

Learning Environmental Parameters for the Design of Optimal English Auctions with Discrete Bid Levels

A. Rogers¹, E. David¹, J. Schiff², S. Kraus³, and N.R. Jennings¹

¹ Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK

² Department of Mathematics, Bar-Ilan University, Ramat-Gan 52900, Israel

³ Department of Computer Science, Bar-Ilan University, Ramat-Gan 52900, Israel
{acr, ed, nrj}@ecs.soton.ac.uk, {schiff@math, sarit@cs}.biu.ac.il

Abstract. In this paper we consider the optimal design of English auctions with discrete bid levels. Such auctions are widely used in online internet settings and our aim is to automate their configuration in order that they generate the maximum revenue for the auctioneer. Specifically, we address the problem of estimating the values of the parameters necessary to perform this optimal auction design by observing the bidding in previous auctions. To this end, we derive a general expression that relates the expected revenue of the auction when discrete bid levels are implemented, but the number of participating bidders is unknown. We then use this result to show that the characteristics of these optimal bid levels are highly dependent on the expected number of bidders and on their valuation distribution. Finally, we derive and demonstrate an online algorithm based on Bayesian machine learning, that allows these unknown parameters to be estimated through observations of the closing price of previous auctions. We show experimentally that this algorithm converges rapidly toward the true parameter values and, in comparison with an auction using the more commonly implemented fixed bid increment, results in an increase in auction revenue.

1 Introduction

The popularity of online internet auctions has increased dramatically over recent years, with total online auction sales currently exceeding \$30 billion annually. This popularity has prompted much research into agent mediated auctions and specifically the development of autonomous software agents that are capable of fulfilling the role of auctioneer or bidder on behalf of their owner. Now, much of the theoretical work on these agent mediated auctions has focused on direct sealed bid protocols, such as the second-price (Vickrey) auction. These protocols are attractive as they are economically efficient and provide simple dominant bidding strategies for participating agents. However, despite these properties, such sealed bid protocols are rarely used in practice [14]. The vast majority of current online and real world auctions implement variants of a single auction protocol, specifically, the oral ascending price (English) auction with discrete bid levels [8]. Under this protocol, the auctioneer announces the price of the next bid and waits until a bidder indicates their willingness to pay this amount. Upon receiving such an indication, the price moves on to another higher discrete bid price, again proposed by the auctioneer. The auction continues until there are no bidders willing to pay the bid price requested by the auctioneer. At this point, the object is allocated to the current highest bidder and that bidder pays the last accepted discrete bid price.

Now, despite its apparent popularity, an auctioneer implementing an English auction with discrete bid levels is faced with two complementary challenges. Firstly, it must determine the actual discrete bid levels to be used. The standard academic auction literature provides little guidance here since it commonly assumes a continuous bid interval, where bidders incrementally outbid one another by an infinitesimally small amount. However, discrete bid levels do have an effect, and have been investigated by Rothkopf and Harstad [13]. They showed that the revenue of the auction is dependent on the number and distribution of discrete bid levels implemented and, in general, the use of discrete bid levels reduces the revenue generated by the auction. Conversely, the discrete bid levels also act to greatly reduce the number of bids that must be submitted in order for the price to reach the closing price. This has the effect of increasing the speed of the auction and, hence, reduces the time and communication costs of both the auctioneer and bidders. By analysing the manner in which the discrete bid auction could close and then calculating the expected revenue of the auctioneer in a number of limited cases (which we detail in section 2), they were able to derive the optimal distribution of bid levels that would maximise this revenue. In previous work, we extended this result to the general case, and we can now determine the optimal bid levels for an auction in which the environmental parameters are given [4]. Specifically, these parameters are the number of bidders participating in the auction and the bidders' valuation distribution .

Thus, performing this optimal auction design introduces the second of the two challenges; that of determining, for the particular setting under consideration, the values of these environmental parameters. While, in some settings these may be well known, in most cases they will not. Thus, in this paper, we tackle the problem of determining the optimal discrete bid levels when these values must be estimated through observations of previous auctions. In so doing, we extend the state of the art in three ways:

1. We extend previous work by deriving an expression that describes the expected revenue of a discrete bid auction when the number of bidders participating is unknown but can be described by a probability distribution.
2. We use this expression to calculate the optimal bid levels that maximise the auctioneers' revenue in this case. We demonstrate that the optimal discrete bid levels produced by this method are dependent on the distribution of the number of participating bidders and on the distribution that describes the bidders' valuations.
3. We show that this expression allows us to use machine learning, and specifically Bayesian inference, in an online algorithm that generates sequentially better estimates for the parameters that describe the two unknown distributions (i.e. the distribution of the number of bidders participating in any auction and the distribution of the bidders' valuation) by observing only the closing price of previous auctions.

The results that we provide may be used in the design of online auctions or may be used by automated trading agents that are adopting the role of an auctioneer within a multi-agent system. In such settings these auction protocols are attractive as they provide a relatively simple bidding strategy for the agents, yet, unlike second price sealed bid auctions, do not require the bidders to reveal their full private information to the auctioneer. In this setting, there is a need to fully automate the design of such auction mechanisms, and the work presented here represents a key step in this direction.

The remainder of the paper is organised as follows: in section 2 we present related work and in section 3 we describe our auction model and present the previously derived results for the expected revenue of this auction (in order to make this paper self-contained). In section 4 we extend this result to the case that the number of bidders participating in the auction is described by a distribution and we use this new result to derive optimal discrete bid levels in this case. In section 5 we present our Bayesian inference algorithm and finally we conclude and discuss future work in section 6.

2 Related Work

The problem of optimal auction design has been studied extensively for the case of auctions with continuous bid increments [12,10]. In contrast, auctions with discrete bid levels have received much less attention, and much of the work that does exist is based on the assumption that there is a fixed bid increment and thus the price of the auction ascends in fixed size steps [15,3,16]. In contrast, Rothkopf and Harstad considered the more general question of determining the optimal number and distribution of these bid levels [13]. They provided a full discussion of how discrete bid levels affect the expected revenue of the auction and they considered two different distributions for the bidders' private valuations (uniform and exponential). In the case of the uniform distribution, they considered two specific instances: (i) two bidders with any number of allowable bid levels, and (ii) two allowable bid levels with any number of bidders. In the first instance, evenly spaced bid levels (i.e. a fixed bid increment) was found to be the optimal. In the second instance, the bid increment was shown to decrease as the auction progressed. Conversely, for the exponential distribution (again with just two bidders), the optimal bid increment was shown to increase as the auction progressed.

In previous work, we extended the analysis of Rothkopf and Harstad [13], and, rather than analyse the ascending price English auction in limited cases, we presented a general expression that relates the revenue to the actual bid levels implemented. For a uniform valuation distribution we were able to derive analytical results for the optimal bid levels, and in general, we were able to numerically determine the optimal bid levels for any bidders' valuation distribution, any number of bid levels and any number of bidders. In addition, we showed that in general, increasing the number of discrete bid levels, causes the revenue to approach that of a continuous bid auction.

In this paper, we extend this previous work and address the problem of estimating the number and valuation distribution of the bidders through observing the closing price of previous auctions. This problem is similar to that studied in the econometrics literature, where it has been used to identify the behaviour of bidders in real world auctions [6]. More recently, it has received attention within electronic commerce, with the goal of determining the reserve price in a repeated procurement auction [2]. Typically, this work uses statistical maximum likelihood estimators to determine the parameters that describe the bidders' valuation distribution through observations of their bidding behaviour. In our case, this task is somewhat different as much of this information is lost in the discretisation of the bids. Thus, we use the expression that we have already derived for the revenue of the discrete bid auction, and use Bayesian inference to infer parameter values through observations of the closing price of previous auctions. This

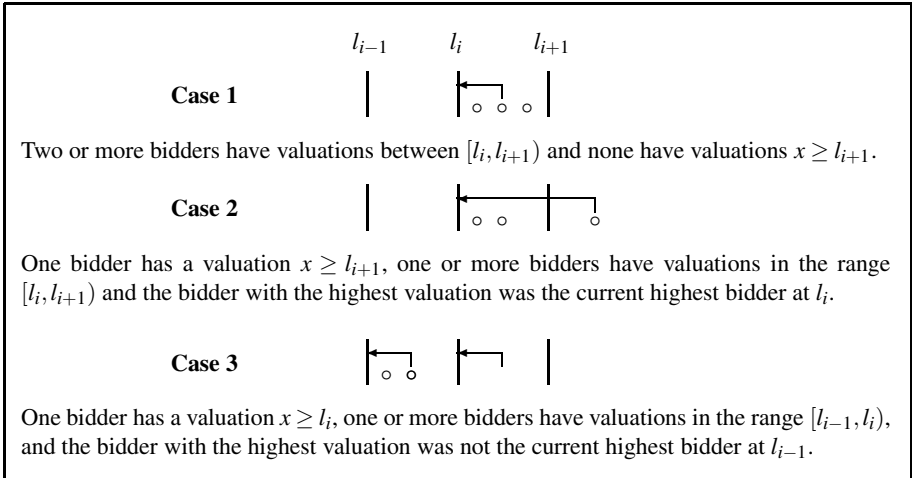


Fig. 1. Diagram showing the three cases whereby the auction closes at the bid level l_i . In each case, the circles indicate a bidder’s private valuation and the arrow indicates the bid level at which that bidder was selected as the current highest bidder.

method is attractive, as rather than providing a single parameter estimate at each iteration, it provides a full distribution that describes the auctioneer’s belief over the entire range of possible parameter values. Thus indicating the confidence that the auctioneer should have in his current estimate [9]. In addition, Bayesian inference tends to be computationally simpler than maximum likelihood methods, since it does not require us to maximise a function over several dimensions [1].

3 Auction Model and Expected Auction Revenue

In this work we consider a common model of an English auction that was used by Rothkopf and Harstad [13]. In this model, n risk neutral bidders are attempting to buy a single item from a risk neutral auctioneer. Bidders have independent private valuations, x_i , drawn from a common continuous probability density function, $f(x)$, within the range $[\underline{x}, \bar{x}]$, and with a cumulative distribution function, $F(x)$, where with no loss of generality, $F(\underline{x}) = 0$ and $F(\bar{x}) = 1$. The bidders participate in an ascending price auction, whereby the bids are restricted to discrete levels which are determined by the auctioneer. We assume there are $m + 1$ discrete bid levels, starting at l_0 and ending at l_m (at this point, we make no constraints on the actual number of these bid levels).

The auction starts with the auctioneer announcing the first discrete bid level (i.e. the reserve price of the auction) and asks the bidders to indicate their willingness to pay this amount. In traditional English auction houses, this indication is normally accomplished by a nod to the auctioneer, while in current online auctions such as www.onsale.com it requires a click of a mouse. If no bidders are willing to pay this amount within a predetermined and publically announced interval, the auction closes and the item remains unsold. However, if a bid is received, the auction proceeds and the auctioneer again requests

bidders willing to pay the next discrete bid level. If no bidders are willing to pay this new price, the auction then closes and the item is sold to the current highest bidder.

Now, in order to determine the optimal bid levels that the auctioneer should announce, an expression for the expected revenue of the auction must be found. Rothkopf and Harstad considered this problem and identified three mutually exclusive cases that described the different ways in which the auction could close at any particular bid level [13]. These cases are shown in figure 1. They then calculated the probability of each case occurring in a number of limited cases. In our earlier work we have been able to use the same descriptive cases, but derive a general result for each probability [4]. Thus we are able to describe the probability of the auction closing at any particular bid level:

$$P_n(l_i) = \begin{cases} [1 - F(l_i)] \left[\frac{F(l_{i+1})^n - F(l_i)^n}{F(l_{i+1}) - F(l_i)} \right] & i = 0 \\ [1 - F(l_i)] \left[\frac{F(l_{i+1})^n - F(l_i)^n}{F(l_{i+1}) - F(l_i)} + \frac{F(l_{i-1})^n - F(l_i)^n}{F(l_i) - F(l_{i-1})} \right] & 0 < i \leq m \end{cases} \quad (1)$$

Note that the subscript in P_n indicates that the expression is in terms of the actual number of bidders, n , who participate in the individual auction, and that we define $F(l_{m+1}) = 1$. Now, the expected revenue of the auctioneer is simply found by summing over all possible bid levels and weighting each by the revenue that it generates:

$$E_n = \sum_{i=0}^m l_i P_n(l_i) \quad (2)$$

Thus, by substituting equation 1 into this expression and performing some simplification, we get the result:

$$E_n = \sum_{i=0}^m \frac{F(l_{i+1})^n - F(l_i)^n}{F(l_{i+1}) - F(l_i)} \left[l_i [1 - F(l_i)] - l_{i+1} [1 - F(l_{i+1})] \right] \quad (3)$$

In our previous work we used this result to generate optimal bid levels when the number of bidders and the bidders valuation distribution are known.

4 Optimising over Uncertainty in the Number of Bidders

Now, we wish to deal with the more general case that the number of bidders participating in the auction is not known by the auctioneer. To do so, we have to carefully define what we mean by participation. Thus, a bidder is said to be participating in (or has entered) the auction, if they have generated a valuation for the item being sold, are present and are prepared to bid. It is this number of bidders (plus their valuation distribution and the discrete bid levels implemented) that determines the expected revenue of the auction (as described in equation 3). However, in the English auction considered here, not all of the bidders who are participating will necessarily submit bids to the auctioneer (i.e. many will find that the other bidders have raised the price beyond their own private valuation and thus they have no opportunity to bid). Thus, the auctioneer is not able to determine the number of bidders who are participating by simply observing the bids.

In addition, in any specific setting, the number of bidders participating in an auction is unlikely to be fixed but will most likely be described by a probability distribution.

Levin and Smith showed this by considering an auction model in which the number of bidders participating was endogenously determined [7]. They modeled a pool of potential bidders, and showed that, at equilibrium, each potential bidder has a fixed probability of actually participating in (or entering) the auction. The number of bidders participating in any auction was thus described by a binomial distribution. Bajari and Hortacsu considered a similar model and compared their model to data collected from eBay auctions selling collectable U.S. coins [1]. They note that in such online auctions, the pool of potential bidders is extremely large. However, the fact that, in general, only a small number of bids are observed, suggests that the probability that a potential bidder participates in any individual auction is very low. Thus, they assume that, in such cases, a Poisson distribution is an appropriate approximation for the binomial proposed by Levin and Smith. In light of this work, we describe the number of bidders participating in any auction by a Poisson distribution and thus the probability that n bidders participate is given by:

$$P(n) = \frac{v^n e^{-v}}{n!} \quad (4)$$

Here the parameter v describes the mean of this distribution and thus represents the expected or average number of participants in any individual auction. Given this distribution, we can extend the results described in the previous section and express the probability of the auction closing at any bid level, in terms of the parameter v , rather than n . To do so, we simply sum the probability given in equation 1 multiplied by the probability of that number of bidders actually occurring:

$$P_v(l_i) = \sum_{n=0}^{\infty} P(n)P_n(l_i) \quad (5)$$

Now substituting equations 1 and 4 into this expression and making use of the identity $\sum_{n=0}^{\infty} v^n/n! = e^v$ allows us to derive the result:

$$P_v(l_i) = \begin{cases} [1 - F(l_i)] \left[\frac{e^{v[F(l_{i+1})-1]} - e^{v[F(l_i)-1]}}{F(l_{i+1}) - F(l_i)} \right] & i = 0 \\ [1 - F(l_i)] \left[\frac{e^{v[F(l_{i+1})-1]} - e^{v[F(l_i)-1]}}{F(l_{i+1}) - F(l_i)} + \frac{e^{v[F(l_{i-1})-1]} - e^{v[F(l_i)-1]}}{F(l_i) - F(l_{i-1})} \right] & 0 < i \leq m \end{cases} \quad (6)$$

Now finally, as before, we are able to perform a weighted sum over all of the discrete bid levels to determine the expected revenue of the auctioneer given the uncertainty in the number of bidders that are participating in any specific auction:

$$E_v = \sum_{i=0}^m \frac{e^{v[F(l_{i+1})-1]} - e^{v[F(l_i)-1]}}{F(l_{i+1}) - F(l_i)} \left[l_i [1 - F(l_i)] - l_{i+1} [1 - F(l_{i+1})] \right] \quad (7)$$

This is a key result. It expresses the expected revenue of the auction in terms of the actual bid levels implemented, the bidders valuation distribution and, v , the mean number of bidders who participate in each auction. We use this result in the next section to derive optimal bid levels in spite of the inherent uncertainty in the number of bidders who will participate in any individual auction.

```

for i=0:m
     $l_i \leftarrow \begin{cases} a + i * (\bar{x} - a) / m & \text{where } a = \max(\underline{x}, \bar{x} / 2) & // \text{uniform} \\ 1 / \alpha + i * (2 / \alpha m) & & // \text{exponential} \end{cases}$ 
d  $\leftarrow \infty$ 
while d > stopping condition,
     $l'_0 \leftarrow \arg \max_{l_0} E_v(l_0, \dots, l_m)$  where  $\underline{x} \geq l_0 < l_1$ 
    for i=1:m-1
         $l'_i \leftarrow \arg \max_{l_i} E_v(l_0, \dots, l_m)$  where  $l_{i-1} < l_i < l_{i+1}$ 
     $l'_m \leftarrow \arg \max_{l_m} E_v(l_0, \dots, l_m)$  where  $l_{m-1} < l_m \leq \bar{x}$ 
    d  $\leftarrow 0$ 
    for i=0:m,
        d  $\leftarrow \max(d, \text{abs}(l'_i - l_i))$ 
         $l_i \leftarrow l'_i$ 

```

Fig. 2. Pseudo-code algorithm for calculating solutions for the optimal bid levels

4.1 Optimal Discrete Bid Levels

The expression presented in the last section describes the expected revenue of the auction when discrete bid levels $l_0 \dots l_m$ are used. Thus, in order to find the optimal bid levels in this case, we must find the values $l_0 \dots l_m$ that maximise this expression. In general, it is not possible to perform this maximisation analytically, so we must use a numerical algorithm. Now, given that there are many numerical multi-dimensional optimisation algorithms available (see Numerical Recipes [11] for examples), two key features of this problem guide our choice. Firstly, since each term in the summation in equation 7 contains only pairs of bid levels (i.e. l_i and l_{i+1}), we note that maximising this expression, and thus solving $\delta E_v / \delta l_i = 0$, is equivalent to solving a tri-diagonal set of $m + 1$ simultaneous equations, that, by denoting $\delta E_v / \delta l_i$ as f_i , we can write as:

$$\begin{aligned}
 f_0(l_0, l_1) &= 0 \\
 f_i(l_{i-1}, l_i, l_{i+1}) &= 0 \quad \text{for } i = 1 \text{ to } m - 1 \\
 f_m(l_{m-1}, l_m) &= 0
 \end{aligned} \tag{8}$$

Secondly, the solutions to these equations are constrained such that their ordering remains constant i.e. $l_{i-1} < l_i < l_{i+1}$. A general purpose optimisation package will fail to exploit the first feature and will be heavily constrained by the second. However, we can produce a simple and efficient numerical algorithm by using Jacobi iteration whereby we iteratively solve the $m + 1$ simultaneous equations [5]. That is, we fix all other bid levels, and we find the value of l_i that maximises equation 7 given that $l_{i-1} < l_i < l_{i+1}$. The expression is well behaved in this range and has a single maximum that can be found using hill climbing or a gradient based method. We update all l_i and then iterate the process until the bid levels converge to the necessary accuracy.

We present this algorithm in figure 2, noting that the expression $E_v(l_0, \dots, l_m)$ represents the revenue expression in equation 7. Whilst we do not prove the convergence properties of this iterative algorithm here, in our experiments it converged reliably given

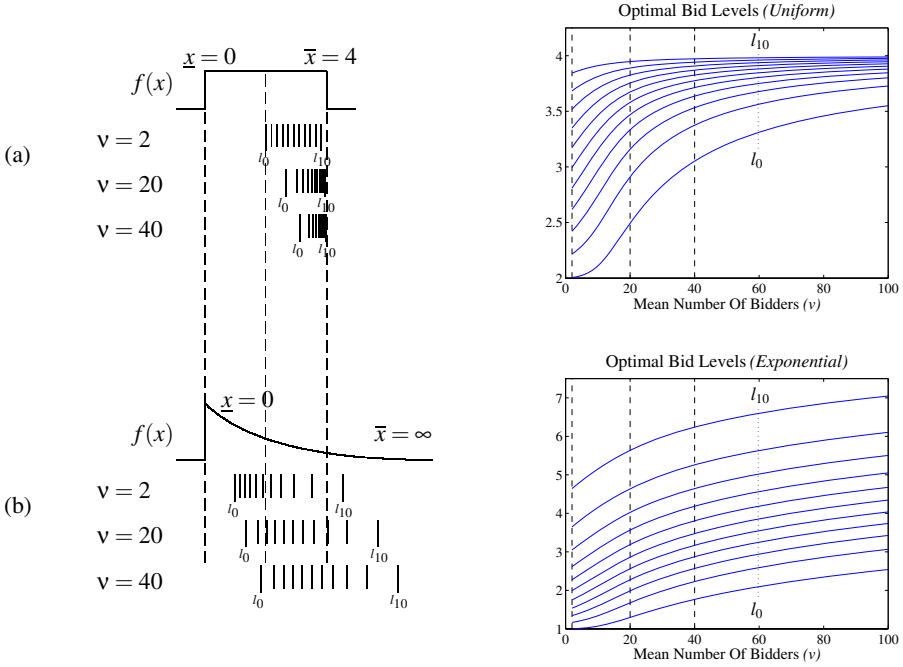


Fig. 3. Optimal bid levels for (a) uniform and (b) exponential valuation distributions

that two starting conditions for l_i were satisfied. Specifically, at the first iteration, no bid level may be outside the upper limit of the bidders’ valuation distribution (i.e. $l_i \leq \bar{x}$) and l_0 must be greater or equal to the reserve price predicted for the equivalent continuous bid auction (i.e. for a uniform bidders’ valuation distribution $l_0 \geq \max(\underline{x}, \bar{x}/2)$). We provide suitable starting conditions for the two valuation distributions that we consider in the next section in the algorithm.

4.2 Comparison of Valuation Distributions

The numerical solution described in the previous sections allows us to calculate the optimal discrete bid levels for any value of v (i.e. the mean number of bidders present in any auction) and any bidders’ valuation distribution. In this section, we compare the optimal bid levels over a range of values of v when two different bidders’ valuation distributions are used. Specifically, we compare the exponential distribution, proposed by Rothkopf and Harstad, with the more common uniform distribution, and, to allow us to compare these two directly, we chose their parameters so that the expected closing price of the auctions are similar in both cases. Thus for the uniform distribution, we consider a range of $[0,4]$ meaning $f(x) = \frac{1}{\bar{x}-\underline{x}}$ and $F(x) = \frac{x-\underline{x}}{\bar{x}-\underline{x}}$ where $\underline{x} = 0$ and $\bar{x} = 4$. For the exponential distribution, we have $f(x) = \alpha e^{-\alpha x}$ and $F(x) = 1 - e^{-\alpha x}$ where $\alpha = 1$. The resulting optimal discrete bid levels are shown in figure 3, for three different mean numbers of bidders ($v = 2, 20$ and 40) and over a continuous range from 2 to 100.

In both cases, we use 10 bid levels (i.e. $m = 10$), as this makes clear the differences between the two cases. Note that whilst changing the number of bid levels does affect their value, it does not affect the general form of the distribution seen in the plot.

Now, Rothkopf and Harstad showed that in the case where there were two bidders, the optimal discrete bid level distribution for the uniform distribution is a fixed bid increment with evenly spaced bid levels. In addition, for an exponential distribution, the optimal bid levels with two bidders is an increasing bid increment with bid levels becoming more widely spaced as the auction progresses. Our results show that in the general case, where there is uncertainty over the number of bidders that are participating, the distribution of the optimal discrete bid levels is complex. For the uniform distribution there is a decreasing bid increment whereby the discrete bid levels become closer together as the auction progresses. While, for the exponential distribution, the bid increment initially decreases, reaches a minimum size and then subsequently increases.

We also see that as the number of bidders increases, the value of l_0 increases. Rothkopf and Harstad fixed the values of the first and last bid levels at the extremes of the valuation distribution. However, we make no such restriction and thus the values of l_0 and l_m are optimised at the same time as the other bid levels. Now, since l_0 is equivalent to the reserve price of the auction (i.e. the item will not sell if there are no bidders willing to pay at least l_0) the results indicate that, in contrast to the literature of optimal auctions with continuous bid increments, the optimal reserve price of an auction with discrete bid levels is dependent on the mean number of bidders. In general, we see that when the number of bid levels is large, or the mean number of bidders is small, the value of l_0 tends toward the continuous result (for the uniform distribution, this is $x^* = \max(\underline{x}, \bar{x}/2)$, and for the exponential distribution it is $x^* = 1/\alpha$ [10]).

Intuitively we can understand these effects by the fact that given a fixed number of bid levels, we should position them closer together in areas where they are most likely to differentiate the bidders with the highest valuations. Thus, for the uniform distribution, the bid levels become closer together nearer to the upper limit of the distribution. Whilst in the exponential distribution, they become closer together where we expect to find the bidder with the second highest valuation. This result suggests that it may be possible to describe the optimal bid levels in terms of the distribution of the expected second highest valuation. However, it has not proved possible to describe the revenue of the discrete bid auction in these terms, so at the moment, this shortcut is not available to us.

5 Estimating Auction Parameters

In the previous sections, we showed that the optimal discrete bid levels, and hence the revenue of the auctioneer, are dependent on the number of bidders that participate and their valuation distribution. Now, when the values of the parameters that characterise these distributions are not known, we must estimate their value through observations of previous auctions. Since, in this paper we have derived an expression for the probability of the auction closing at any particular bid level (given these parameter values) it is natural to use Bayesian inference to perform this task. That is, having observed an auction closing at a certain bid level, we calculate our belief that a particular set of parameter values gave rise to this event. This method contrasts with statistical maximum

likelihood techniques since rather than simply providing a single ‘most likely’ parameter value, we derive a distribution that describes our belief over all possible values.

To illustrate this process, we describe a general setting, in which an auctioneer implements a regularly repeating auction, and in each auction a single identical item is sold. As described earlier, we assume that there is a large pool of potential bidders, who have private independent valuations that are drawn from a common distribution. Each potential bidder has a small probability of actively participating in any auction, and thus each repeated auction faces a number of bidders that is described by the Poisson distribution shown in equation 4. Note that whilst their numbers are similar, these bidders are different individuals with different valuations and, since we are explicitly interested in the actions of the auctioneer, we assume that their bidding behaviour is unaffected by their own observations of previous auctions¹. Thus, our goal is to estimate the typical number of bidders who participate in each auction, v , and also the parameters that describe their common valuation distribution. These estimated parameter values can then be used to calculate optimal discrete bid levels in subsequent auctions.

5.1 Estimating the Mean Number of Bidders

We first consider an example in which the bidders’ valuation distribution is known, but, v , the parameter that characterises the Poisson distribution and represents the mean number of bidders participating in each repeated auction, is unknown. Thus, if at time t the auctioneer implemented an auction that used the discrete bid levels $\mathbf{I} = \{l_0^t \dots l_m^t\}$ and closed at bid level l_w^t , we wish to find the value v that best explains this outcome. In other words, we wish to calculate the probability distribution $P(v|l_w^t, F(x), \mathbf{I})$. Now, in equation 6 we have already derived the probability of the auction closing at any bid level, in terms of the mean number of bidders, the bidders’ valuation distribution and the actual bid levels implemented. Thus, in the notation we are using here, we have already derived $P(l_w^t|v, F(x), \mathbf{I})$. With this expression, we can use Bayes’ theorem in order to calculate the required result:

$$P(v|l_w^t, F(x), \mathbf{I}) = \frac{P(l_w^t|v, F(x), \mathbf{I})P(v)}{P(l_w^t|F(x), \mathbf{I})} \quad (9)$$

Now, this described the case where the auctioneer has made an observation of a single auction. In general, if t such auctions have been observed, the auctioneer can use all of this evidence to improve its estimate. Thus if the bid levels used in these auctions were $\mathbf{L} = \{\mathbf{I}^1, \dots, \mathbf{I}^t\}$, and the observed closing prices were $\mathbf{l}_w = \{l_w^1, \dots, l_w^t\}$, we have:

$$P(v|\mathbf{l}_w, F(x), \mathbf{L}) = \frac{\prod_{i=1}^t P(l_w^i|v, F(x), \mathbf{I}^i) P(v)}{Z} \quad (10)$$

In this expression, Z is a normalising factor that ensures that $P(v|\mathbf{l}_w, F(x), \mathbf{L})$ sums to one over the range of possible values of v . Now, $P(v|\mathbf{l}_w, F(x), \mathbf{L})$ is a continuous prob-

¹ This assumption is reasonable in circumstances where historical auction data is not available to the bidders. However, we intend to investigate the full implications of this assumption in future work.

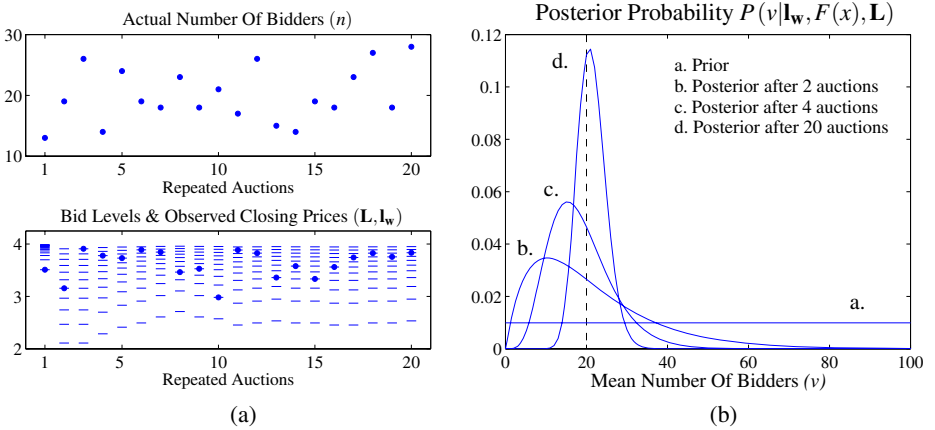


Fig. 4. Plots showing (a) the actual number of bidders that participated in the auction (unknown to the auctioneer) and the actual bid levels and closing prices observed by the auctioneer and (b) the prior and posterior belief distributions of the auctioneer after 2, 4 and 20 repeated auctions

ability distribution. However, for our purposes, we calculate it as a discrete probability distribution over a suitable range. In this example, we calculate $P(v | l_w, F(x), L)$ for integer values of v from \underline{v} to \bar{v} . Thus, this normalising factor is given by:

$$Z = \sum_{v=\underline{v}}^{\bar{v}} \left[\prod_{i=1}^t P(l_w^i | v, F(x), l^i) P(v) \right] \quad (11)$$

Finally, $P(v)$ represents the auctioneers' prior belief; an initial assumption as to which values of v are most likely to occur, before any observations have been made. If no such intuition is available (as in our simulations here), the prior can simply be initialised as a uniform distribution, and it will have no effect on the estimates generated.

Thus the procedure adopted by the auctioneer is as follows: it first uses its prior belief (i.e. an initial guess) to calculate the bid levels for the first auction. Having observed the closing price of this auction, the expression in equation 10 is used to calculate the probability distribution that describes its updated belief in the parameter v . This probability distribution is then used to choose the value of v for the calculation of the optimal bid levels to be implemented in the next auction. There are two ways in which this choice can be made, either: (i) the most likely value of v can be used (i.e. the value of v where the probability distribution has a maximum), or (ii) a value of v may be sampled from this probability distribution. The first option is identical to a statistical maximum likelihood estimator. However the second option ensures more rapid convergence in cases where the auctions that occur early in the learning process represent extreme events (i.e. when many of the bidders have extremely high or low valuations or the auction happens to have many more or many less bidders than is typical).

Simulation results for this procedure are shown in figure 4. Here, we consider the same uniform valuation distribution as discussed in section 5 (i.e. $\underline{x} = 0$ and $\bar{x} = 4$). The real value of v in this case is 20, whilst the auctioneer's prior belief is that it lies somewhere between 0 and 100 (i.e. $P(v)$ is a uniform distribution over this range). In

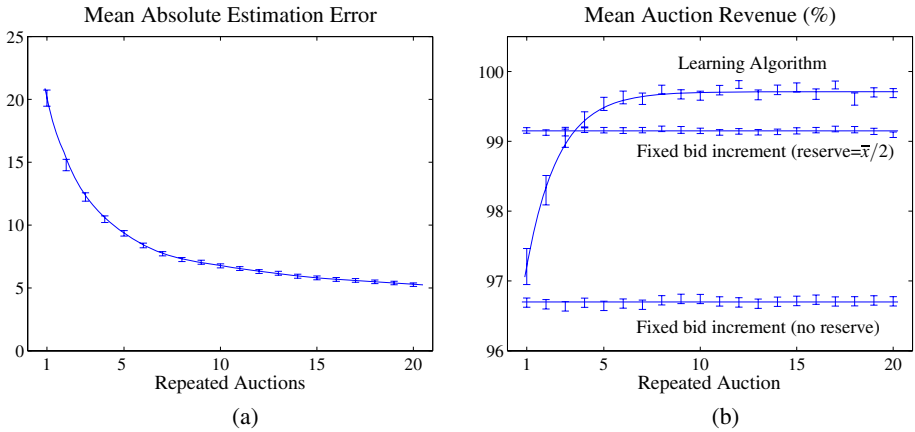


Fig. 5. Plots showing (a) the converging estimates generated by the learning algorithm, and (b) how this results in improvements in the auctioneer’s revenue

figure 4a we show the actual number of bidders that participated (unknown to the auctioneer) and the bid levels that were implemented in each repeated auction, along with the actual bid level at which the auction closed (denoted by a filled circle on the appropriate bid level and observed by the auctioneer). In figure 4b, we show the probability distribution, $P(v|\mathbf{l}_w, F(x), \mathbf{L})$, that describes the auctioneers’ belief in the values of v (shown after 2, 4 and 20 auctions). The variance in the observed auction closing prices is driven by the stochastic nature of the number of bidders, their valuations and also the changing auction bid levels. However, despite this variance, the auctioneers’ belief in the most likely value of v converges rapidly to the true value. Thus the bid levels implemented by the auctioneer also converge to the those that generate the maximum revenue.

To demonstrate the convergence of this algorithm, after each repeated auction we calculate the error in the estimate that it produced (i.e. the difference between the estimated value and the true value). We repeat the process 1000 times using the same parameter values (i.e. $v = 20$ and a uniform bidders’ valuation distribution where $\underline{x} = 0$ and $\bar{x} = 4$) and average over the results. Figure 5a shows the mean absolute estimation error plotted against the number of repeated auctions. The plot shows that the estimates improve rapidly after the first few auctions and then converge to the true value.

Figure 5b shows the improvement in revenue that results from more accurately estimating the mean number of bidders who are participating in the auctions, and then use this result to optimise the discrete bid levels used in subsequent auctions. For the same simulation runs presented in figure 5a, we show the efficiency of the auction, calculated in terms of the percentage of the second highest bidder’s valuation that the auction was able to extract. We compare this revenue to that which would have been achieved with an auction that used the more commonly implemented fixed bid increment, with and without setting a reserve price. Clearly, as the estimates of the auction parameters improve, so the revenue of the auctioneer increases. Significantly, the greatest improve-

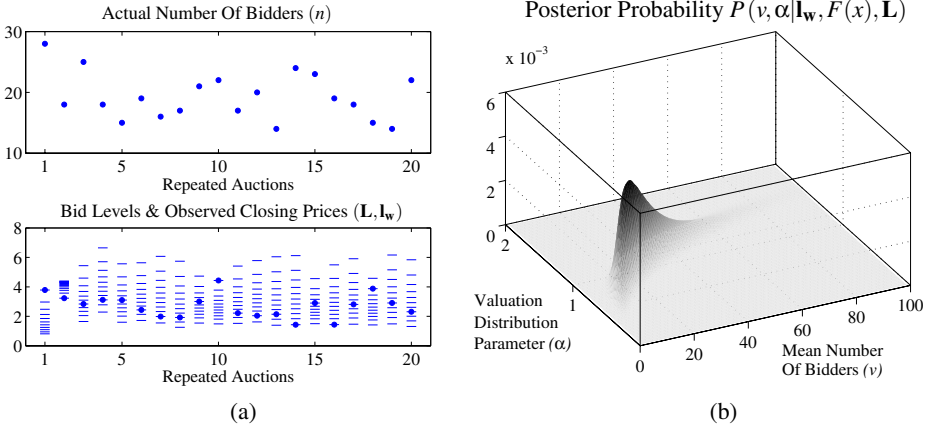


Fig. 6. Plots showing (a) the actual number of bidders that participated in the auction (unknown to the auctioneer) and the actual bid levels and closing prices observed by the auctioneer and (b) the joint posterior belief distributions of the auctioneer after 20 repeated auctions

ment is realised after the first few auction and after this point, the revenue exceeds that generated with fixed bid increments.

5.2 Estimating Multiple Parameters

The algorithm that we have presented here is certainly not restricted to learning single parameters. In figure 6 we present a second example, this time for the exponential valuation distribution presented in section 4.2. In this case we infer both the value of parameter that describes the distribution of the number of bidders, v , and the value of the parameter that describes the bidders' exponential valuation distribution, α . Thus we must calculate the two-dimensional joint probability distribution $P(v, \alpha | \mathbf{l}_w, F(x), \mathbf{L})$. Again, despite the stochastic nature of the auction process, after twenty repeated auctions the probability distribution shows a clear peak around the true values of $v = 20$ and $\alpha = 1$, and thus the bid levels converge toward the true optimal bid levels. Space does not allow us present a full analysis of the convergence, however, in general, increasing the number of parameters that are learnt reduces the convergence rate.

We can extend this method to estimate more parameters, by simply calculating larger joint probability distributions in more dimensions. However, in so doing, the cost of performing this exact calculation increases geometrically. Fortunately Bayesian inference is a well developed field with sophisticated methods that allow us to approximate these distributions. For example, variational methods (which we intend to explore in the future) allow us to approximate the full n -dimensional joint distribution as the product of n independent distributions, with a corresponding computational saving [9].

6 Conclusions

In this paper we considered the optimal design of English auctions with discrete bid levels and our aim was to automate their configuration to generate the maximum revenue

for the auctioneer. To this end, we extended earlier work and derived an expression for the revenue of the auction under uncertainty in the number of bidders who are participating in the auction. We used this result to numerically calculate optimal bid levels under this uncertainty and showed that the value and distribution of these optimal bid levels are highly dependent on both the mean number of bidders and the bidders' valuation distribution. Finally, we considered the case in which these environmental parameters are unknown to the auctioneer, and used Bayesian inference to estimate these parameters through observations of the closing price of previous auctions. We showed that despite the stochastic nature of the auctions, the estimates generated by this algorithm rapidly converged to the true values. In addition, we showed that by correctly estimating the true values of these parameters, the auctioneer is able to bid levels that result in an increase in auction revenue.

Our future work in this area consists of extending the auction model to incorporate an explicit expression of the auctioneer's costs (rather than the explicit bound on the maximum number of bid levels that we have presented here). In addition, we intend to extend the inference method that we have presented here, and in particular, we would like to use these techniques to perform model identification and selection. Thus, we would infer the full parameters of several different valuation distributions (using variational methods to minimise the computational cost of this task) and then infer which of these distributions best explains the closing prices that were observed (also considering the effect that an incorrect assumption will have). In so doing, we believe these techniques will significantly contribute toward our goal of automating the mechanism design of optimal discrete bid auctions.

Acknowledgments

This research was funded by the DIF-DTC project on Agent-Based Control and the ARGUS II Defence and Aerospace Research Partnership. This is a collaborative project involving BAE SYSTEMS, QinetiQ, Rolls-Royce, Oxford University and Southampton University, and is funded by the industrial partners together with the EPSRC, Ministry of Defence (MoD) and Department of Trade and Industry (DTI). Sarit Kraus is also affiliated with the University of Maryland Institute for Advanced Computer Studies (UMIACS) and this work was part supported by NSF Grant IIS-0208608.

References

1. P. Bajari and A. Hortacsu. The winners curse, reserve prices, and endogenous entry: empirical insights from eBay auctions. *RAND Journal of Economics*, 34(2):329–355, 2003.
2. M. Bichler and J. Kalagnanam. A non-parametric estimator for setting reserve prices in procurement auctions. In *ACM Conference on Electronic Commerce 2003*, pages 254–255, 2003.
3. M. S.-Y. Chwe. The discrete bid first auction. *Economics Letters*, 31:303–306, 1989.
4. E. David, A. Rogers, J. Schiff, S. Kraus, and N. R. Jennings. Optimal design of english auctions with discrete bid levels. In *Proceedings of ACM Conference on Electronic Commerce*, pages 98–107, 2005.

5. L. A. Hageman and D. M. Young. *Applied Iterative Methods*. Academic Press, 1981.
6. J.-J. Laffont, H. Ossard, and Q. Vuong. Econometrics of first-price auctions. *Econometrica*, 63(4):953–980, 1995.
7. D. Levin and J. L. Smith. Equilibrium in auctions with entry. *American Economic Review*, 84(3):585–99, 1994.
8. D. H. Lucking-Reiley. Auctions on the internet: What’s being auctioned, and how? *Journal of Industrial Economics*, 48(3):227–252, 2000.
9. D. J. C. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2003.
10. R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
11. W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, 1992.
12. J. G. Riley and W. F. Samuelson. Optimal auctions. *American Economic Review*, 71:381–392, 1981.
13. M. H. Rothkopf and R. Harstad. On the role of discrete bid levels in oral auctions. *European Journal of Operations Research*, 74:572–581, 1994.
14. M. H. Rothkopf, T. J. Teisberg, and E. P. Kahn. Why are Vickrey auctions rare? *Journal of Political Economy*, 98(1):94–109, 1990.
15. B. S. Yamey. Why 2,310,000 [pounds] for a Velazquez? An auction bidding rule. *Journal of Political Economy*, 80:1323–1327, 1972.
16. J. Yu. *Discrete Approximation of Continuous Allocation Mechanisms*. PhD thesis, California Institute of Technology, Division of Humanities and Social Science, 1999.