Learn While You Earn: Two Approaches to Learning Auction Parameters in Take-it-or-leave-it Auctions

(Short Paper)

Archie C. Chapman, Alex Rogers and Nicholas R. Jennings
{ac05r,acr,nrj}@ecs.soton.ac.uk
Electronics and Computer Science
University of Southampton
Southampton, SO17 1BJ, UK.

ABSTRACT

Much of the research in auction theory assumes that the auctioneer knows the distribution of participants' valuations with complete certainty. However, this is unrealistic. Thus, we analyse cases in which the auctioneer is uncertain about the valuation distributions; specifically, we consider a repeated auction setting in which the auctioneer can learn these distributions. Using take-it-or-leave-it auctions (Sandholm and Gilpin, 2006) as an exemplar auction format, we consider two auction design criteria. Firstly, an auctioneer could maximise expected revenue each time the auction is held. Secondly, an auctioneer could maximise the information gained in earlier auctions (as measured by the Kullback-Liebler divergence between its posterior and prior) to develop good estimates of the unknowns, which are later exploited to improve the revenue earned in the long-run. Simulation results comparing the two criteria indicate that setting offers to maximise revenue does not significantly detract from learning performance, but optimising offers for information gain substantially reduces expected revenue while not producing significantly better parameter estimates.

Categories and Subject Descriptors

J.4 [**Social and Behavioural Sciences**]: Economics

General Terms

Design, Economics

Keywords

Optimal auction design, Bayesian experimental design

1. INTRODUCTION

The use of agent mediated electronic commerce has grown rapidly over recent years and represents a vast potential market. This popularity has prompted much research into agent mediated auctions and, specifically, the development of autonomous software agents that fulfil the role of auctioneer or bidder on behalf of their owner. However, much of this work makes strong assumptions about the type of information available to the auction designer. In particular, it is often assumed that, although individual bidders' valuations are private, the auctioneer knows the distribution from which they are drawn (e.g. [5, 2, 8]). In this paper we relax this somewhat unrealistic assumption to analyse cases where the auctioneer is uncertain about the distribution of bidders' valuations, but we consider a repeated auction setting so the auctioneer has an opportunity to learn the distribution.

The subject of our investigation is the optimal design of one particular variant of the take-it-or-leave-it auction protocol [8], a single-offer take-it-or-leave-it auction with symmetric bidders (which we abbreviate as TLA). Under this protocol, each bidder is made an *offer* for the single item. If they accept the offer, the auction closes and the item is allocated to them. If they reject the offer, they leave the auction and the next bidder is made a new offer. The auction continues until the item is sold or all potential bidders reject the offers made to them. Consequently, bidders have a dominant strategy to accept an offer if it is less than their valuation. This protocol is used as a canonical setting in which we can examine the fundamentals of our problem, without being distracted by the complications of other auction protocols.

In this situation, the auctioneer faces two complementary challenges. The first is to determine the actual values of the offers made to the bidders, which brings us to our first extension to the state of the art. We derive an expression for the expected revenue of an auction given uncertainty over the parameters that describe the bidders' valuation distribution, and describe a procedure for maximising this expression. This is achieved by using reasonable priors over parameter values to represent the uncertainty, and integrating over these priors to find the expected revenue of the auction. The second challenge is determining appropriate priors over the values of these parameters for use in optimising the auction design. To do this, we consider repeated auction setting, where identical items are sold to a large pool of bidders, giving the auctioneer an opportunity to learn these values. An online Bayesian inference algorithm is then used to refine the estimates of the parameters, which are used as priors for the next auction. The auctions are designed to maximise revenue, so by following the procedure described above the auctioneer "learns while it earns".

However, learning from the outcomes of auctions designed to maximise expected revenue does not necessarily imply efficient, or even rapid, learning. In particular, we expect that if an auctioneer were to have more accurate estimates of the unknown parameter values earlier in a sequence of auctions, the average total revenue of subsequent auctions could be improved. This may, in turn, increase the total revenue of the complete sequence of auctions.

Cite as: Learn While You Earn: Two Approaches to Learning Auction Parameters in Take-it-or-leave-it Auctions (Short Paper), A. C. Chapman, A. Rogers and N. R. Jennings, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008,Estoril,Portugal,pp. 1561-1564.

To test this hypothesis, and as a second extension to the literature, we derive an expression for the additional information elicited by running the auction, and procedure for maximising this expression. We use the expected *Kullback-Liebler* (KL) *divergence* between all possible posterior distributions and the prior distribution of the unknowns as an experimental design utility function, just as one would use expected revenue [1]. In the context of Bayesian inference, KL divergence is the natural measure of information gain, so using it should improve the learning performance of our Bayesian inference algorithm.

Our third extension to the state of the art is a comparison of the two procedures. Using the expressions described above as design criteria, we apply them to the problem of learning unknown parameter values from the closing price of TLAs, and compare their learning performance and revenue. Our results show that using offers that maximise expected information gain produce only a marginally better learning performance than compared to using offers that maximise expected revenue, but at a large cost in terms of foregone revenue. The implication of this result is that, although not specifically targeted at maximising the information gained from an auction, offers that maximise expected revenue do indeed provide sufficient information for the Bayesian machine learning algorithm to learn effectively.

The paper progresses as follows. Next, we describe the extended TLA model that considers uncertainty in the parameters of the bidders' valuation distribution. Then, in section 3, we present the first auction design criterion, which sets offers to maximise the expected revenue of the auction given uncertain parameter values. Section 4 details the online Bayesian machine learning algorithm used by our auctioneer. In section 5, we present our second design criterion, which sets offers to maximise the expected information gained from an auction. Then, in section 6, we compare how our Bayesian inference algorithm learns from offers generated by each approach and how using each criterion affects the revenue generated in a series of auctions. The final section summarises the paper and discusses future work.

2. THE AUCTION MODEL

In this paper we consider the model of a TLA introduced by Sandholm and Gilpin [8]. In particular, we investigate the single-offer variant as it is the most readily applicable within a real world scenario and it avoids the problem of having to decide how many offers to make to each potential bidder. We consider a series of such auctions as a model of a situation where the auctioneer can learn the values of the parameters of the bidders' valuation distributions.

Under the single-offer TLA protocol, an auctioneer has an indivisible good, which it values v_0, that it can allocate to any one of a set of n risk neutral bidders. The auctioneer approaches each potential bidder in sequence and proposes one price for the item. Let $\mathbf{o} = o_1, \ldots, o_n$ be the sequence of n offers. If bidder i accepts the offer, the item is allocated to that bidder at price $o_{w=i}$, the auction finishes, and the seller gains utility $U = o_{w=i} - v_0$. If the bidder rejects the offer they leave the auction and the next bidder is offered the item. The auction continues until the item is sold or all bidders reject the offer made to them, at which point the auctioneer's utility for holding the auction is zero.

We analyse *symmetric* TLAs, where each bidder has a valuation v_i for the good that is independently drawn from a common distribution with cumulative density $F(v)$. Making this restriction removes the problem of ordering offers, allowing us to focus solely on the offer levels implemented by the auctioneer. The bidders' valuation distribution is itself characterised by a vector of parameters θ which are, in the full information case, known to the auctioneer. In the extended model considered here, this assumption is relaxed to analyse the effects of uncertainty over the parameters.

2.1 The Full Information Case

In the single-offer TLA protocol, the decision to accept an offer is dependent only on a bidder's own valuation. As such, a bidder's dominant strategy is to accept any offer that is less than its valuation. Thus, the probability of the auction closing on a particular offer is the probability that the current bidder i has a valuation v_i greater than o_i, multiplied by the probability that all previous offers were rejected. The probability of closing on offer $P(o_{w=i})$, given that all $i-1$ previous offers have been rejected, is:

$$P(o_{w=i}) = [1 - F(o_i)] \prod_{j=1}^{i-1} F(o_j), \quad 1 \leq i \leq n. \tag{1}$$

Then, the probability that the item is not allocated is the probability that every bidder has a valuation lower than the offer it is made:

$$P(o_{w=0}) = \prod_{i=1}^{n} F(o_i). \tag{2}$$

Under Sandholm and Gilpin's model (recalling that $o_0 = v_0$), the expected utility of an auctioneer setting offers at the beginning of an auction to maximise revenue is:

$$\mathbb{E}[U(\mathbf{o})] = \sum_{i=0}^{n} P(o_{w=i})\, o_i. \tag{3}$$

Now, working backwards from the n^{th} offer, see that the expected value of an offer is determined only by its own value and that of the offers after it. Using this, Sandholm and Gilpin construct a simple algorithm to solve for the optimal offers. To begin, the auctioneer sets a virtual reserve price, π, equal to its own valuation (i.e. $\pi = v_0$). It then sets the last offer, o_n, to maximise its expected revenue given π, recalculates the virtual reserve price using this offer, and computes the next highest offer o_{n-1}. By backward induction, as each offer is optimal at each step of the auction, the algorithm produces optimal offers, and offers naturally decrease over time.

2.2 The Incomplete Information Case

The full information model assumes the auctioneer has perfect knowledge of the parameters of the bidders' valuation distribution. The model we consider is of the more general case where the parameters, θ, describing the bidders' valuation distribution are not known. We represent the auctioneers' initial uncertainty about θ with a prior $P(\theta)$, which represents an initial assumption as to which values of θ are most likely to occur.

We illustrate this general approach by considering two example valuation distributions. Firstly, the bidders' valuations may be drawn from a uniform distribution:

$$F(v) = \frac{v-a}{b-a}$$

with lower and upper supports a and b, and where the auctioneer's initial uncertainty about b is represented by a uniform prior with support $[\underline{b}, \overline{b}]$.[1] Secondly, the bidders' valuations may be drawn from an exponential distribution, where the unknown parameter is the mean of the distribution, $1/\alpha$:

$$F(v) = 1 - e^{-\alpha v}.$$

[1]Our use of uniform priors is reasonable because in economic scenarios, there are almost always bounds on the likely valuations bidders will hold for an item. This information is captured in the range of the prior.

```
d ← ∞
while d > stopping condition,
  for i=n:1
    o'_i ← argmax_{o_i} E[U(o)]
  d ← 0
  for i=1:n,
    d ← max(d, abs(o'_i − o_i))
    o_i ← o'_i
```

Figure 1: Algorithm for optimal offers in the extended model.

Again, the auctioneer's initial uncertainty about the value of α is also represented by a uniform prior, $\alpha \in [\underline{\alpha}, \overline{\alpha}]$. These two examples have been chosen because they frequently appear in the literature (e.g. [7, 3, 6]), and uncertainty in these distributions is easily represented by uncertainty in a single parameter. However the general approach can be used to learn distributions with multiple parameters, or may be used in conjunction with Bayesian model selection techniques [4].

3. EXPECTED REVENUE CRITERION

Uncertainty in the value of θ requires a new method to maximise the expected revenue of the auction. First, a change of notation is necessary to explicitly consider uncertainty in the parameters. Thus, given a valuation distribution $F(v,\theta)$ and a set of offers $\mathbf{o}$, let the probability of closing on offer o_i, after $i-1$ other offers have been rejected, or of not closing, o_0, be written as:

$$P(o_{w=i}|\theta,\mathbf{o}) = \begin{cases} [1-F(o_i,\theta)]\prod_{j=1}^{i-1}F(o_j,\theta), & 1 \le i \le n, \\ \prod_{i=1}^{n}F(o_i,\theta) & i = 0. \end{cases} \quad (4)$$

To maximise expected revenue, the auctioneer can integrate out uncertainty in the parameter values using $P(\theta)$ to moderate the value of the expected revenue for each value of θ. Thus, the expected utility of an auctioneer interested in maximising revenue is:

$$\mathbb{E}[U(\mathbf{o})] = \int P(\theta)\sum_{i=o}^{n}P(o_{w=i}|\theta,\mathbf{o})\, o_i\, d\theta. \quad (5)$$

Note that in practice we perform these calculations numerically by considering a discrete approximation of $P(\theta)$ from $\underline{\theta}$ to $\overline{\theta}$.

In order to find the optimal offers, we must set $\mathbf{o}$ to maximise equation (5). As we were unable to find an analytical solution to this expression, we use a numerical algorithm based on Jacobi iteration. Specifically, while fixing all other offers, we find the value of o_i that maximises equation (5) using a one-dimensional search method (such as golden section search). We sequentially update all o_i and iterate the process until the offer levels converge to the necessary accuracy.

We present this algorithm in pseudo-code in figure 1, and note that $\mathbb{E}[U(\mathbf{o})]$ represents the revenue expression shown in equation (5). Whilst our purpose is not to prove the convergence properties of this algorithm, in our experiments it was found to converge reliably and rapidly, with the only constraints on starting conditions being that the offers are in the correct order (i.e. descending).

4. ESTIMATING AUCTION PARAMETERS

In section 2.2 of this paper we extended the model of TLAs to consider uncertainty in the parameters that describe the bidders' valuation distribution. Now, under the commonly made assumption that the participants in each auction are drawn from a large pool of potential bidders whose valuations are described by the fixed distribution $F(v)$ (e.g. [7, 3, 6]), the auctioneer can use the information contained in the closing price to refine its estimate of the unknown auction parameters. In order to compare our results to Sandholm and Gilpin's, we only consider the case where the auctioneer learns at the end of each auction. Given this, we use Bayesian inference to estimate parameter values via the expressions for the probability of closing on a particular offer given in equation (4). Bayesian inference provides a full distribution that describes the auctioneer's belief over the entire range of possible parameter values. The shape of this distribution indicates the confidence that the auctioneer should have in his current estimate [4].

By this approach, the auctioneer updates its joint distribution over the unknown parameters θ using Bayes' theorem. In general, if T auctions have been observed, the auctioneer can use all of this evidence to improve its estimate. Thus if the offer levels used in auction $t \in T$ were $\mathbf{o}^t$, all the levels used in T auctions were $\mathbf{O} = \{\mathbf{o}^1, \ldots, \mathbf{o}^T\}$, and the sequence of observed closing prices were $\mathbf{o}_w = \{o^1_{w=i}, \ldots, o^T_{w=i}\}$, we have:

$$P(\theta|\mathbf{o}_w,\mathbf{O}) = \frac{\prod_{t=1}^{T}P(o^t_{w=i}|\theta,\mathbf{o}^t)\,P(\theta)}{\int\prod_{t=1}^{T}P(o^t_{w=i}|\theta,\mathbf{o}^t)\,P(\theta)d\theta}. \quad (6)$$

Again, in this expression the denominator is a normalising factor that ensures that $P(\theta|\mathbf{o}_w,\mathbf{o}^t)$ sums to one.

The auctioneer adopts the following procedure to learn the unknown parameter values. Using its prior belief, it selects offers for the first auction using equation (5). After observing the closing price of this auction, it uses the expression in equation (6) to update its belief over the parameters θ. This belief distribution is then used as $P(\theta)$ to calculate the offers for the next auction. The process of refining the estimate of θ and implementing new offer levels is repeated for subsequent auctions.

5. INFORMATION GAIN CRITERION

In section 4, the auctioneer learns θ through observing the behaviour of bidders in auctions designed to maximise revenue. However, if the auctioneer had better estimates of θ earlier in the sequence of auctions, it could implement better offers and earn greater revenue in subsequent auctions. To this end, in this section, the problem of setting offers specifically so the auctioneer can learn in the most effective way possible is addressed. By doing so, we can then answer the associated question of whether more rapid learning in earlier stages of the series of auctions will increase the revenue earned over the longer term.

The gain in information about parameters θ obtained by selecting an experimental (auction) design $\mathbf{o} = \{o_1, \ldots o_n\}$, observing an outcome $o_{w=i}$ which is then used to update the estimate of θ, may be measured by the KL divergence between the prior and the posterior distributions of θ:

$$D_{KL}[P(\theta|o_{w=i},\mathbf{o})||P(\theta)] = \int P(\theta|o_{w=i},\mathbf{o})\log\left(\frac{P(\theta|o_{w=i},\mathbf{o})}{P(\theta)}\right)d\theta.$$

However, as the outcome has not yet been observed, the expected amount of information gained by holding an auction is the average KL divergence over all possible outcomes in $\mathbf{o}$, as each outcome will result in a different posterior estimate. Thus, the expected utility of an auctioneer interested in maximising the information

Average revenue - uniform valuation distribution

Offers maximising	1^{st} auction	2^{nd} auction
D_{KL}, revenue	0.5804	0.6846
Revenue, revenue	0.6126	0.6846

Average revenue - exponential valuation distribution

Offers maximising	1^{st} auction	2^{nd} auction
D_{KL}, revenue	1.393	1.723
Revenue, revenue	1.472	1.720

Figure 2: Revenue comparison. First auction expected revenue for offers maximising D_{KL} or revenue; Second auction revenue from offers maximising revenue using posterior from outcome of first auction based on D_{KL} or revenue maximising offers.

it gains about the bidders' valuation distribution is:

$$\mathbb{E}[U(\mathbf{o})] = \sum_{i=0}^{n} P(o_{w=i}|\mathbf{o}) \int P(\theta|o_{w=i},\mathbf{o}) \log\left(\frac{P(\theta|o_{w=i},\mathbf{o})}{P(\theta)}\right) d\theta. \tag{7}$$

We now describe how we implement this criterion. Equation (7) is in a clumsy form that requires computing all potential posterior distributions generated by the different outcomes of the auction. By applying Bayes' theorem twice the following expression can be implemented in our maximum expected KL divergence algorithm:

$$\mathbb{E}[U(\mathbf{o})] = \sum_{i=0}^{n} \int P(o_{w=i}|\theta,\mathbf{o}) P(\theta) \log\left(\frac{P(o_{w=i}|\theta,\mathbf{o})}{P(o_{w=i}|\mathbf{o})}\right) d\theta \tag{8}$$

One convenient aspect of equation (8) is that it may also be maximised using the iterative algorithm used to maximise the expected revenue criterion (figure 1, discussed in section 3). In this case, $\mathbb{E}[U(\mathbf{o})]$ represents the expected KL divergence given by equation (8).

6. COMPARING THE CRITERIA

In this section we compare how setting offers using either the expected revenue or the expected information gain criterion affects the performance of the Bayesian machine learning algorithm and the revenue generated by the auction. We implement offers in a simulated TLA and use the results to test whether or not the information gain–maximising offers produce a posterior estimate that can be used to generate greater revenue in subsequent auctions, compared to simply optimising offers for revenue.

We addresses this question by looking at the effects of learning from either approach on the revenue generated if the auction were run a second time. We consider the case of 8 bidders. In the first auction, offers are generated from an ignorant prior to either maximise expected revenue or expected information gain. In the second auction, both sets of offers maximise expected revenue, however one set is generated using the posterior of the information gain maximising auction, while the other uses the posterior of the revenue maximising auction. The rationale for this experiment is that it is the limiting case: In the first auction you have the most to learn, while the second auction presents the greatest potential to exploit the additional information to earn additional revenue.

Figure 2 shows that in the first auction, the revenue maximising offers generate significantly more revenue than those set to maximise information gain. The expected revenue for the uniform distribution is 0.6126 for the revenue maximising offers and 0.5804 for the information gain maximising offers. For the exponential distribution, the values are 1.472 and 1.393, respectively. Thus, in setting the offers to maximise expected information gain, the auctioneer has to forgo a significant amount of revenue (5%), and this is particularly so in the case of exponential valuation distributions.

To justify this action, subsequent auctions must be able to make up the difference. However, figure 2 also shows that the revenue generated by the second auction is only slightly affected by differences in the parameter estimates used (i.e. at the fourth significant figure). That is, the refinements to the prior produced by using information gain maximising offers in the first auction do not generate enough revenue in the second auction than if revenue maximising offer levels had been used in both auctions. Furthermore, even if the extremely small benefit is maintained in subsequent auctions, unless a very long horizon is used (i.e. > 40 repetitions, without discounting), the future benefits of more rapid learning performance will not overcome the revenue foregone in the very first auction. On the other hand, the results presented above indicate that setting offers in TLAs to maximise revenue does not significantly detract from learning performance.

7. CONCLUSIONS

In this paper we extended an existing model of TLAs to consider uncertainty in the value of parameters describing the bidders' valuation distribution. We derived two criteria; one that maximised the expected revenue of the auction given this parameter uncertainty, and one that maximised the expected information gained about the unknown parameter values by holding the auction. We used the criteria to test whether, by adopting a strategy of learning the unknown parameters more quickly, an auctioneer could increase the long-term revenue generated by the entire series of auctions.

Our results show that a TLA optimised to earn revenue reveals close to the same amount of information as a TLA designed specifically for this purpose, so also allows the auctioneer to come close to learning as rapidly as if they had optimised the auction specifically to learn. However, the benefits of learning more quickly early in a series of auctions do not manifest themselves in significantly more revenue. As such, the revenue forgone by an auctioneer who implements offer to maximise the expected information gain of a TLA is not recouped by any additional revenue in future auctions.

8. ACKNOWLEDGMENTS

This research was undertaken as part of the ALADDIN (Autonomous Learning Agents for Decentralised Data and Information Systems) project and is jointly funded by a BAE Systems and EPSRC (Engineering and Physical Research Council) strategic partnership (EP/C548051/1).

9. REFERENCES

[1] J. M. Bernardo. Expected information as expected utility. *The Annals of Statistics*, 7:686–690, 1979.

[2] J. Bredin and D. C. Parkes. Models for truthful online double auctions. In *Proceedings of the 21th Annual Conference on Uncertainty in Artificial Intelligence (UAI-05)*, pages 50–59, Arlington, Virginia, 2005. AUAI Press.

[3] V. Conitzer and N. Garera. Online learning algorithms for online principal–agent problems (and selling goods online). In *Proceedings of the 23rd International Conference on Machine Learning (ICML '06)*, 2006.

[4] D. J. C. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2003.

[5] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[6] D. Pardoe, P. Stone, M. Saar-Tsechansky, and K. Tomak. Adaptive mechanism design: A metalearning approach. In *Proceedings of the Eighth International Conference on Electronic Commerce (ICEC '06)*, 2006.

[7] A. Rogers, E. David, J. Schiff, S. Kraus, and N. R. Jennings. Learning environmental parameters for the design of optimal English auctions with discrete bid levels. In H. L. Poutré, N. Sadeh, and J. Sverker, editors, *Agent-mediated Electronic Commerce, Designing Trading Agents and Mechanisms: AAMAS 2005 Workshop, AMEC*, pages 1–15, Utrecht, Netherlands, July 25 2005. Springer.

[8] T. Sandholm and A. Gilpin. Sequences of take-it-or-leave-it offers: Near-optimal auctions without full valuation revelation. In *Proceedings of AAMAS'06, May 8-12, Hakodate, Hokkaido, Japan*, 2006.