Computing Reserve Prices and Identifying the Value Distribution in Real-world Auctions with Market Disruptions

William E. Walsh
CombineNet, Inc.
Fifteen 27th St.
Pittsburgh, PA, USA
wwalsh@combinenet.com

David C. Parkes
SEAS
Harvard University
Cambridge, MA, USA
parkes@eecs.harvard.edu

Tuomas Sandholm
Computer Science Dept.
Carnegie Mellon University
Pittsburgh, PA, USA
sandholm@cs.cmu.edu

Craig Boutilier
Dept. of Computer Science
University of Toronto
Toronto ON, Canada
cebly@cs.toronto.edu

Abstract

Single-good ascending auctions, including the English Auction and its close variants (e.g. eBay), are the most widely used type of auction. Effective strategies for buyers and sellers in these auctions can have an enormous economic impact. We present a system that, relying on minimal assumptions, computes reserve prices for real-world ascending auctions by inferring the value distribution from past auction data. Our contributions include improvements on previous methods for estimating the bidder value distribution, and novel Bayesian methods for adapting to disruptions in market conditions. We demonstrate the effectiveness of our system in both simulated auctions and in real Internet auctions that we conducted in the domain of selling off returned merchandise.

Introduction

Single-good ascending auctions, including the English Auction and its close variants (e.g. eBay), are the most widely used type of auction. Traditionally used to sell art and livestock, ascending auctions have become widespread on the Internet in the past decade. Given the prevalence of ascending auctions, effective strategies for buyers and sellers can have an enormous economic impact. For sellers, an appropriate reserve price is essential for maximizing profit, as it always determines the auction price when there is one bidder, and it can determine the price even if there are many bidders. Under certain conditions—namely, in isolated auctions where bidders have static, independent, identically distributed, private values—it is well-understood how to compute the optimal reserve price as long as the distribution of bidders' valuations is known. However, these assumptions rarely hold in practice. First, the value distributions are not known *ex ante*, although, fortunately, for many goods we can observe the bids from previous auctions. Still, estimating the value distribution is non-obvious and non-trivial because the distribution of bids is not the same as distribution of values [3, 5]. Second, thanks in no small part to their proliferation on the Internet, auctions rarely run in isolation. Except for rare and unique items, there are often multiple channels, both auctions and fixed-price, for purchasing a good [6]. Third, value distributions are typically non-stationary. Buyer values may exhibit seasonality effects (e.g. higher before Christmas and lower afterward) and may change according to a variety of other factors, including product obsolescence (e.g. for electronics and fashion) and economic downturns.

Others have addressed the issue of estimating value distributions from observed bids. Jiang and Leyton-Brown [5] used an expectation maximization algorithm to address the issue of how "hidden bids" in online auctions can skew the bidding distribution away from the underlying value distribution. Although they focused on using the value distributions to inform bidding strategies, their method could also be used to compute reserve prices. Their approach was able to effectively infer the value distribution when given the parameterized form. Haile and Tamer [3] developed a method for inferring bounds on the distribution and the optimal reserve price, making no *a priori* assumptions about the bidder valuations and very minimal assumptions about bidder behavior. However, their approach was not complete, as it did not provide a way to choose the reserve price within the bounds.[1] Additionally, neither of the aforementioned approaches addressed the issue of non-isolated auctions or non-stationary distributions. Addressing these issues for real-world auctions is a significant challenge, as the interconnectedness of information and commerce in the Internet age suggests an expansive and complex model. But performing a full game-theoretic analysis in setting reserve prices in this dynamic context is certainly prohibitive.[2]

In this paper we present an automated methodology and system for computing reserve prices for ascending auctions with heterogeneous, but related goods, and describe the results of a live test of the system in real Internet auctions. Our basic approach is to estimate a value distribution, adjust it to account for competing channels and non-stationarities, and compute a reserve price from the distribution. Throughout, we make minimal modeling assumptions and use only that information which is readily available from past auctions.

Our initial system is based on the approach of Haile and Tamer [3], but with the addition of our own technique for computing a specific distribution within their bounds. We

[1] Haile and Tamer compute *bounds* on the optimal reserve price in an *ex post* analysis of data from timber-harvesting auctions, but do not estimate a specific optimal reserve price, nor do they actually test their approach for live auctions.

[2] Juda and Parkes [6] address multiple auctions in the context of eBay for the purpose of value distribution identification, generalizing Haile and Tamer somewhat to allow for bidders that participate in multiple auctions. However, their focus is not on setting reserve prices, and they do not consider problems related to non-stationary distributions.

adjust the distribution to account for the historically observed influence of known simultaneous auctions and seasonal variance in bidding. We fielded our system in a live two-month trial on real auctions, which demonstrated the effectiveness of our approach, but also revealed its inability to adapt to a drop in bidding coinciding with the appearance of a competitor. In response, we developed a new Bayesian technique for adjusting to sudden shocks in the market without explicitly modeling their cause.

The Core Method

Consider an isolated English auction in which bidders have independent private values for a single good. A bidder's value is drawn from distribution $f_0(v)$ with cumulative density function $F_0(v)$, bounded in the range $[\underline{v}, \overline{v}]$. The seller places value v_0 on the good, which can represent her value for keeping it, or, alternatively, the price for which she can sell it through another channel. A bidder may place a bid for any price at least as high as the *price quote*, which is the maximum of the reserve price r and the highest standing bid, plus a bid increment. The auction continues until no bidder increases their bid, in which case the highest bidder wins at his bid price.

Haile and Tamer's Bounds In this section we briefly review the approach of Haile and Tamer [3] for computing bounds on the value distribution. We refer the reader to their work for further details and proofs of correctness.

An appropriately set reserve price can improve revenue for the seller, as it determines the price when there is only one bidder.[3] If $(p-v_0)[1-F_0(p)]$ is strictly pseudo-concave on $[\underline{v}, \overline{v}]$, then the optimal reserve is given by

$$\arg\max_p (p - v_0)[1 - F_0(p)]. \qquad (1)$$

Note that this is also the optimal reserve price for the second-price auction. As such, it is the optimal reserve for eBay-style auctions, under the assumptions above.[4]

Note that, if past auctions had a reserve price r^0, then nothing about $F_0(v)$ is revealed below that point. Hence, Haile and Tamer consider $F(v) = [F_0(v) - F_0(r^0)][1 - F_0(r^0)]$, the distribution $F_0(v)$ truncated at r^0, as the primitive of interest. However, they show that if the past reserve price was below the optimum, this is sufficient to extract bounds on the optimal reserve price.

They construct upper and lower bounds, F_U and F_L, such that, for all v, $F_L(v) \leq F(v) \leq F_U(v)$, and prove that they are theoretically correct. They make no assumptions about the form of the distribution and make only minimal assumptions about bidding behavior, namely 1) bidders do not bid more then they are willing to pay, and 2) bidders do not allow an opponent to win at a price they are willing to beat. In particular, unlike previous nonparametric approaches (e.g., [7, 2, 1]), equilibrium behavior is *not* assumed.

For a given v and a history of auctions, $F_L(v)$ and $F_U(v)$ can be estimated using relatively straightforward averaging and binary search techniques. Note that, although the bounds are theoretically correct for an infinitely large history, the estimated bounds may be in error with small sample sizes.[5]

For English auctions, the method for computing bounds considers only the highest bids.[6] One important point to note is that the method for computing the bounds uses information only from auctions with at least two bids. The reason is that, for only one bidder we can infer only that her value is above the reserve, and if there are no bidders, we can infer only that no bidder happened to have a value as high as the reserve. Haile and Tamer [3] explain why these facts are not useful in computing bounds. This limitation will be of importance when we must estimate the distributions when the past reserve prices were too high.

Haile and Tamer suggest methods for estimating the bounds from past auctions for related, but heterogeneous, goods. If we wish to compute the reserve price for a good x, we use a weighted average of past auction data when computing F_U and F_L. Given a past auction for good x', we weight the data from that auction based on the *similarity* of x' to x. One possible similarity measure, as suggested by Haile and Tamer, is a product of Gaussian kernels [4] on a vector of relevant attributes.

Computing a Specific Distribution Although Haile and Tamer described how to compute bounds on the value distribution, they did not describe how to select a particular distribution within the bounds. Our first contribution is a method for estimating a particular distribution $\hat{F}(v)$ within the bounds. Our system then uses $\hat{F}(v)$ in place of $F_0(v)$ in Eq. 1 to compute a reserve price.[7]

We first compute $F_L(v)$ and $F_U(v)$ for a fine mesh of discrete values of v, according to Haile and Tamer's method.[8] We then pick a particular distribution $\hat{F}(v) = F_L(v) + \alpha(F_U(v) - F_L(v))$, $\alpha \in [0,1]$ by simulating the revenue and finding a best fit to the actual revenue R. To estimate the revenue $\hat{R}_\alpha$ for a candidate α, we first compute the empirical distribution $\tilde{N}$ of the number of bidders in an auction, according to the numbers of bidders actually observed. Then we perform a number of simulations, and compute $\hat{R}_\alpha$ as the average revenue from the simulations. For a given simulation, we sample a number n from $\tilde{N}$, and then sample n values v_i from $\hat{F}(v)$. We evaluate the revenue for the simulation as the second highest of the values v_i. In order to avoid the possible distortion of reserve prices, we use data

[3] It can also increase revenue if there is a large difference between the values of the highest and second highest bidders.

[4] eBay auctions are English auctions with fixed end times, and often bidders wait until the last minute to bid. But, because eBay has a proxy bidding system that automatically bids up to a bidder's stated value, but no more than necessary to win, an auction where bidders submit their valuations to the proxy agents is effectively a second-price sealed-bid auction.

[5] The technique makes uses of a relationship between the order statistics of the bids and $F(v)$. Assumption (1) is used to derive F_U and assumption (2) is used to derive F_L.

[6] If the auctions are eBay-style auctions with proxies, we would use only the publicly available *revealed* bids, not the value reported to the proxy.

[7] Our estimated distribution could also be used to inform bidding strategies, as in [5].

[8] We also developed some technical improvements for handling sparse data that we report in the full paper.

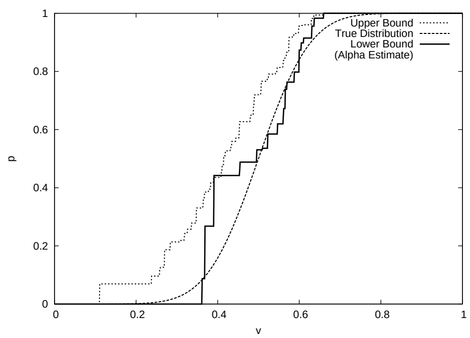

Figure 1: Estimated distribution based on 25 auctions.

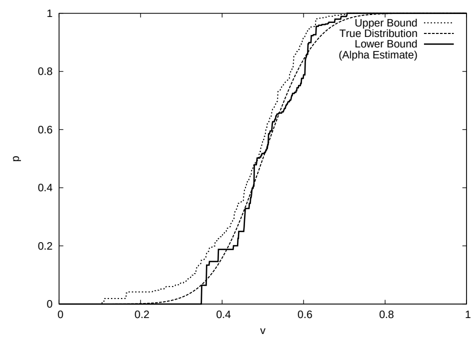

Figure 2: Estimated distribution based on 100 auctions.

only from auctions with two or more bidders in computing $\tilde{N}$ and R.

To determine the correct α, we exploit the fact that $\hat{R}_\alpha$ must be non-decreasing in α, and perform a binary search within $[0, 1]$ until $|\hat{R}_\alpha - R|$ is sufficiently small.

Verification on Simulated Data To illustrate our method, we show the results of applying it to simulated auction data. We simulated T English auctions in which bidders had valuations drawn from from a Gaussian with $\mu = 0.5$, $\sigma = 0.1$, but bounded on $[0, 1]$. Each auction had a number n of potential bidders drawn from a Gaussian with $\mu = 6$, $\sigma = 1$, but with at least one bidder. The auctions had a reserve price of 0.1 and a bid increment of 0.001. As with many Internet auctions, we assume that bidders use proxies. However, as is often observed in Internet auctions, bidders do not necessarily report their true value to the proxy initially. If a bidder with value v is not currently winning, and the current price quote is q, then the bidder randomly reports a value uniformly on $[q, v]$. We simulate biddings in rounds, with the next bidder chosen randomly from the losing bidders with values above q.

As shown in Fig. 1, $\hat{F}(v)$ roughly tracks the true distribution for $T = 25$, while Fig. 2 shows that $\hat{F}(v)$ is an excellent estimation of $F(v)$ for $T = 100$. In the figures $\hat{F}(v)$ is labeled as "Alpha Distribution" and is at the lower bound $F_L(v)$. A visual inspection confirms that, in both cases, $F_L(v)$ is the closest to the true distribution, within the family of distributions characterized by $F_L(v)+\alpha(F_U(v)-F_L(v))$, $\alpha \in [0, 1]$. For $T = 25$ the computed reserve price is below the true optimum by 15.6%. The estimate improves significantly for $T = 50$, giving a reserve price that is above the optimum by only 2.1%. This error of the estimate improves slightly to 2.0% for higher T.

A Live Test of Reserve Pricing

We developed an automated reserve pricing system that uses the method described above and tested it in a two-month live trial. During the trial, the system was used by a company that auctions returned goods in bulk to wholesalers on the Internet. Each auctioned good was a bundle of a variety of electronics in various conditions. 176 auctions were conducted in the trial, producing over $270,000 in revenue.

Modeling the Value Distributions A significant challenge to computing reserve prices was that no two auctions sold the exact same bundle of items. We handled this heterogeneity with two methods. First, we performed all computations in terms of *recovery*, that is the dollar value of a bid normalized by the bundle's wholesale cost. Historically, all bundles sold for a recovery of less than 1, allowing us to bound the search for $\hat{F}(v)$ and the reserve price. Second, we identified various attributes that affected price and weighted past data according to a product of Gaussian kernels on the attributes.

The company was already aware of a strong seasonality effect on auction prices. A nonlinear regression confirmed that final prices were higher in the months before Christmas, when people purchase for the holidays, and lower in January and early February, when people return goods after the holidays. Thus we used day of the year as the attribute in one of the Gaussian kernels.

Ideally, we would have used a detailed description of the composition of each bundle as a basis for further attributes in the kernels. However, the descriptions were too nonstandard for this to be automatizable effectively. Instead we identified easy-to-compute proxies that signaled the relative value of different bundles. Based on a non-linear regression analysis, we found that the final prices tended to increase with the the total wholesale cost of the bundle and also with the average wholesale cost of individual items in the bundle. We used both of these measures as attributes in the product of Gaussian kernels.

Another issue was that the company ran multiple auctions simultaneously, and the goods in simultaneous auctions could be considered (partial) substitutes to some of the bidders. A complete model of valuations and the optimal reserve price would account for these substitutes, but we believe that such a model would be prohibitively complex. Instead, we limited our focus to the isolated auction model, but with the additional narrow assumption that the number of simultaneous auctions would tend to affect the revenue in the auction. In fact, a regression showed that revenue tended to decrease with the number of simultaneous auctions. Thus we used the number of simultaneous auctions as the attribute in a Gaussian kernel.[9]

[9]There is a subtle point in how we are addressing simultaneous auctions. The existence of simultaneous auctions for substitute goods does not actually decrease a bidders' underlying *value* for the goods. Rather, it will tend to decrease the resulting *bids* for the goods because there are multiple options. However, rather than modeling the substitutabilities and the resulting effect on bidding, we instead pretend as if the presence of simultaneous auctions actu-

Note that we did not encode the actual trend of valuations with respect to the attributes. That is, the system did not know the actual effects on value due to seasonality, cost, average cost, or number of simultaneous auctions *a priori*. Rather, when computing the reserve price for a good with particular attributes, the system simply weighted the data from a past auction more or less heavily based on the similarity of the past auction attributes to the current one (according to a product of Gaussian kernels). The actual effect of the attributes on the value distribution was automatically inferred by the system.

The Auction Setup The auctions were eBay-style auctions—that is English auctions with fixed end dates and proxy bidding. The auctions were run by the company on it's web site, and were accessible only to registered buyers with salvage licenses. The auctions were started on Mondays, Wednesdays, and Fridays, and lasted for two days each. On average, 6.5 auctions were run at the same time.

Before our trial, the company always set the reserve price to a recovery of 0.1 (i.e., 10% of the retail price), which was clearly suboptimal since they were able to sell the goods for an average of 0.18 through a separate, fixed-price channel. Because the company had historically set the reserve price so low, the auctions had always received bids above the reserve. Since our system would compute significantly higher reserve prices we had to contend with unsold goods. We decided with the company that, if a good didn't sell at auction, it would be reposted one more time to auction at the same reserve price. If it did not sell in that second auction either, it would be combined with less desirable items into a much larger bundle and sold through a fixed-price channel to one of a small set of select buyers (in order to minimize the effect that a bidder in the auction would feel that he could buy the same bundle through a secondary channel if the item did not sell in the auction). Since we knew that historically goods could be sold through the fixed-price channel for an average recovery of 0.18, we took that to be our v_0.

Since we allowed unsold goods to be reposted, a higher reserve price than specified by Eq. 1 would be optimal. We first computed an initial value using Eq. 1, then searched for the optimum at higher prices. To estimate the revenue that would be obtained at a candidate reserve r, we performed simulations in the same manner as we did to compute $\hat{F}(v)$, but with the reserve set to r and with the auction rerun one more time, with a new set of bidders, if the reserve was not met. In our live trials we used one year of past auction data as the historical data. We did not update that data set during the trial. We performed 2000 simulations each in the computation of $\hat{F}(v)$ and the reserve price adjusted for reposting. For each Gaussian kernel, we used a bandwidth of 0.3 times the standard deviation of the attribute. In evaluating the effectiveness of our system, we compared the recovery obtained using reserve prices from our system with the recovery obtained during the same time period the previous year.[10]

ally changes the underlying values for an individual good, and that bidding proceeds as if the auction is run in isolation. Although this is admittedly an approximation, the success of our live trial bears out its effectiveness.

Results from the Trial Our system was quite effective during the first month of the trial. The recovery increased by 5.6%, from 0.303 to 0.32, as compared to the same period during the previous year.[11] This improvement includes the price of goods sold at auction as well as unsold goods later sold at a fixed price averaging 0.18. The reserve price varied, with different attributes of the goods, in the range $[0.26, 0.33]$. The recovery of goods successfully sold at auction was in the range $[0.285, 0.433]$ and 10% of the auctions did not receive bids exceeding the reserve. The experimental results also agreed well with our simulation, which suggested a recovery improvement of 4.8% and that 9.2% would not receive bids above the reserve.

But something nefarious occurred during the second month of the trial: the performance of our system dropped dramatically, resulting in a decrease in recovery of 19%, as compared to the previous year. During this period, 76% of the auctions did not receive bids exceeding the reserve price. Investigations by the company revealed that a competing seller entered the market during the second month, and we believe that the increased competition resulted in the drop in bid prices. To test this, we turned off the reserve pricing system after the second month and gathered data for another month. The recovery during the third month decreased by 13% as compared to the same period during the previous year. We conclude that the 19% decrease in the second month was due largely to entry of the competitor, with some portion of the decrease due to a reserve price that was too high in light of the new market environment.

Adapting to Market Disruptions

Our deployed system failed to adapt to the changing market conditions in the live trial because it did not incorporate the data from auctions run during the trial (as explained above, it only used the historical data set as the basis for setting reserve prices). A straightforward solution would have been to simply include auction results generated during the trial, perhaps with a heavier weighting on more recent data. Although this would allow our system to adapt if bids had increased, it could not not have been effective for decreasing bids. The problem is that the core system makes use of data only from auctions with at least two bids. When the reserve price is too high, as was most surely the case in the second month, the only data available is from the highest value bidders. As a result, when bids decrease, the estimate of $\hat{F}(v)$, and hence the reserve price, is increasingly skewed upward. If our system had incorporated the new auction data in a straightforward fashion, it would likely have performed worse.

We need a method for inferring something about the por-

[10] We compared recovery, rather than actual dollar revenue, because the wholesale cost of the goods sold changed during this time.

[11] This amount of improvement was good considering that the returned electronics auctions had historically been quite competitive. We would expect an even greater improvement for other types of goods that tend to get fewer bidders.

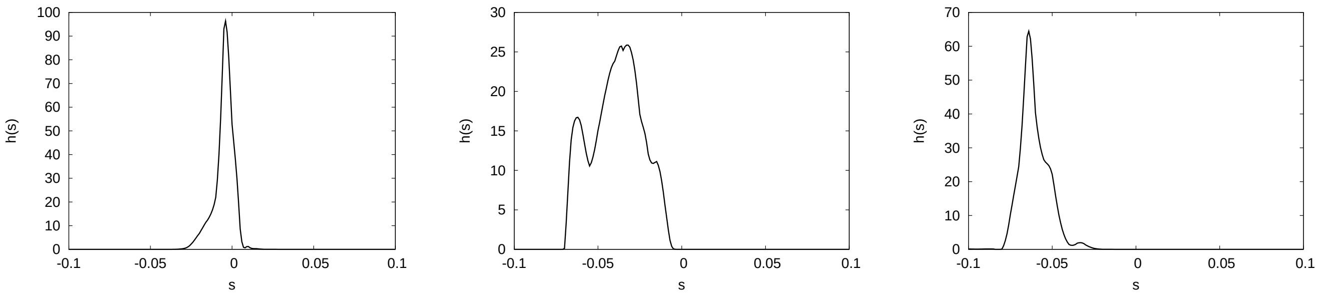

(a) (b) (c)

Figure 3: Distribution $h(s)$ of shift s in the trial as collected: (a) at the end of the first month, (b) in the middle of the second month, (c) at the end of the second month.

tion of the distribution that is hidden by the reserve price.[12]

Modeling Market Changes Changes in observed bidding behavior can potentially arise from a variety of underlying causes. Because of the complexity of modeling multiple effects in the world market, we choose a simplified, unified model whereby we assume changes in bidding reflect fundamental changes in the value distribution. Our approach is to model changes in bidding as a *shift* in the underlying value distribution. That is, a drop in bidding is indicative of a leftward shift of the value distribution while an increase in bidding is indicative of a rightward shift. To estimate the amount of shift, we look at a window of w previous auctions, and assume that a shift may have occurred immediately before the auctions were run. Using the predicted probability of an auction receiving no bids above the reserve, and the actual history of auctions receiving no bids, we compute the posterior distribution of the amount of shift in the value distribution. We then shift the value distribution by the estimated amount and compute the reserve price from the shifted distribution.

Computing the Distribution of Shift Assume for now that all auctions have the same attributes. Let $\hat{F}(v)$ be our current estimate of the cumulative density function of values, as determined with our initial method (i.e., not taking into account any shifting). We assume that the true distribution $F(v)$ is equal to $\hat{F}_s(v)$, the distribution $\hat{F}(v)$ shifted by $s \in \Re$. We use $s > 0$ to indicate a rightward shift, and $s < 0$ to indicate a leftward shift. Let $\delta(s)$ be the event that the shift is s and let $g(s)$ be the prior probability density function of $\delta(s)$. Let r_i be the reserve price for past auction $i \in w$, and let $\pi_{i,s}$ be the probability that i would receive no bids above r_i, given s. For a given s, it is straightforward to determine $\hat{F}_s(v)$, and, as indicated earlier, we can compute $\pi_{i,s}$ via simulation using $\hat{F}_{i,s}(v)$ and the observed distribution of the number of bidders. Let $\gamma_{i,s}$ indicate the probability of the actual outcome of auction i, given r_i and s. That is, if i received at least one bid, then $\gamma_{i,s} = 1 - \pi_{i,s}$, otherwise $\gamma_{i,s} = \pi_{i,s}$. To compute the posterior distribution $h(s)$ of s, we use Bayes' rule:

$$h(s) = \frac{\Pi_{i \in [1,w]} \gamma_{i,s} g(s)}{\int_{-\infty}^{\infty} \Pi_{i \in [1,w]} \gamma_{i,s} g(s) ds}. \quad (2)$$

It is not immediately clear how to estimate $g(s)$, given that it is a prior on the error of our estimate. One possibility would be to substitute the prior distribution over the amount of shift that occurs in the actual market during any given short period of time. For the present work, we simply assume an uninformative prior and let the data fully determine $h(s)$. To compute $h(s)$, we first compute each $\gamma_{i,s}$ for discrete values of s within a reasonable range. In our real-world trial, we normalized values into $[0, 1]$, so we can safely assume that $s \in [-1, 1]$. Then, for any desired value of $s \in [-1, 1]$, we compute $h(s)$ numerically using Romberg integration and polynomial interpolation on $\gamma_i(s)$ [8].

Our goal is to estimate the optimal reserve price for a new auction. We can compute $\hat{F}(v)$ using the initial method and $h(s)$ given the Bayesian method above. Ideally, we would compute the reserve r^* that maximizes the expected revenue. Unfortunately, this would require estimating the revenue for a large number of (r, s) tuples, which could be quite expensive.[13] A less expensive approach would be to compute from $h(s)$ the expected shift $\bar{s}$, shift $\hat{F}(v)$ by $\bar{s}$, and then compute the optimal reserve price from the shifted distribution. This would give a particularly good approximation when $h(s)$ is unimodal with small variance. If auctions have different attributes, then we assume that the goods are similar enough that the amount of the shift is the same for all goods. For each past auction i we compute a separate $\hat{F}_{i,s}(v)$ based on its attributes, and use that to determine $\gamma_{i,s}$. Then we compute $\hat{F}(v)$ for the good we are auctioning presently, shift it by $\bar{s}$, and compute the reserve price.

Adaptation Experiment To demonstrate our adaptation method, we applied it to the data obtained in our live test. At three points in time during the trial, we computed the unshifted value distribution $\hat{F}(v)$, using the attributes of one auction run at that time, and the optimal reserve, r, given $\hat{F}(v)$. Using a window of 30 auctions we computed $h(s)$, $\bar{s}$, $\hat{F}_{\bar{s}}(v)$ (the distribution $\hat{F}(v)$ shifted by $\bar{s}$) and the optimal reserve price $r_{\bar{s}}$ for $\hat{F}_{\bar{s}}(v)$. We compute a product of

[12] Jiang and Leyton-Brown's [5] addressed the problem of how to infer the distribution when low-value bidders do not bid because the auction price has already exceeded their value (because high-value bidders bid first). However, their method, like our first method, ignores the potentially valuable information available from auctions that receive no bids at all.

[13] For every s in our domain, and every reasonable r, we would have to estimate the revenue, via simulation.

Day	r	$\bar{s}$	$r_{\bar{s}}$	Day	r	$\bar{s}$	$r_{\bar{s}}$
4	0.35	0.38	0.73	30	0.48	−0.09	0.39
6	0.35	0.06	0.41	32	0.44	−0.14	0.37
8	0.35	0.22	0.57	34	0.52	−0.18	0.34
10	0.48	−0.09	0.39	36	0.37	−0.18	0.36
12	0.50	−0.08	0.43	38	0.36	−0.19	0.37
14	0.51	−0.06	0.45	40	0.36	−0.31	0.35
16	0.48	−0.06	0.42	42	0.39	−0.16	0.23
18	0.63	−0.05	0.58	44	0.40	−0.13	0.29
20	0.48	−0.08	0.42	46	0.31	−0.08	0.31
22	0.51	−0.12	0.47	48	0.31	−0.08	0.32
24	0.59	−0.11	0.47	50	0.31	−0.05	0.30
26	0.48	−0.11	0.38	52	0.31	−0.06	0.27
28	0.47	−0.09	0.41	54	0.33	−0.06	0.28

Table 1: Results from a simulation of market adaptation. The true μ is 0.5, for the results on the left. The true μ dropped to 0.3 for the results on the right.

Gaussian kernels using attributes as in the live test. Unlike in our live test, we now use all data obtained up to a given date. Figs. 3 (a)–(c) show that $h(s)$ does indicate a shift in the distribution during the second month of the trial. Furthermore $r_{\bar{s}}$ decreases from r in response to the shift. At the end of the first month, $r = 0.30$, $\bar{s} = -0.005$, and $r_{\bar{s}} = 0.30$. In the middle of the second month, $r = 0.25$, $\bar{s} = -0.04$, and $r_{\bar{s}} = 0.24$. Finally, at the end of the second month, $r = 0.23$, $\bar{s} = -0.061$, and $r_{\bar{s}} = 0.23$. (Note that r is different at each point because the auction attributes are different.)

To investigate the interaction of the augmented system with auctions and market changes, we tested it on simulated data. We ran a sequence of simulated auctions in the manner described in the "Verification of Simulated Data" section. We first ran 100 auctions with a reserve price of 0.1. Thereafter, every two simulated days we computed r, then $\bar{s}$ and $r_{\bar{s}}$ using a window of 30 auctions, then ran six auctions with the reserve set to $r_{\bar{s}}$. We used all past data to compute $\hat{F}(v)$, and used recency as an attribute, with a bandwidth of 0.3 for the Gaussian kernel. Bidders initially had values drawn from a Gaussian with $\mu = 0.5$, $\sigma = 0.1$ and bounded in $[0, 1]$, but after 29 days of simulation we decreased μ to 0.3. The optimal reserve for $\mu = 0.5$ is 0.43, and the optimal reserve for $\mu = 0.3$ is 0.30. The results are shown in Table 1.

The first estimate is done based on the 100 auctions with a low reserve price. We see that the system initially, and incorrectly, estimates a large positive shift. Subsequently, after incorporating data from auctions run with reserve prices determined by the system, $|\bar{s}|$ is small and $r_{\bar{s}}$ is significantly more accurate. Since recency is used as an attribute, $\hat{F}(v)$ does tend to drift during the first 30 days, as does the unshifted estimate, as indicated by r. However, the estimate of $\bar{s}$ corrects this drift, giving good estimates for $r_{\bar{s}}$ overall. The change in the true values occurs immediately before day 30, so the estimate on that day does not include auction data reflecting the market shift. Later, estimates of $\bar{s}$ and $r_{\bar{s}}$ decrease appropriately as the new data is incorporated. By the end of the simulation, the estimate of $\bar{s}$ is -0.06. This is not the actual shift in the distribution, but by this point, the reduced reserve price has allowed the system to collect bids from lower-value bidders. As a result, $\hat{F}(v)$ is actually more accurate than immediately after the market change, thus less shift needs to be applied to make it accurate.

Day	r	Day	r	Day	r	Day	r
4	0.35	16	0.51	30	0.56	42	0.48
6	0.35	18	0.52	32	0.48	44	0.48
8	0.35	20	0.51	34	0.55	46	0.48
10	0.35	22	0.51	36	0.51	48	0.48
12	0.50	24	0.51	38	0.61	50	0.48
14	0.50	26	0.51	40	0.61	52	0.48
		28	0.51			54	0.48

Table 2: Results from a simulation without adaptation. The true μ is 0.5 for the results on the left of the double bar, and the true μ dropped to 0.3 for the results to the right of the double bar.

We also simulated our reserve pricing system, using the same sequence of bidders as before, but without the adaptation method. As shown in Table 2, the estimate of the reserve price drifts upward before the shock and remains too high after the shock. The average revenue obtained (excluding the first 100 days) is 0.32. In contrast, with the adaptation method we get a 22% higher revenue of 0.38.

Conclusions

We presented a methodology and a complete system for estimating bidder value distributions and computing reserve prices for ascending auctions with heterogeneous, but related goods, using minimal assumptions about bidder behavior and broader market interactions. The system infers the distributions from the history of past auctions and can adapt to market disruptions, including non-stationary distributions and changes in market competition. We demonstrated the effectiveness of the system in simulations and in a real-world trial of Internet auctions. The live trial of our initial system showed an improvement of 5.6% in revenue before a change in market competition. We augmented the approach with a Bayesian inference approach to determining distributional shifts within the marketplace and demonstrated its efficacy in simulation.

Automated reserve pricing is emerging as an approach that is used in (Internet) auctions in a variety of industries. Our methodology provides a scalable way to compute appropriate reserve prices in complex settings while using only inputs that are readily available in practice.

References

[1] Susan Athey and Philip A. Haile. Identification of standard auction models. *Econometrica*, 70:2107–2104, 2002.

[2] Emmanuel Guerre, Isabelle M. Perrigne, and Quang Vuong. Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68:525–574, 2000.

[3] Philip A. Haile and Elie Tamer. Inference with an incomplete model of English auctions. *J. of Political Economy*, 111(1):1–51, 2003.

[4] W. Härdle. *Applied Nonparametric Regression*. Cambridge University Press, 1990.

[5] Albert Xin Jiang and Kevin Leyton-Brown. Bidding agents for online auctions with hidden bids. *Machine Learning J.*, 67(1-2):117–143, 2007.

[6] Adam I. Juda and David C. Parkes. The sequential auction problem on eBay: An empirical analysis and a solution. In *ACM EC*, pages 180–189, 2006.

[7] Tong Li, Isabelle M. Perrigne, and Quang Vuong. Structural estimation of the afficliated private value auction model with an application to ocs auctions. *Rand J. of Economics*, 33:171–193, 2002.

[8] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes in C: Second Edition*. Cambridge University Press, 1992.