Win %' is the probability of winning one out of the three auctions. The competitor's utility is his budget minus the price paid if he won. The results indicate that as the competitors budget increases, strategy IPA causes the agent to spend more per auction on average and win less often. This is because the two bidders are often bidding in the same auction, which the agent will win since it has the greater of the two budgets. Note however that when the competitor's budget is equal to the agent's budget, the agent wins less often than in other scenarios, but also spends less. This is because the agent is only winning in auctions that have a low listed price and that the competitor has not bid in. These results indicate that the two strategies are similar.

Competitor's Budget	*IPA* Win %	*IPA* Avg. Utility	*IPA* Avg. Spent per Win	Competitor Win %	Competitor Avg. Utility	Competitor Avg. Spent per Win
100	89.1	38.6	106.7	10.9	1.5	86.2
110	88.8	33.8	111.9	20.4	3.0	95.3
120	84.8	28.0	116.9	31.8	5.6	102.8
130	83.7	23.7	121.6	41.6	8.7	109.1
140	78.6	19.5	125.2	50.9	13.2	114.1
150	45.9	16.8	113.4	87.1	17.4	130.0

Table 18: Performance of bidding against Policy 'Bid budget at T-1 in lowest listed price auction' in 3 auctions, for Palm Pilot IIIs, the agent's budget is 150.

Probability of Competitor's Entrance	*IPA* Win %	*IPA* Avg. Utility	*IPA* Avg. Spent per Win	Competitor Win %	Competitor Avg. Utility	Competitor Avg. Spent per Win
0.0	78.6	19.5	125.2	50.9	13.2	114.1
0.25	77.7	15.6	129.9	35.2	12.5	104.4
0.50	76.1	11.8	134.5	33.3	13.1	100.5
0.75	78.0	11.9	134.7	36.9	14.4	101.0
1.0	85.7	15.8	131.5	54.5	18.8	105.5

Table 19: Performance of bidding against Policy 'Bid budget of 140 at T-1 in lowest listed price auction' in 3 auctions, for different anticipated probabilities of entrance into single auction, for Palm Pilot IIIs, the agent's budget is 150.

Table 19 shows how IPA performs bidding for Palm Pilots in three auctions when IPA was constructed using a combination of the population's bid and the competitor's bid, as in Section 2.4. Note that since IPA solves auctions independently, we assume the entrance probability is the same for each auction and independent of other auctions. In simulations, the competitor is always present and bids in the auction with the lowest listed price at time $T-1$. In this example, the competitor has a budget of 140. For the case when the entrance probability is less than one, IPA performs worse than if it does not know of the competitor. This occurs because IPA has trouble deciding between committing its budget to one auction in order to beat the competitor, and not bidding at all in order to keep the prices low. We notice improvement in IPA's win percentage when the assumed entrance probability is one. These results show that while IPA is able to handle multiple auctions with a large enough budget, the algorithms inability to account for future bidding constraints can cause it to have difficulty in bidding against competitive agents bidding in multiple auctions. These results also demonstrate, however, that incorporating some information of the competitor's presence does increase IPA's performance by certain measures.

4 Multiple overlapping auctions

In this section, we extend our methods to the more general setting of a bidder interested in bidding simultaneously in multiple auctions, not all ending at the same time. The set of auctions we consider is fixed, that is we do not consider prospective auctions which are not already in process. In Bertsimas et. al. [5] we consider the problem of dynamically arriving auctions. Due to the high dimensionality required from an exact dynamic programming based approach, we focus on the integer programming approximation method IPA, as this was the method that gave the best results in the simultaneous auctions case.

Suppose there are currently N auctions currently in process. Let x^i, h^i, t^i be the listed price, proxy bid, and time remaining, respectively, in auction i. Let A be the amount of the budget remaining, and A_i be the amount we are willing to spend in auction i. The state space then becomes $(\boldsymbol{x}, \boldsymbol{h}, \boldsymbol{t}, A) = (x^1, \ldots, x^N, h^1, \ldots, h^N, t^1, \ldots, t^N, A)$. By solving a single auction problem using exact dynamic programming, we calculate the quantities $d^i_{t^i}(x^i_{t^i}, h^i_{t^i}, A, j)$, the expected utility of bidding j, in auction i, with t^i time remaining and a total budget of A to spend. Let t be the current time. We use the decision variables $u_i(j,t)$, which is equal to one if the agent bids at least j in auction i at time t, and zero, otherwise.

The agent solves Problem (25) with a slightly modified objective function as follows:

$$\text{maximize} \quad \sum_{i=1}^{N}\sum_{j=0}^{A_i} u_i(j,t)(d_{ti}^i(x_{ti}^i, h_{ti}^i, A, j) - d_{ti}^i(x_{ti}^i, h_{ti}^i, A, j-1)).$$

This objective accounts for the fact that different auctions need different durations until their completions.

5 Summary and conclusions

We have provided an optimal dynamic programming algorithm for the problem of optimally bidding in a single on-line auction. The proposed algorithm was tested in simulation with real data from eBay, and it clearly outperforms in simulation static widely used strategies. We have also used the proposed algorithm to buy over one hundred stamp collections and a Palm Pilots III at attractive prices. The first author has applied the algorithm for a single item in over one thousand auctions for stamps and stamp collections. While it is difficult to assess scientifically the effects, the first author feels the algorithm contributed to (a) increasing the probability of winning and (b) decreasing by 20% the amount paid per win. We have also provided several approximate algorithms when bidding on multiple simultaneous auctions under a common budget. We have found that a method based on combining the value functions of single auctions found by dynamic programming using an integer programming framework produces high quality solutions fast and reliably. The method also extends to the problem of multiple auctions ending at different times. We feel that this method applies more generally to dynamic programming problems that are weakly coupled.

Acknowledgments

We would like to thank the Associate Editor and the referees for their insightful comments and suggestions that have helped us improve both the exposition as well as the content of this paper. We would also like to thank Nicolas Stier for his help in the collection of the data. This research was partially supported by MIT's E-business center, the Singapore-MIT alliance, and PECASE award DMI-9984339 from the National Science Foundation.

References

[1] Avery, Christopher. 'Strategic Jump Bidding in English Auctions.' *Review of Economic Studies*, 65:185–210, 1998.

[2] Bajari, Patrick and Ali Hortacsu. 'The Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights From eBay Auctions'. Working paper, 2002.

[3] Bertsekas, Dimitri and John Tsitsiklis. *Neuro-Dynamic Programming*, Athena Scientific, Belmont, MA, 1996.

[4] Bapna, Ravi, Paulo Goes and Alok Gupta. 'A Theoretical and Empirical Investigation of Multi-item On-line Auctions.' *Information Technology and Management*, 1(1):1–23, 2000.

[5] Bertsimas, Dimitris, Jeffrey Hawkins and Georgia Perakis. 'Optimal Bidding in Stochastically Arriving Auctions.' Working paper.

[6] Boutilier, Craig, Moises Goldszmidt and Bikash Sabata. 'Continuous Value Function Approximation for Sequential Bidding Policies.' *Proceedings of The 15th International Conference on Uncertainty in Artificial Intelligence (UAI99)*, 1999.

[7] Dewan, Sanjeev and Vernon Hsu. 'Trust in Electronic Markets: Price Discovery in Generalist Versus Specialty On-line Auctions.' Working paper, January 2001.

[8] Hahn, Jungpil. 'The Dynamics of Mass On-line Marketplaces: A Case study of an On-line Auction.' *Proceedings of the CHI 2001 Human Factors in Computing Systems*, 3(1), 2001.

[9] Hasker, Kevin, Raul Gonzalez and Robin C. Sickles. 'An Analysis of Strategic Behavior and Consumer Surplus in eBay Auctions.' Working paper, 2001.

[10] Klemperer, Paul. 'Auction Theory: A guide to the literature.' *Journal of Economic Surveys*, 13:227–286, 1999.

[11] McAfee, R. Preston and John McMillan. 'Auctions and Bidding'. *Journal of Economic Literature,* 25:699–738, 1987.

[12] Milgrom, Paul. 'Auctions and bidding: a primer.' *Journal of Economic Perspectives,* 3:3–22, 1989.

[13] Landsburg, Steven E. (1999). 'My Way to eBay: An economics professor's strategy for winning on-line auctions,' *Slate* , Thursday, April 8, 1999,

[14] Lucking-Reiley, David, Doug Bryan, Naghi Prasad and Daniel Reeves. 'Pennies from eBay: the Determinants of Price in On-line Auctions.' Working paper, 2000.

[15] Mizuta, Hideyuki and Ken Steiglitz. 'Agent-Based Simulation of Dynamic On-Line Auctions.' *Proceedings of the 2000 Winter Simulation Conference,* 2000.

[16] Ockenfels, Axel, and Alvin E. Roth. 'The Timing of Bids in the Internet Auctions: Market Design, Bidder Behavior, and Artificial Agents.' *Artificial Intelligence Magazine*, forthcoming.

[17] Oren, Shmuel S. and Michael H. Rothkopf. 'Optimal Bidding in Sequential Auctions.' *Operations Research*, 23(6):1080–1090, 1975.

[18] Roth, Alvin E., and Axel Ockenfels. 'Last Minute Bidding and the Rules for Ending Second-Price Auctions: Theory and Evidence from a Natural Experiment on the Internet'. Working paper, 2000.

[19] Stone, Peter and Amy Greenwald. 'The First International Trading Agent Competition: Autonomous Bidding Agents'. *Electronic Commerce Research: Special Issue on Dynamic Pricing,* forthcoming.

[20] Vickrey, William. 'Counterspeculation, Auctions, and Competitive Sealed Tenders.' *Journal of Finance,* 16:8–37, 1961.

[21] Zheng, Mingli. 'Bidding Behavior with Competing Auctions: Evidence from eBay.' Working paper.