Optimal Design of English Auctions with Discrete Bid Levels

ESTHER DAVID, ALEX ROGERS, and NICHOLAS R. JENNINGS
University of Southampton
JEREMY SCHIFF and SARIT KRAUS
Bar-Ilan University
and
MICHAEL. H. ROTHKOPF
Rutgers Business School and RUTCOR: the Rutgers Center for Operations Research

This article considers a canonical auction protocol that forms the basis of nearly all current online auctions. Such *discrete bid* auctions require that the bidders submit bids at predetermined discrete bid levels, and thus, there exists a minimal increment by which the bid price may be raised. In contrast, the academic literature of optimal auction design deals almost solely with continuous bid auctions. As a result, there is little practical guidance as to how an auctioneer, seeking to maximize its revenue, should determine the number and value of these discrete bid levels, and it is this omission that is addressed here. To this end, a model of an ascending price English auction with discrete bid levels is considered. An expression for the expected revenue of this auction is derived and used to determine numerical and analytical solutions for the optimal bid levels in the case of uniform and exponential bidder's valuation distributions. Finally, in order to develop an intuitive understanding of how these optimal bid levels are distributed, the limiting case where the number of discrete bid levels is large is considered, and an analytical expression for their distribution is derived.

Categories and Subject Descriptors: I.2.11 [**Artificial Intelligence**]: *The Multiagent systems*

General Terms: Design, Economics, Theory

Additional Key Words and Phrases: Discrete bids, English auction, optimal auction design

This research was partially funded by the DIF-DTC project (8.6) on Agent-Based Control and the ARGUS II DARP (a collaborative project involving BAE SYSTEMS, QinetiQ, Rolls-Royce, Oxford University and Southampton University, funded by the industrial partners together with the EPSRC, Ministry of Defence and Department of Trade and Industry). Sarit Kraus is also affiliated with the University of Maryland Institute for Advanced Computer Studies (UMIACS) and this work was supported in part by NSF Grant IIS-0208608.
Author's address: A. Rogers, Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK; email: acr@ecs.sofon.ac.uk.
A preliminary version of this work appeared in David et al. [2005].

ACM Reference Format:
David, E., Rogers, A., Jennings, N. R., Schiff, J., Kraus, S., and Rothkopf, M. H. 2007. Optimal design of English auctions with discrete bid levels. ACM Trans. Intern. Tech. 7, 2, Article 12 (May 2007), 34 pages. DOI = 10.1145/1239971.1239976 http://doi.acm.org/10.1145/1239971.1239976

1. INTRODUCTION

Recent years have seen a rapid increase in the number of online auction sites that allow both private individuals and businesses to trade goods within a virtual worldwide market (prominent examples included eBay, uBid, and Yahoo!). While there are many minor implementation differences between these online auctions (e.g., the availability of proxy bidding services, the use of a fixed or variable auction duration, and the ability to set both starting and reserve prices), these auctions have been modelled on real-world counterparts, and thus, in general, they all share two common features [Lucking-Reiley 2000]. First, they are predominantly based on the ascending price English auction, whereby bidders submit bids to an auctioneer in an open fashion and the auction price increases until no bidder is willing to bid higher[1]. Second, they typically exhibit discrete bid levels, whereby the bids that the bidders may submit within the auction are restricted to certain levels either through a minimum bid increment that the next bid must exceed (as in eBay) or by forcing the auction price to increment through a set of predetermined price levels (as in the popular Israeli auction site `www.olsale.com`).

In contrast, the academic literature on auction theory has almost solely considered auctions in which the bid increment is continuous, and thus bidders may submit extremely small increments in order to outbid the current highest bidder. As such, it has typically been assumed that neither the bidders nor the auctioneer have any time constraints and that bidding is not a costly process. However, the prevalence of the discrete bid protocols within both real-world and online auctions challenges both these assumptions. More significantly, the existence of discrete bid levels causes many of the well-known results from the continuous bid auction literature to fail. For example, the bidders within the auction no longer have a dominant bidding strategy as they must decide whether or not to bid at each bid level [Yu 1999]. In addition, as the item is no longer guaranteed to be allocated to the bidder with the highest valuation, the *revenue equivalence theorem*[2] no longer holds, and thus the revenue of the

[1]The multiple-item Dutch auctions of eBay and the sealed bid Name-Your-Own Price auctions of Priceline are prominent exceptions to this rule.

[2]The revenue equivalence theorem states that all feasible efficient auction protocols will yield the same revenue at equilibrium. For example, consider a group of bidders with private valuations drawn from a common distribution attempting to buy an item within an auction. If the auctioneer implements a second-price sealed bid auction, then there exists a Nash equilibrium in dominant strategies at which each bidder truthfully bids their valuation. Alternatively, if the auctioneer implements a first-price sealed bid auction, then there exists a Bayes-Nash equilibrium at which each bidder shades their bid by an amount dependent on their common beliefs regarding the number of bidders participating and the distribution that describes their valuations. In both cases, the auctions' outcomes are efficient (i.e. they are guaranteed to allocate the item to the bidder with the highest valuation), and thus both auction protocols generate the same expected revenue for the auctioneer.

auction will be dependent on the specific implementation details such as the number and distribution of the discrete bid levels [Chwe 1989].

Thus, despite the widespread use of discrete bid levels, the standard academic auction literature provides little insight or guidance for an auctioneer attempting to maximize its revenue. In real-world auctions, an auctioneer typically uses a combination of intuition and historical experience to adjust the bid increments (and hence the discrete bid levels through which the auction price increases) during the course of the auction [Cassidy 1967]. However, since our ultimate goal is to automate the configuration of online auctions such human interventions are not applicable, and we require a more theoretical understanding of how the discrete bid levels affect the properties of the auction. What little work has been done in this area has addressed the question for very limited cases (see Section 2 for more details). For example, Rothkopf and Harstad [1994] considered several cases where the number of bidders or the number of discrete bid levels was restricted to two. In the case of two bidders with valuations that are independently drawn from a uniform distribution, they showed that it was optimal to use a fixed bid increment with evenly spaced bid levels. However, it proved difficult to generalize these results to instances with larger numbers of bidders whose valuations were drawn from arbitrary distributions.

Against this background, it is our aim to address this lack of guidance. Specifically, we seek to determine the optimal auction design for English auctions with discrete bid levels. As described previously, this represents a canonical auction protocol on which nearly all current online auction protocols are based[3]. In particular, we aim to determine both the reserve price of the auction and the number and distribution of the discrete bid levels that yield the maximum auction revenue in the general case of an arbitrary number of bidders whose valuations are drawn from arbitrary distributions. In so doing, we extend the state-of-the-art in this area in four key ways.

(1) We consider the same model of an ascending price auction with a bounded number of discrete bid levels that was proposed by Rothkopf and Harstad [1994]. But rather than considering particular instances with limited numbers of bidders or bid levels, we derive for the first time a general expression for the expected revenue of the auction. This expression relates the expected auction revenue to the specific discrete bid levels used in that auction and is valid for any number of bidders and any distribution of bidders' private valuations.

(2) We demonstrate how this expression is used to determine the optimal bid levels analytically, and, in addition, we present an algorithm to calculate them numerically. In order to compare our results with the earlier work of Rothkopf and Harstad, we consider two example cases where the bidders' valuations are drawn independently from a uniform and an exponential distribution. In the case of the uniform distribution, we prove that when there

[3]To fully characterise a specific implementation of this auction protocol, it is also necessary to consider all of the additional implementation details. For example, see Rogers et al. [2007] for a full discussion of how the proxy bidding system, the minimum bid increment, and the fixed auction closing time combine to affect the properties of the eBay auction protocol.

are more than two bidders participating within the auction, a decreasing bid increment is optimal and thus the interval between bid levels decreases with each bid level. For the first time, we are able to calculate both analytically and numerically how this decrease should proceed for any number of bid levels and for any number of bidders. In addition, in the case of the exponential distribution, we are able to calculate numerically the optimal distribution of discrete bid levels.

(3) Building on this analysis, we extend the initial auction model to consider two additional cases that extend its scope and realism. First, we consider the more realistic case that the number of bidders within the auction is not a known fixed value but is described by a Poisson distribution whose mean the auctioneer knows (or can estimate). Second, we explicitly include within the model an expression that describes the auctioneer's costs (e.g., an incremental cost for each bid level that the auction progresses through). As before, we derive expressions for the expected revenue of the auction in both cases and numerically solve for the optimal discrete bid levels in the case of uniform and exponential bidders' valuation distributions.

(4) Finally, in order to develop an intuition into the optimal distribution of the discrete bid levels, we consider the distribution of these discrete bid levels when their number approaches infinity (and assuming that in this limit the bid levels get closer and closer together). In this case, we are able to derive an analytic expression that describes the density of the discrete bid levels. We show that this expression is similar (but not identical) to the distribution of the expected closing price of the auction. However, we show that the later distribution (which is easier to estimate from historical auction data) can serve as a good estimate for the former distribution.

The remainder of the article is organized as follows. Section 2 presents related work, Section 3 introduces the initial auction model that we consider, Section 4 derives a general expression for the expected revenue of the auction, and, in Section 5, this is used to show how the optimal bid levels can be derived analytically and determined numerically. Section 6 extends the initial model to cover the two new cases discussed previously, and we use our numerical algorithm to calculate the optimal discrete bid levels in these cases. Section 7 considers the limiting behavior of these discrete bid levels and derives an analytical expression for their density. Finally, Section 8 concludes and discusses future work.

2. RELATED WORK

The problem of optimal auction design has been studied extensively for the case of auctions with continuous bid increments and independent private valuations [Riley and Samuelson 1981; Myerson 1981]. In such auctions, the *revenue equivalence theorem* states that all feasible efficient auctions generate the same revenue, thus the interesting design question concerns the reserve price of the auction (i.e., in continuous English auctions, the price at which the bidding commences). In general, setting a reserve price increases the revenue of the auction and, thus, optimal auction design has typically been concerned with finding the reserve price that maximizes the expected revenue of the auctioneer.

A key finding in this respect is that the optimal reserve price of the auction is independent of the number of bidders and is solely dependent on the bidders' valuation distribution.

In contrast to the literature of continuous bid auctions, the case of discrete bid levels has received little attention although some preliminary works exists. Much of this work is based on the assumption that there is a fixed bid increment and thus the price of the auction ascends in fixed-size steps [Yamey 1972; Chwe 1989; Yu 1999; Bapna et al. 2002, 2003]. In more detail, Yamey [1972] first considered this scenario and commented that such bidding rules appear to have the effect of speeding up the auction proceedings and hence reduce the costs of both the auctioneer and the bidders. He concluded that if the fixed bid increment is small, the expected revenue of the auction will approximate the second highest price.

Chwe [1989] also assumed fixed bid increments but considered a first-price sealed bid auction where bidders' independent valuations were uniformly distributed. He showed that a symmetric unique Nash equilibrium bidding strategy exists and that this equilibrium converges to the equilibrium of the continuous bid auction as the bid increment reduces to zero. In addition, he showed that the expected revenue of the discrete bid auction is always less than that of the equivalent continuous bid auction. Thus, the auctioneer has an incentive to make the bid increments as small as possible, assuming that the time and communication costs of the bidding can be ignored.

Yu [1999] also considered auctions with fixed bid increments but studied each of the four common auction protocols: the first-price sealed-bid, second-price sealed-bid, English, and Dutch auctions. Extending Chwe's result, Yu showed that in each of the auction protocols a symmetric pure strategy equilibrium exists. Specifically, no dominant strategy was identified for the English protocol. However, for each of the protocols, she proved that as the number of bid levels become very large (i.e., the bid increment becomes small), the equilibrium bids converge to those of the corresponding continuous bid auction.

In contrast to this work, Rothkopf and Harstad [1994] considered the more general question of determining the optimal number and value of the bid levels. The authors provided a full discussion of how the discrete bid levels affect the expected revenue of the auction, and they considered two different distributions for the bidders' private valuations: a uniform and an exponential distribution. 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 and any number of bidders. In the first instance, even spacing of bid levels (i.e., a fixed bid increment) was found to be optimal. In the second instance, the optimal bid increment was shown to decrease as the auction progressed (this decrease was described analytically). In the case of the exponential distribution of bidders' valuations, the instance of just two bidders was again considered, and the optimal bid increment was shown to increase as the auction progressed.

In this article, we extend the work of Rothkopf and Harstad [1994]. We initially consider the same model of the ascending price auction but derive the optimal bid levels in the general case with any distribution of bidders' valuations, any number of bid levels, and any number of bidders. Moreover, we then extend this model to incorporate the more realistic case that there is uncertainty in the number of bidders who may enter the auction. In addition, we explicitly consider the costs of the auctioneer, and, in both cases, we are able to determine optimal bid levels.

Our work is also closely related to recent results that have been presented in the context of auctions with severely bounded communication, that is, auctions in which the bidders must communicate their bid using just $t = \log_2(k)$ bits, and thus, can express just k discrete bid values. In this context, Blumrosen et al. [2007] have presented protocols based on modified priority games that are superficially similar to sealed bid auctions. They have shown that in the case that bidders have an a priori priority, they can achieve a dominant strategy equilibrium whereby each bidder adopts a set of valuation intervals (each one associated with one of the k bids), and then bids the interval into which their valuation falls. When compared to the optimal case, they show that these games incur a loss of revenue of order $\frac{1}{k^2}$. Similarly, Kress and Boutilier [2004] have considered a class of incremental limited-precision auctions which are more similar to the ascending price auctions that we consider here. In order to enforce dominant bidding strategies, they introduce an auction closing rule that is similar to our own, and using a bid increment that decreases over time, they empirically demonstrate their protocol generates revenues close to the optimal. However, they do not formally analyse the expected revenue of their protocol but it appears that the analysis that we present here for the ascending price English auction would be directly applicable within their auction protocol.

3. AUCTION MODEL

Initially, we consider an auction model where n risk-neutral bidders are attempting to buy a single item from a risk-neutral auctioneer. The bidders have independent private valuations, v_i, drawn from a common continuous probability density function, $f(v)$ within the range $[\underline{v}, \overline{v}]$. This probability density function has a cumulative distribution function, $F(v)$, and, with no loss of generality, we can state that $F(\underline{v}) = 0$ and $F(\overline{v}) = 1$. These 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, $l_0 < l_1 < \ldots < l_m$, and we note that the value of m explicitly bounds the time and costs of the auctioneer (in Section 6, we extend this model to relax this constraint and explicitly consider the costs of the auctioneer). At this point, we make no constraints on the actual number of these bid levels nor on the intervals between them.

In the work of Rothkopf and Harstad [1994], a standard oral English auction was considered. That is, the auctioneer proposes each bid level, and the first bidder to indicate to the auctioneer his willingness to bid this amount becomes the current highest bidder[4]. In traditional English auction houses, this

[4]This auction protocol is implemented in many online auctions including the Israeli site `www.olsale.com`. In this setting, bidders are presented with a bid button that advances the auction price by an amount predetermined by the auctioneer each time it is clicked.

indication is normally accomplished by raising a paddle or by a prearranged signal to the auctioneer. However, as discussed earlier, there is no dominant bidding strategy within this protocol, and bidders must strategize over whether or not to bid at each bid level. To simplify their analysis, Rothkopf and Harstad assumed that the bidders did not attempt to strategize but instead adopted the simple policy of *pedestrian bidding*, that is, bidders sequentially raised the bid price through the discrete bid levels until their own private valuation was exceeded. Indeed, they showed that in the case of two bidders whose valuations are drawn from any nonincreasing distribution, such as the uniform and exponential distributions considered here, this policy is an equilibrium.

However, in our work, we modify this standard auction protocol to improve its applicability within online auctions. Thus, under our protocol, the auction commences with the auctioneer announcing the first discrete bid level, and all the bidders have a fixed predetermined time interval in which to indicate their willingness to pay this bid level. Having received indications from all willing bidders, the auctioneer then randomly selects one of these bidders and nominates this bidder as the provisional winner. The auction proceeds with the price ascending through the discrete bid levels proposed by the auctioneer, and a provisional winner is randomly selected at each level until either just one bidder is willing to pay the offered bid price or no bidders are willing to pay the offered bid price[5]. At this time, the auction closes. If there is just a single bidder willing to pay this bid level, then this bidder is the winner. If there are no bidders willing to pay this bid level, then the item is allocated to the provisional winner from the previous bid level. To ensure that bidders do not need to strategize over whether or not to bid at any level, we introduce an additional clearing rule. Should a bidder find that they are the only bidder willing to pay the current bid level, and they were also nominated as the provisional winner at the previous bid level, then the price they pay is that of the previous bid level (i.e., bidders do not pay more when they find that they have needlessly outbid themselves)[6].

Our modified auction protocol has three key properties which make it particularly attractive within online settings where bidders are increasingly likely to be automated trading agents rather than humans. First, unlike the standard oral auction, bidders within our protocol have a simple weakly dominant bidding strategy; they should continue to participate in the auction, and thus bid at each bid level until the current bid level exceeds their private valuation. Note that this weakly dominant bidding strategy is particularly attractive within online auction settings since each bidders' strategy in unaffected by assumptions regarding the rationality of the other bidders. In addition, the bidders' strategies are also unaffected by the knowledge of whether or not they have been nominated as the provisional winner at each level. Thus, this nomination

[5]This protocol is somewhat similar to the Japanese auction except that these auctions typically operate with a continuous bid increment that precludes the possibility that there is no unique highest bidder [Cassidy 1967].

[6]In their study of incremental limited-precision auctions, Kress and Boutilier [2004] employ a more complex mechanism whereby one bidder is randomly selected and is held out from the next round of bidding. Their mechanism and ours are interchangeable, however, we believe that ours is easier to implement in real online auctions.

need not be made public and need not actually be made until the final winner of the auction is determined (in the case that there are multiple bidders to select between), or the payment of the winner is determined (in the case that there is just a single bidder willing to pay the current bid level)[7]. Second, as the rounds of the auction have a predetermined and fixed duration, there is no advantage in attempting to submit a bid earlier than an opponent, and thus, bidders with greater computational or communication resources cannot gain an unfair advantage. Third, bidders may enter and leave the auction at any time and need not be present at the commencement of the auction (an important consideration where online auctions are subject to communication drop-outs)[8]. Crucially, however, the analysis of how an auction using our modified protocol closes at a particular discrete bid level (presented in the next section) is identical to the analysis of the standard oral auction performed by Rothkopf and Harstad [1994]. The extended results that we show are all still applicable within this original model.

4. THE AUCTION REVENUE

In order to calculate the optimal bid levels, we must first find an expression for the expected revenue of the auctioneer, given the specific discrete bid levels used in that auction. Following the work of Rothkopf and Harstad [1994], we can describe the probability of the auction closing at any particular bid level by considering three exhaustive and mutually exclusive cases. These three cases are shown in Figure 1, and they describe all the possible configurations of bidders' valuations that lead to the auction closing at a bid level of l_i. In the diagram, the valuations of the bidders are shown as circles, and the arrows indicate which bidder was nominated as the current highest bidder at each bid level. We can describe each case as follows.

Case 1. Two or more bidders have valuations greater than bid level l_i, but none of these bidders have valuations greater than l_{i+1}. Thus, once the bid price has reached l_i, no bidder is willing to increase the bid any further, and the item is allocated to the current highest bidder. In this case, the revenue earned by the auctioneer is less than that which would have been earned in a continuous auction (i.e., the second-highest valuation), and the outcome may be inefficient as the item is not necessarily allocated to the bidder with the highest valuation.

Case 2. Two or more bidders have valuations between l_i and l_{i+1} and a single bidder has a valuation greater than l_{i+1}. As this single bidder was also nominated when the bid level reached l_i, none of the other bidders have valuations sufficient to raise the bid to l_{i+1}. Thus, the auction closes at the price l_i, and the

[7] Making the nomination public does improve the transparency of the auction since all participants can understand the outcome of the auction and the amount that the winner pays. This is often an important reason for using an open ascending price auction as opposed to a sealed bid auction. However, many successfully online auctions (most notably the proxy bidding service of eBay) operate in a nontransparent way (since the amount that the winner pays is determined by the bid of the second-highest bidder, and this is not disclosed during the course of the auction).

[8] Although clearly they may miss the opportunity to buy the item at a low price should there be few other bidders present.

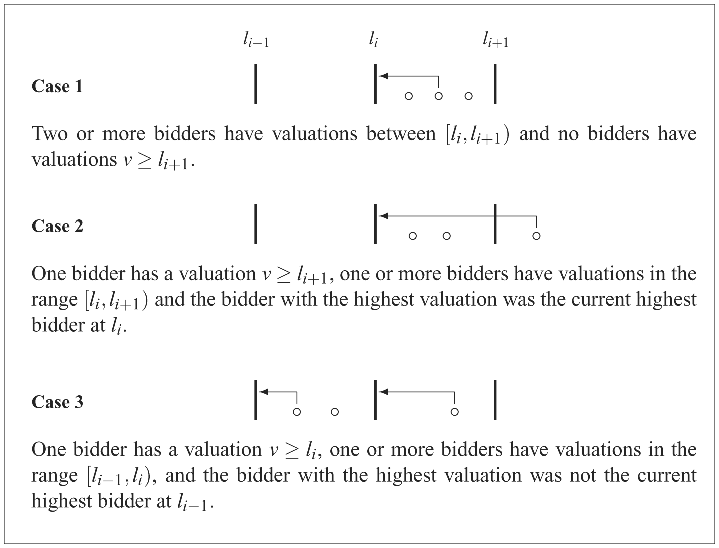

Fig. 1. Diagram showing the three cases whereby the auction closes at the bid level $\mathbf{l_i}$. In each case, the circles indicate a bidder's private valuation, and the arrow indicates the bid level at which that bidder was selected as the current highest bidder.

item is allocated to the bidder with the highest valuation. Again, the revenue earned by the auctioneer is less than that which would have been earned in a continuous auction, but the outcome is allocatively efficient.

Case 3. Two or more bidders have valuations between l_{i-1} and l_i, a single bidder has a valuation greater than l_i, but, unlike. Case 2, this bidder was not nominated when the bid level reached l_{i-1}. Thus, this bidder is forced to raise the bid level, and the auction closes at l_i rather than at l_{i-1}. Again this case is allocatively efficient, however, the revenue earned by the auctioneer is actually greater than that earned in a continuous auction.

The expected revenue of the auction is dependent on the probability of each of these three cases occurring. Each of these probabilities can be described in terms of the cumulative distribution function of the bidders' valuations, $F(v)$. We write $P(\text{Case1}, l_i)$ for the probability that Case 1 occurs, and that the auction closes at bid level l_i. This probability can be computed by considering the probability of having k bidders with valuations between bid levels l_i and l_{i+1} (this happens with probability $[F(l_{i+1}) - F(l_i)]^k$) while the other $n-k$ bidders have valuations below l_i (this happens with probability $F(l_i)^{n-k}$). Summing over all possible values of k gives:

$$P(\text{case1}, l_i) = \sum_{k=2}^{n} \binom{n}{k} F(l_i)^{n-k} [F(l_{i+1}) - F(l_i)]^k. \tag{1}$$

We can perform a similar calculation for Case 2, where we have k bidders with valuations between l_i and l_{i+1}, one bidder with a valuation greater than l_{i+1}, and $n - k - 1$ bidders with valuations below l_i. In this case, we must also consider the probability that the bidder with the highest valuation is the current highest bidder. Under our assumption that this selection is random, this probability is simply given by $\frac{1}{k+1}$, and thus the whole expression is:

$$P(\text{case2}, l_i) = \sum_{k=1}^{n-1} \binom{n-1}{k} \frac{n}{k+1} F(l_i)^{n-k-1} [F(l_{i+1}) - F(l_i)]^k [1 - F(l_{i+1})]. \quad (2)$$

Finally, we consider Case 3, which is identical in form to Case 2 with the exception that the bidder with the highest valuation was not nominated as the current highest bidder at bid level l_{i-1} and must thus raise the price to l_i. The probability of this occurring is $\frac{k}{k+1}$, rather than $\frac{1}{k+1}$ as in Case 2. Note that this description implies that there exists a bid level below l_i and thus the expression that we derive is only valid for bid levels $l_1, \ldots, l_m$. In order to include the instance in which the auction closes at the bid level l_0, we note that this requires all but one bidder have valuations below l_0. Thus, the final expression is:

$$P(\text{case3}, l_i) = \begin{cases} nF(l_0)^{n-1}[1 - F(l_0)] & i = 0 \\ \sum_{k=1}^{n-1} \binom{n-1}{k} \frac{kn}{k+1} F(l_{i-1})^{n-k-1} [F(l_i) - F(l_{i-1})]^k [1 - F(l_i)] & i > 0. \end{cases} \quad (3)$$

As these three expressions completely describe all the possible ways in which the auction may close at any particular bid level, we can find the expected revenue of the auctioneer by simply summing over all possible bid levels and weighting each by the revenue that it generates, l_i. Thus the expected revenue of the auction is given by:

$$E = \sum_{i=0}^{m} l_i [P(\text{case1}, l_i) + P(\text{case2}, l_i) + P(\text{case3}, l_i)]. \quad (4)$$

The resulting expression at this stage is extremely complex due to the combinatorial sums in Equations (1), (2) and (3). However, as detailed in Appendix A, it is possible to simplify this expression significantly (noting, in so doing, that, with no loss of generality we can define $F(l_{m+1}) = 1$), to give the final result:

$$E = \sum_{i=0}^{m} \frac{F(l_{i+1})^n - F(l_i)^n}{F(l_{i+1}) - F(l_i)} [l_i(1 - F(l_i)) - l_{i+1}(1 - F(l_{i+1}))]. \quad (5)$$

This expression is a key result, and many of the results that we present in this article stem from the fact that we have been able to express the revenue of the auction in a relatively compact form. Unlike previous work that has considered simple instances of the auction, for example, those with just two bidders or two bid levels, this expression is for the general case. It relates the revenue of the auction to the actual bid levels used and is valid for any number of bid levels, any number of bidders, and for any valuation distribution function which is described by $F(v)$. Also, unlike the earlier work, we make no assumptions about the positions of the first and last bid levels. Whereas Rothkopf and Harstad [1994] fixed these at the extremes of the bidders' valuation distribution (i.e., $l_0 = \underline{v}$ and $l_m = \overline{v}$), we make them free parameters and allow them to take any value. Since l_0 is equivalent to the reserve price of the auction, we thus determine the optimal reserve price and the optimal bid levels by the same process.

5. OPTIMAL AUCTION DESIGN

The expression derived in the last section describes the expected revenue of the auction when discrete bid levels $l_0, \ldots, l_m$ are used. Given the constraint of this fixed number of bid levels, our goal is to attempt to determine their actual values such that the revenue of the auctioneer is maximised. Initially we present analytical results applying this methodology to a uniform bidders' valuation distribution. However, since it is not always possible to derive analytical results, we also present a numerical algorithm that is applicable in the general case.

5.1 Analytical Solutions

Now, in order to solve for the discrete bid levels that generate the maximum expected revenue for the auctioneer, we must find the partial derivatives of the revenue expression given in Equation (5) with respect to each individual bid level l_i. We can then solve the equations $\partial E/\partial l_i = 0$ to find the values of l_i that maximize the revenue.

To perform this differentiation, we must note that each l_i occurs in the summation of Equation (5) twice. For example, the bid level l_5 occurs in the summand when $i = 5$, as $F(l_i)$ and also in the preceding term when $i = 4$, as $F(l_{i+1})$. Thus, for a uniform bidders' valuation distribution, we substitute the analytical expression $F(l_i) = \frac{l_i - \underline{v}}{\overline{v} - \underline{v}}$ into these two terms and differentiate to give:

$$\frac{\partial E}{\partial l_i} = \frac{(l_{i+1} - \underline{v})^n - (l_{i-1} - \underline{v})^n}{(\overline{v} - \underline{v})^n} + \frac{n l_{i-1}(l_i - \underline{v})^{n-1} - n l_{i+1}(l_i - \underline{v})^{n-1}}{(\overline{v} - \underline{v})^n}. \qquad (6)$$

In order to find the value of l_i that maximizes the revenue, we can then simply make this partial derivative equal to zero (i.e. $\partial E/\partial l_i = 0$) and solve the resulting expression[9]. Doing so gives:

$$l_i = \underline{v} + \sqrt[n-1]{\frac{(l_{i+1} - \underline{v})^n - (l_{i-1} - \underline{v})^n}{n(l_{i+1} - l_{i-1})}}. \qquad (8)$$

This expression relates any individual optimal bid level to the bid levels on either side of it. Thus, if we consider the specific case where $n = 2$, we can simplify this expression to:

$$l_i = \frac{l_{i-1} + l_{i+1}}{2}. \qquad (9)$$

[9]Note that the second derivative is given by:

$$\frac{\partial^2 E}{\partial l_i^2} = \frac{n(n-1)(l_{i-1} - l_{i+1})(l_i - \underline{v})^{n-2}}{(\overline{v} - \underline{v})^n}. \qquad (7)$$

Since $l_{i+1} \geq l_{i-1}$ and $l_i \geq l_{i-1} \geq \underline{v}$, then $\frac{\partial^2 E}{\partial l_i^2} \leq 0$. Thus, if solving $\partial E/\partial l_i = 0$ yields a solution, then this must be a unique solution that maximizes the revenue.

Thus, the value of l_i is midway between l_{i-1} and l_{i+1}, and as this is true for all l_i, the optimal discrete bid levels are evenly spaced with a fixed bid increment. This results confirms the analysis of Rothkopf and Harstad [1994] who considered exactly this two-bidder case. However, given our general model, we can also consider the case of more bidders (i.e., when $n > 2$), and, in this case, we can show that:

$$l_i > \frac{l_{i-1} + l_{i+1}}{2}. \tag{10}$$

Again this is true for all l_i so the optimal distribution of bid levels consists of a decreasing bid increment, whereby the bid levels become closer together as the auction progresses (see Appendix B for a proof of this result). Thus, perhaps rather surprisingly, the commonly used fixed bid increment is in fact only optimal in one very limited case (i.e., when two bidders with uniform valuation distributions participate).

5.2 Numerical Solutions

When we apply the analytical method previously presented to arbitrary bidders' valuation distributions, we often find that solving Equation (6) is intractable. Thus, since we would like to compare the optimal discrete bid levels in the more general case of arbitrary bidders' valuation distributions, we must adopt a numerical approach to maximizing the expected auction revenue. There are many numerical optimization algorithms available (see Press et al. [1992] for examples), but two key features of this problem guide our choice. First, since each term in the summation in Equation (5) contains only pairs of bid levels (i.e., l_i and l_{i+1}), we note that maximizing this expression or solving $\partial E/\partial l_i = 0$ is equivalent to solving a tridiagonal set of $m+1$ simultaneous equations that by denoting $\partial E/\partial l_i$ as $\mathcal{G}_i$, we can write as:

$$\begin{aligned} \mathcal{G}_0(l_0, l_1) &= 0 \\ \mathcal{G}_i(l_{i-1}, l_i, l_{i+1}) &= 0 \quad \text{for } i = 1 \text{ to } m-1 \\ \mathcal{G}_m(l_{m-1}, l_m) &= 0 \end{aligned} \tag{11}$$

Second, the solutions to these equations are constrained by requiring their ordering to be fixed (i.e., $l_{i-1} \le l_i \le l_{i+1}$). Typically, a general-purpose optimization 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 implementing a version of the Jacobi iteration for solving iteratively a system of $m+1$ simultaneous equations [Hageman and Young 1981]. That is, while fixing all other bid levels, we find the value of l_i that maximizes Equation (5), allowing l_i to vary in the range $l_{i-1} \le l_i \le l_{i+1}$. As shown in the previous section, between these limits, the expression is well behaved and has a single maximum that can be found using hill climbing or any other gradient-based method. We apply exactly the same procedure for the boundary conditions, allowing l_0 to vary in the range $\underline{v} \le l_0 \le l_1$, and l_m to vary in the range $l_{m-1} \le l_m \le \overline{v}$. Thus, we sequentially update all l_i and iterate the process until the bid levels converge to the necessary accuracy.

for i=0:m

$$l_i \leftarrow \begin{cases} \underline{v} + i * (\overline{v} - \underline{v})/m & \text{// uniform} \\ i/m/\alpha & \text{// exponential} \end{cases}$$

$d \leftarrow \infty$

while $d >$ *stopping condition,*

$\quad l'_0 \leftarrow \arg\max_{l_0} E(l_0, \ldots, l_m) \quad$ *where* $\quad \underline{v} \leq l_0 < l_1$

$\quad$ *for i=1:m-1*

$\quad\quad l'_i \leftarrow \arg\max_{l_i} E(l_0, \ldots, l_m) \quad$ *where* $\quad l_{i-1} < l_i < l_{i+1}$

$\quad l'_m \leftarrow \arg\max_{l_m} E(l_0, \ldots, l_m) \quad$ *where* $\quad l_{m-1} < l_m \leq \overline{v}$

$\quad d \leftarrow 0$

$\quad$ *for i=0:m,*

$\quad\quad d \leftarrow \max(d, \text{abs}(l'_i - l_i))$

$\quad\quad l_i \leftarrow l'_i$

Fig. 2. Pseudocode for a numerical algorithm based on Jacobi iteration to calculate solutions for the optimal bid levels with arbitrary bidders' valuation distributions.

We present this numerical algorithm in pseudocode in Figure 2 and note that the expression $E(l_0, \ldots, l_m)$ represents the revenue expression shown in Equation (5). While our purpose here is not to prove the convergence properties of this iterative algorithm, in our experiments, it was found to converge reliably and rapidly, given that a single condition for the initial values of l_i is satisfied. Specifically, no bid level may be greater than the upper limit of the bidders' valuation distribution (i.e., $l_i \leq \overline{v}$). In the first two lines of the algorithm, we provide suitable starting conditions for the two valuation distributions that we consider in the next section. These starting conditions simply uniformly distribute the discrete bid levels over the range where we expect to find bidders' valuations.

5.3 Comparison of Bidders' Valuation Distributions

The numerical solution described in the previous section allows us to calculate the optimal discrete bid levels for any value of n (i.e., the 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 n for two different bidders' valuation distributions; the exponential distribution proposed by Rothkopf and Harstad [1994] and the uniform distribution. To allow us to compare these two directly, we chose their parameters so that the expected closing prices of the auctions are similar. Thus, in the case of a uniform distribution, we consider a range of [0,1] meaning $f(v) = \frac{1}{\overline{v}-\underline{v}}$ and $F(v) = \frac{v-\underline{v}}{\overline{v}-\underline{v}}$, where $\underline{v} = 0$ and $\overline{v} = 1$. For the exponential distribution, we take $f(v) = \alpha e^{-\alpha v}$ and $F(v) = 1 - e^{-\alpha v}$, where $\alpha = 4$. The resulting optimal discrete bid levels are shown in Figure 3 for three different numbers of bidders ($n = 2$, 10, and 20) and over a continuous range from 2 to 20. In both cases, we use 10 bid levels (i.e., $m = 10$) as this makes clear the differences between the two cases[10].

[10]Note that while changing the number of bid levels does affect their value, it does not affect the general form of the distribution seen in the plot.

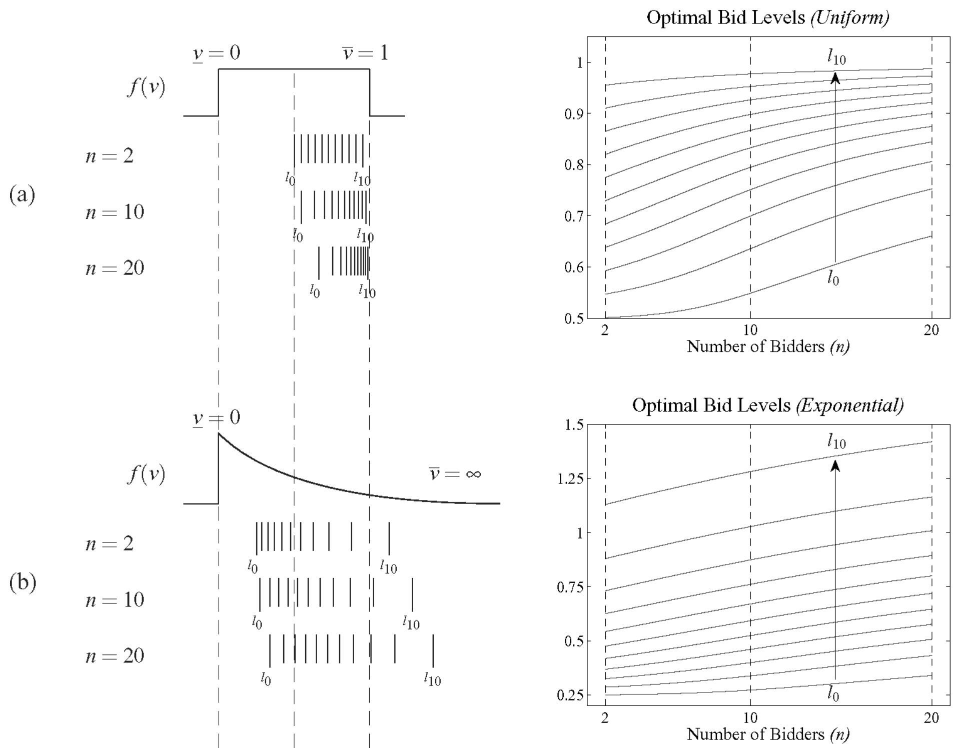

Fig. 3. Optimal bid levels for (a) uniform and (b) exponential bidders' valuation distributions for the initial auction model.

Rothkopf and Harstad [1994] and our preceding analytical analysis proved that when there are two bidders whose valuations are drawn from a uniform distribution, the optimal discrete bid levels are evenly spaced with a fixed bid increment. This same result is observed in our numerical results . In addition, when there are more than two bidders, we observe that the optimal bid levels become closer together, and the bid increment decreases as the bid price increases as proved in the previous section. The case of the exponential valuation distribution is more complex. When there are two bidders, we see an increasing bid increment as was shown by Rothkopf and Harstad. However, as the number of bidders increases, we observe that, rather than increasing, the bid increment initially decreases, reaches a minimum size, and then increases again.

We also observe that, in both cases, as the number of bidders increases, the value of the first bid level, l_0, increases. Rothkopf and Harstad [1994] fixed the values of the first and last bid level at the extremes of the valuation distribution (i.e., for the uniform case, $l_0 = \underline{v}$ and $l_m = \overline{v}$). However, we make no such restriction, and thus the values of l_0 and l_m are optimized at the same time as the other bid levels[11]. Since l_0 is equivalent to the reserve price of the auction

[11]This is particularly important in the case of the exponential distribution where the upper limit of the probability distribution is infinity. We do not force $l_m = \infty$ and thus make more efficient use of the constrained number of bid levels.

(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 continuous bid auctions, the optimal reserve price of an auction with discrete bid levels is dependent on the number of bidders. In general, we see that when the number of bid levels is large, or the number of bidders is small, the value of l_0 tends toward the optimal reserve price of the equivalent continuous bid auction (i.e., given by Riley and Samuelson [1981]: $v^* = [1 - F(v^*)]/F'(v^*)$. For the uniform valuation distribution, $v^* = \max(\underline{v}, \overline{v}/2)$, and for the exponential valuation distribution $v^* = 1/\alpha$).

Naïvely, our intuition would guide us to suggest 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, in the case of the uniform distribution, the bid levels become closer together nearer to the upper limit of the distribution, while in the exponential distribution, they become closer together in the area where we expect to find the bidder with the second-highest valuation. In Section 7, we show that this intuition is only partially correct, and more accurately, we analytically calculate the density of bid levels by considering the limiting case where the number of bid levels becomes large.

5.4 Auction Properties

Having shown that we can derive both numerical and analytical solutions for the optimal bid levels, we consider how these optimal bid levels affect the properties of the auction. We consider three properties (i) the expected revenue of the auction (i.e., the property that we have maximized in the derivation of the optimal bid levels), (ii) the expected duration of the auction (measured in terms of the number of bid levels that the price has been raised through), and (iii) the allocative efficiency of the auction expressed as the probability that the item is sold to the bidder with the highest private valuation. Using these measures, we compare the properties of an auction using optimal discrete bid levels to one implementing fixed bid increments (i.e., where we use the same number of discrete bid levels but simply evenly distribute them between $\underline{v}$ and $\overline{v}$). Note that this comparison can only be made in the case of a bidders' valuation distribution with a finite upper limit (such as the uniform distribution considered here) since otherwise we have no way of determining the actual value of the fixed bid increment to use.

To this end, we consider bidders' private valuations drawn uniformly on $[0, 1]$, and we assume bid levels from l_0 to l_{10}. For each number of bidders in the range 2 to 20, we use our numerical algorithm to find the optimal bid levels and then, given these bid levels, we calculate the expected revenue, duration, and efficiency of the auction[12]. We then compare these measures to those calculated when a fixed bid increment is used and present the results in Figure 4.

[12]The first measure is simply calculated using the revenue expression shown in Equation (5). See Appendix C for details of the calculation of the other two measures.

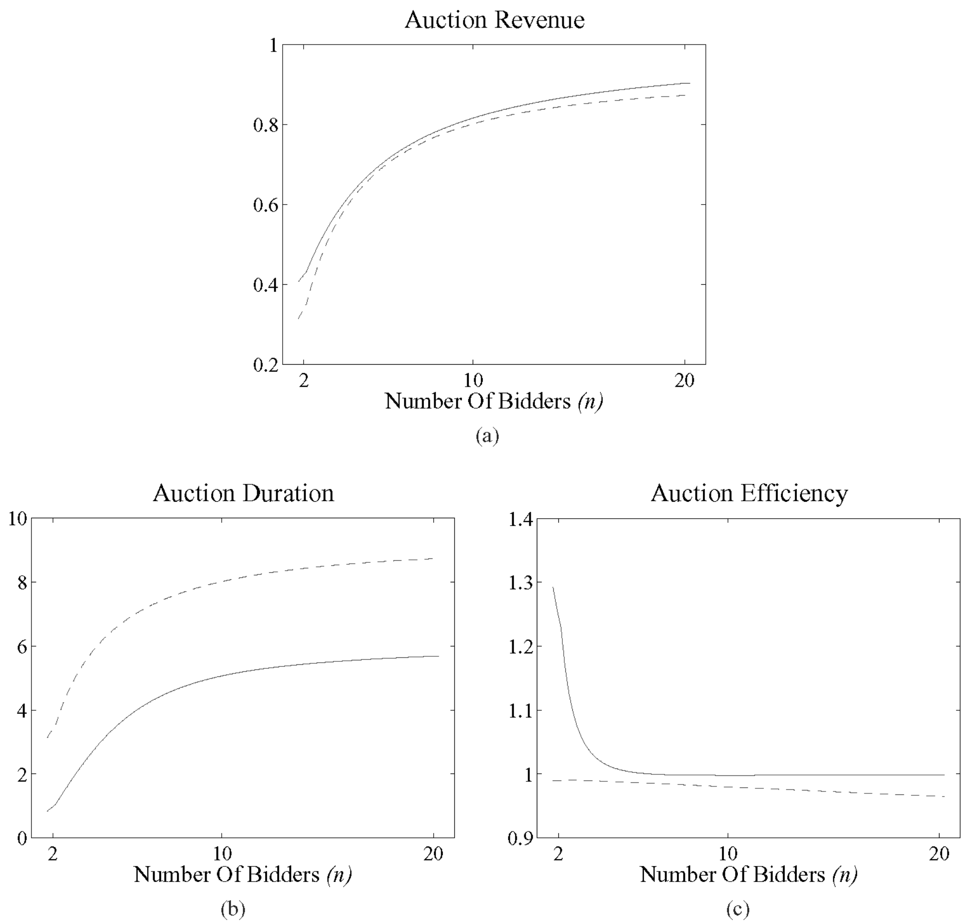

Fig. 4. Comparison of the expected revenue and duration of the auction, and the probability that the bidder with the highest valuation actually wins the auction when optimal bid levels (solid lines) and fixed bid increments (dashed lines) are used. Bidders valuations are drawn uniformly on [0,1] and $m = 10$.

If we first consider the auction with fixed bid increments, we see that the revenue of the auction increases as the number of bidders increases (Figure 4(a)). Thus, the auction closes at a higher bid level, and we also see an increase in the auction duration since bidders must raise the price through more bid levels in order to reach this closing price (Figure 4(b)). We also see a large loss in the allocative efficiency of the auction (Figure 4(c)). As the number of bidders increases, it becomes more likely that the valuations of the highest bidders fall between the discrete bid levels (Case 1 as described in Section 4). The item is then allocated randomly to one of these bidders with the corresponding loss of allocative efficiency and auction revenue.

When optimal discrete bid levels are used, the probability that the highest bidders fall between bid levels is reduced. The auction is thus able to differentiate between the bidders with the highest valuations, and we see a substantial improvement in the allocative efficiency of the auction. As expected, this efficiency improvement also results in an increase in the expected revenue of the auction. However, the increase is relatively small since clearly the most that we would expect to improve the revenue of the auction would be of the same order as the difference between discrete bid levels. In addition, the initial widely-spaced bid increments and optimal reserve price ensure that the bidders do not have to raise the auction price too many times before it approaches the likely closing price. And so we also see a reduction in the expected duration of the auction.

Note that the improvements in auction revenue, duration and efficiency need to be set against the potential disadvantages of an auction protocol that is unfamiliar and more complex to describe (compared to one implementing a fixed bid increment). However, we believe that this disadvantage will be short-lived since in the near future much of the actual bidding within online auctions will be performed by automated bidding agents acting on behalf of their owners [Anthony and Jennings 2003]. Even if this is not the case and a fixed bid increment is actually favored, it is equally important to be able to calculate and understand the loss in revenue that this explicit design choice has incurred.

6. EXTENSIONS TO THE INITIAL AUCTION MODEL

Having derived these results for our initial auction model, we consider two incremental extensions to it that increase its realism and extend its applicability. We first consider the more general setting in which the number of bidders participating in the auction is not fixed but is described by a probability distribution. We then explicitly incorporate a model of the costs of the auctioneer into the revenue calculation. In both cases, we are able to derive an expression for the expected revenue of the auctioneer and thus use the numerical algorithm described in the previous section to calculate representative results.

6.1 Uncertainty in the Number of Bidders

The initial auction model that we considered assumed that the number of bidders in the auction, n, is fixed and known to the auctioneer. In some settings this may be the case, and thus the auction can be designed using this specific knowledge of the number of bidders who will participate. However, in general, this it is not so. It is more likely that while the auctioneer may have an estimate of the number of bidders who will participate, it will be described by a probability distribution[13].

There are a number of candidates for this probability distribution. Levin and Smith [1994] considered an auction model in which the number of bidders participating was endogenously determined. They modeled a pool of potential bidders and showed that at equilibrium each potential bidder has a fixed

[13]Indeed, in the case of the standard oral auction, it is not possible to determine the number of bidders who are participating even once the auction has commenced since, in general, the number of received bids can be less than the number of participants.

probability of actually participating in (or entering) the auction. The number of bidders participating in any auction was thus described by a binomial distribution. More recently, Bajari and Hortacsu [2003] considered a similar model and compared their model to data collected from eBay auctions selling collectible U.S. coins. 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. They deduce that, in such cases, a Poisson distribution is an appropriate approximation for the binomial proposed by Levin and Smith. This observation, was confirmed by Jiang and Leyton-Brown [2005] who compared parameterized models to real eBay auction data and found that the number of bidders within the auctions was well described by such a Poisson distribution.[14]

In light of this work, we describe the number of bidders participating in any auction by a Poisson distribution where the probability that n bidders participate is:

$$P(n) = \frac{\lambda^n e^{-\lambda}}{n!}. \tag{12}$$

The parameter λ describes the mean of this distribution and therefore represents the expected number of participants in any individual auction. Given this distribution, we can extend the expression derived in Section 4 and describe the expected revenue of the auction in terms of the parameter λ rather than n. To do so, we simply sum the product of the expected revenue of the auctioneer, given a fixed number of bidders, E_n, and the probability of that number of bidders actually occurring:

$$E_\lambda = \sum_{n=0}^{\infty} P(n) E_n. \tag{13}$$

Substituting Equations (5) and (12) into this expression and making use of the identity $\sum_{n=0}^{\infty} \frac{\lambda^n}{n!} = e^\lambda$ allows us to derive:

$$E_\lambda = \sum_{i=0}^{m} \frac{e^{\lambda[F(l_{i+1})-1]} - e^{\lambda[F(l_i)-1]}}{F(l_{i+1}) - F(l_i)} [l_i[1 - F(l_i)] - l_{i+1}[1 - F(l_{i+1})]]. \tag{14}$$

As before, we can use this expression within the numerical algorithm presented in Section 5.2 to calculate the optimal discrete bid levels. Again, we compare uniform and exponential valuation distributions, and, in Figure 5, we show these results for a range of values of λ from 2 to 20. In each case, λ represents the expected number of bidders, but the actual number of bidders who participate in an auction is described by the Poisson distribution.

The figure shows that when λ is large, there is little difference between this case and the case of a fixed number of bidders considered in the previous sections. This is simply due to that fact that a large value of λ results in a Poisson distribution with a peak close to λ and a relatively small standard deviation

[14]Note that Jiang and Leyton-Brown [2005] also attempted to infer the presence of bidders who participated but did not have valuations sufficient to allow them to bid.

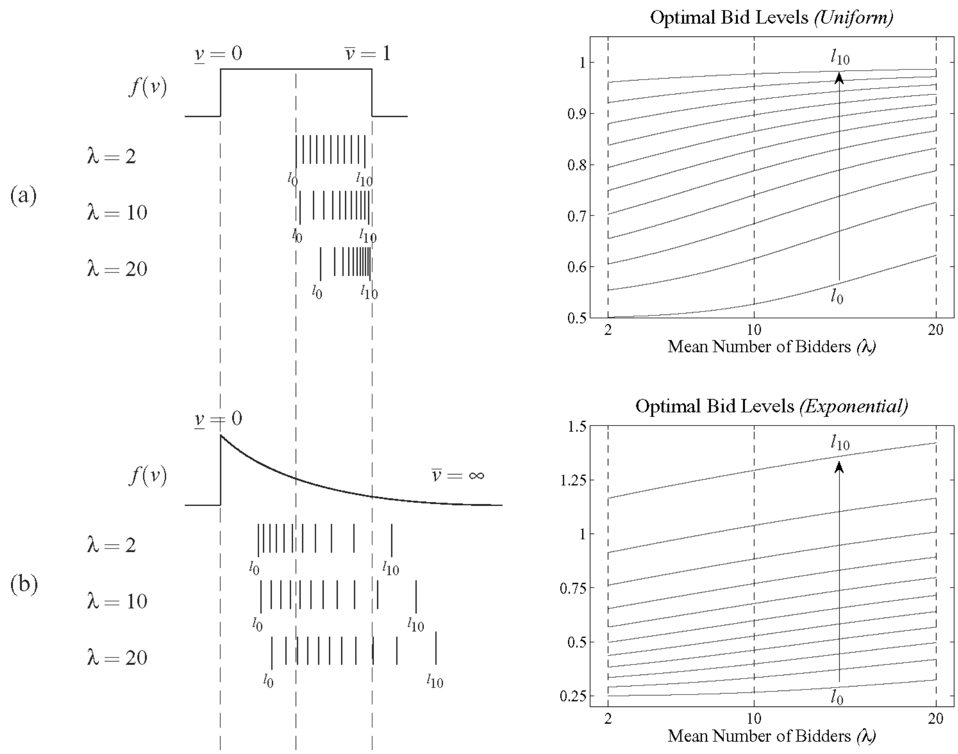

Fig. 5. Optimal bid levels for (a) uniform and (b) exponential valuation distributions where the number of bidders participating within the auction is described by a Poisson distribution who expected value is λ.

$\sqrt{\lambda}$. Thus the results are little different from those with an equivalent value of n. However, when λ is small, the standard deviation of the Poisson distribution is larger relative to the mean value, and there is a significant probability that auctions occur in which the number of bidders differs from the expected value by a large percentage. It is the effect of having fewer bidders that dominates, and therefore, we observe that these optimal discrete bid levels are more evenly spaced, and hence closer to a fixed bid increment. In addition, the value of l_0 is lower, and is closer to the optimal reserve price of the equivalent continuous bid auction.

6.2 Explicit Auctioneer Costs

In the previous analysis, we assumed that a fixed number of discrete bid levels have been implemented by the auctioneer. This fixed number of discrete bid levels places a strict bound on the time and costs of the auction, and our goal has been to calculate the value of these bid levels in order to maximize the expected revenue of the auctioneer. However, we have not explicitly included the costs that each bid level incurs when calculating this revenue. Here we do so by assuming that each bid level that the auction proceeds through costs the

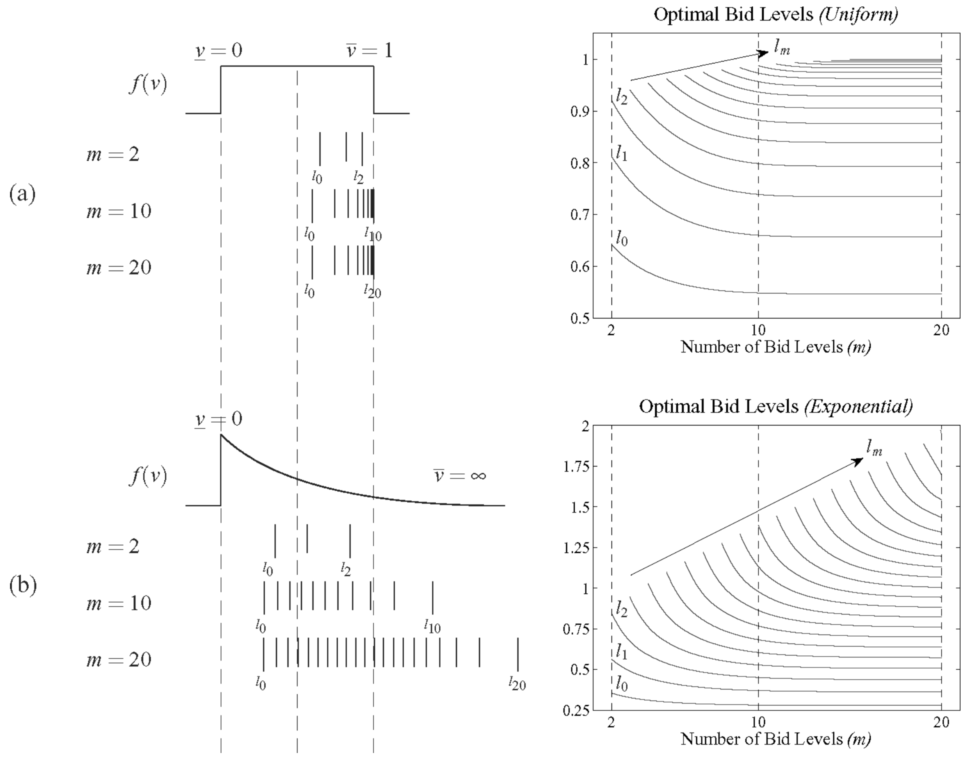

Fig. 6. Optimal bid levels for (a) uniform and (b) exponential valuation distributions for an auction where we explicitly model the costs of the auctioneer. In this case, $\lambda = 10$, and the cost to the auctioneer of the auction progressing through each district bid level, $c = 0.005$.

auctioneer a fixed amount that represents both the time and communication costs of the auctioneer[15]. This cost is denoted by c, and it is now simple to extend the expression for the expected revenue derived in the previous section to deduct this cost, giving:

$$E = \sum_{i=0}^{m} \frac{e^{\lambda[F(l_{i+1})-1]} - e^{\lambda[F(l_i)-1]}}{F(l_{i+1}) - F(l_i)} \times [[l_i - c(i+1)](1 - F(l_i)) - [l_{i+1} - c(i+2)](1 - F(l_{i+1}))]. \quad (15)$$

We can use this expression within our numerical algorithm to calculate the optimal discrete bid levels. In Figure 6, we compare uniform and exponential valuation distributions, but this time, we fix the expected number of bidders and the auctioneer's cost per-bid level (taking $\lambda = 10$ and $c = 0.005$), and we vary the number of discrete bid levels (taking $m = 2$, 10 and 20). Including this explicit cost has more effect than the previous extension to the model. First, it results in a small increase in the value of all the discrete bid levels, since the costs of the auction progressing through these bid levels must be recovered.

[15]Note that each auction could also incur a fixed overhead cost, but since this will not affect the optimal bid levels, we do not incorporate it.

More significantly, while in the previous examples increasing the number of bid levels would have resulted in levels that were closer together and an increase in the expected revenue of the auctioneer, this is no longer true. Now, increasing the number of bid levels could potentially result in a net loss since each level that the auction progresses through incurs an additional cost. Thus, as we increase the number of bid levels, these extra bid levels are positioned so that they are unlikely to affect the outcome of the auction. With the uniform valuation distribution, these bid levels appear close together in the extreme of the distribution (i.e., close to $\overline{v}$). With the exponential distribution, they again appear in the extreme of the distribution, but since this distribution is unbounded (i.e., $\overline{v} = \infty$), they do not appear close together but have incrementally increasing values.

Most notable, however, is that due to this effect, as we increase the number of bid levels that we calculate (i.e., increase m), the values of the first few bid levels rapidly reach static values that do not change. It is the position of these bid levels that determine the expected revenue of the auctioneer (it is extremely unlikely that the bidders will have high enough valuations to continue bidding past these bid levels), and so we can calculate these significant bid levels for an arbitrarily large total number of levels, and in so doing, we remove the parameter m from the design of the auction.

In the results we have presented, we have assumed a relatively small value for this cost per-bid level (i.e., $c = 0.005$). If we decrease the cost, the results approach those presented earlier (i.e., the case in which the auction costs are not calculated directly, and the expected revenue is determined by the strict bound of the number of discrete bid levels, m). At the other extreme, if we increase this cost further, it becomes increasingly difficult to recover the costs of the auction process through the appropriate distribution of discrete bid levels, and thus the expected bid level at which the auction closes approaches l_0 (i.e., the auction closes at the first announced price). At this point, there is effectively no longer an auction as such, and when we calculate the value of l_0, we are effectively calculating the optimal single take-it-or-leave-it price that the auctioneer should announce to the bidders.

7. DISCRETE BID-LEVEL DENSITY FUNCTION

In Section 5, we demonstrated that in our basic auction model, we can calculate the optimal discrete bid levels, using a numerical algorithm for any number of bidders and any bidders' valuation distribution. In addition, we then showed that the same numerical algorithm can be applied to various extensions of the model. In this section, in order to develop an intuitive understanding of these numerical results, we go back to consider what happens in our basic model when the number of discrete bid levels becomes very large. In the previous section, we have seen that when we do this in an extended model with a cost term, the optimal arrangement is for most of the bid levels to be clustered toward the upper end of the valuation range—adding low-bid levels is likely to just increase the cost of the auction without increasing the sale price. However, this is not the case when there is no cost term. For typical valuation distributions, including the uniform and exponential distributions we have considered here, discrete bid auctions yield a lower revenue than continuous bid auctions (since, in the language of Section 4, the losses due to Cases 1 and 2 exceed the gains due to Case 3). Thus we expect that as the number of permitted bid levels m is increased, the optimal bid levels should get closer together. In calculating the details of this effect, we can derive an analytic expression for the density of the discrete bid levels, and gain an intuitive understanding of how they are distributed that can then be applied in more general cases.

Suppose that for some large value of m the bid levels are closely spaced. Then we can find a smooth function $l(t)$, defined over the range $0 \leq t \leq 1$, such that the actual bid levels l_i are equally-spaced samples of this function:

$$l_i = l\left(\frac{i}{m}\right) \qquad \text{for} \quad i = 0, \ldots, m. \tag{16}$$

where the endpoints of the range are given by $l_0 = l(0)$ and $l_m = l(1)$. By substituting this term into the revenue expression given in Equation (5) and expanding in inverse powers of m (see Appendix C for more details on this calculation), we obtain an approximation to the auction revenue:

$$\begin{aligned} E \approx\ & n \int_{l(0)}^{l(1)} F(v)^{n-1}[vF'(v) + F(v) - 1]dv + l(1)(1 - F(l(1))^n) \\ & - \frac{n(n-1)}{12m^2} \int_0^1 F(l(t))^{n-2}[2F'(l(t))^2 + (1 - F(l(t)))F''(l(t))]l'(t)^3 dt. \end{aligned} \tag{17}$$

The terms in the first line of this expression give the revenue in a continuous price auction with n bidders, a starting price $l(0)$ and a maximum possible price $l(1)$ (see e.g., Riley and Samuelson [1981] and note that since it is typically assumed that $F(l(1)) = 1$, the second term vanishes). The terms in the second line give the loss in revenue due to discrete bid levels (note that for common distributions, including the uniform and exponential distributions used previously, the quantity $2F'^2 + (1 - F)F''$ is positive). As noted by Rothkopf and Harstad [1994], the loss due to discrete bid levels is of the order $\frac{1}{m^2}$, and thus as the number of bid levels increases, the revenue of the auction rapidly approaches that of the equivalent continuous bid auction[16].

Thus, in order to maximize the revenue of the auction, it is necessary to minimize the integral on the second line in the previous expression. This is a standard problem in the calculus of variations (and details are again provided in Appendix C) with the result that $l(t)$ must satisfy:

$$\frac{1}{l'(t)} = \frac{[F(l(t))^{n-2}(2F'(l(t))^2 + (1 - F(l(t)))F''(l(t)))]^{1/3}}{C}, \tag{18}$$

where C is a constant[17]. Since $l(t)$ describes the values of the bid levels in terms of the continuous variable t, a plot of $1/l'(t)$ against $l(t)$ represents the

[16]Blumrosen et al. [2007] show a similar result in their analysis of auctions with severely-bounded communication.

[17]Note that all the considerations of this section can be extended to the model of Section 6.1 in which the number of bidders in the auction is not fixed but is a Poisson process with mean λ. The

optimal density of the discrete bid levels as a function of the bid level (all this, we reiterate, is valid when m is large). Thus:

$$\text{Discrete Bid Level Density} \propto [F(v)^{n-2}(2F'(v)^2 + (1 - F(v))F''(v))]^{1/3}. \tag{19}$$

In order to compute the proportionality constant, we need to know the value of the minimum and maximum bid levels, $l(0)$ and $l(1)$. To a first approximation, these can be taken to be the values required to give optimal revenue in a continuous price auction. Thus, as the number of discrete bid levels becomes larger, the maximum bid level is simply the maximum of the valuation distribution[18], and thus $l(1) = \overline{v}$. The minimum bid level, is simply the optimal reserve price of the equivalent continuous auction, v^* and, as shown in Riley and Samuelson [1981], is determined by solving the expression $v^* = [1 - F(v^*)]/F'(v^*)$.

Thus, for the uniform valuation distribution that we used previously, $v^* = \max(\underline{v}, \overline{v}/2)$, and, for the exponential valuation distribution, $v^* = 1/\alpha$.

In Figure 7, we show the resulting density functions for both the uniform and exponential bidders' valuation distributions. For the case of the uniform distribution when $n = 2$, the density of bid levels is constant across the range, indicating as before that in this case a fixed bid increment is optimal. As the number of bidders increases, the density of bid levels at the upper extreme of the distribution increases, confirming the previous numerical results whereby the value of l_0 increases and the bid increment decreases as the auction price increases. In the case of the exponential distribution when $n = 2$, we see a decreasing density as the bid price increases, and as before, we observe an increasing bid increment. However, for larger values of n, we see a clear peak in the discrete bid density that corresponds to the bid increment decreasing, reaching a minimum size, and then increasing again. The position of this maximum bid density increases as the number of bidders increases.

In previous sections, we observed that the density of the bid levels increases in the region where we expect the auction to close (i.e., in the region where we expect to find the bidder with the second-highest valuation). However, the analytical expression we have just derived shows that this intuition is only approximately correct. The probability of the continuous bid auction closing price at any price, v, is given by:

$$\text{Probability of Auction Closing} \propto F(v)^{n-2}F'(v)(1 - F(v)). \tag{20}$$

---

large m approximation to the revenue is:

$$E = \int_{l(0)}^{l(1)} \lambda e^{\lambda(F(v)-1)}[vF'(v) + F(v) - 1]dv + l(1)(1 - e^{\lambda(F(l(1))-1)}) - \frac{\lambda^2}{12m^2}\int_0^1 e^{\lambda(F(l(t))-1)}[2F'(l(t))^2 + (1 - F(l(t)))F''(l(t))]l'(t)^3 dt + O(m^{-3}).$$

The optimal bid level distribution satisfies:

$$\frac{1}{l'(t)} = \frac{[e^{\lambda(F(l(t))-1)}(2F'(l(t))^2 + (1 - F(l(t)))F''(l(t)))]^{1/3}}{C}.$$

[18]Note that for the exponential distribution, this cannot be attained, and it is necessary to consider the first-order correction to get a finite answer for $l(1)$.

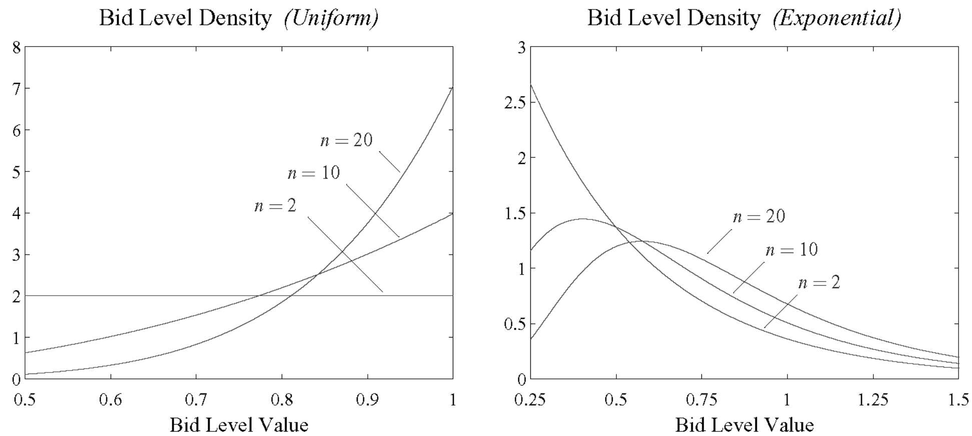

(a) (b)

Fig. 7. Analytically calculated results showing the optimal density of discrete bid levels for the case where bidders valuations are drawn from (a) uniform and (b) exponential distributions (as before, the uniform distribution has support [0, 1] and the exponential distribution is determined by $\alpha = 4$.)

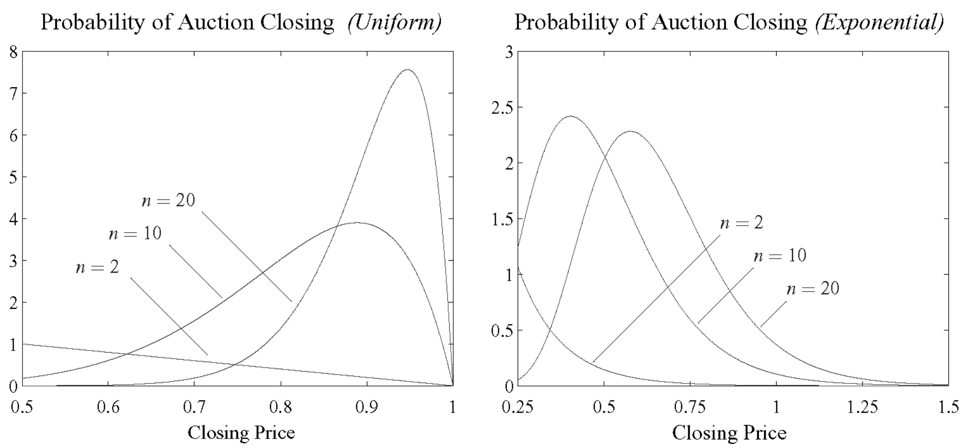

(a) (b)

Fig. 8. Calculated results showing the distribution of the expected continuous bid auction closing price when bidders' valuations are drawn from (a) uniform and (b) exponential distributions (as before, the uniform distribution has support [0, 1] and the exponential distribution is determined by $\alpha = 4$).

Clearly this expression is similar to that shown in Equation (19) with both dominated by the $F(v)^{n-2}$ term. However, the plot of this distribution in Figure 8 shows the differences clearly (and enables a comparison with the optimal bid level density shown in Figure 7). For the uniform distribution, while the probability of the auction closing in the upper extreme of the distribution decreases to zero, it is exactly this area where the density of discrete bid levels reaches a maximum. For the exponential distribution, both expressions simplify and are

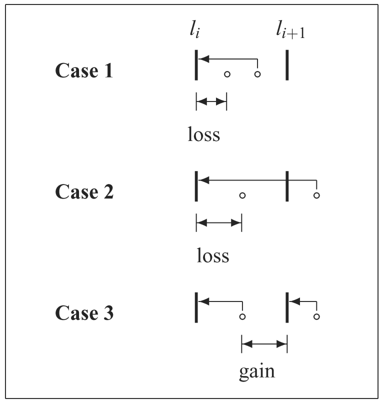

Fig. 9. Diagram showing the three cases in the large m limit, whereby the auctioneer makes a gain or loss, compared to the equivalent continuous auction.

identical apart from an additional 1/3 power. This causes the bid level density function to be much flatter than the auction closing price distribution, and thus the bid levels are more evenly distributed than our initial intuition suggested.

We can intuitively understand the form of the expression for the density of bid levels by reconsidering the three cases that we initially presented in Section 4. Now, in these original case descriptions, we considered the possibility that two or more bidders' valuations fell between any two bid levels. However, as we increase the number of bid levels, we can assume that, at most, two bidders' valuations fall between any two bid levels. Thus, in Figure 9, we can represent the three cases described in Section 4 in terms of the gain or loss that each incurs when compared to a continuous auction which closes at a bid price equal to the valuation of the second-highest bidder. Thus, we see that Case 1 describes an event whereby when two bidders' valuations do in fact fall between two discrete bid levels and that this event generates a net loss in revenue for the auctioneer. The probability of this occurring is described by the first term in the discrete bid level density expression, $2F'(v)^2F(v)^{n-2}$. In addition, we see that now Cases 2 and 3 occur with approximately equal probability, and while Case 2 generates a loss in revenue, Case 3 generates a gain in revenue. The net effect of Cases 2 and 3 depends on the bidder's valuation distribution. In the case of a uniform distribution, Rothkopf and Harstad [1994] noted that the loss of Case 2 exactly cancels the gain of Case 3. However, when the distribution is not uniform (or more accurately when $F''(v) \neq 0$) this does not happen, and the second term in the discrete bid level density expression, $(1-F(v))F''(v)F(v)^{n-2}$ exactly captures the probability of a net gain or loss to the auctioneer due to these two cases. Hence, the discrete bid levels are distributed in order to both minimize an expression that contains two terms: the probability of Case 1 occurring, and the difference in probabilities of Cases 2 and 3.

However, despite the differences between the expressions for the optimal discrete bid level density and the expected auction closing price (shown in Equations (19) and (20) the comparison between them is valuable. It is the probability distribution of the closing price of the auction that is most easily estimated from historical auction data; estimating the number of bidders and the valuation distributions which gave rise to these closing prices is a much harder task. Thus, we can use the expected auction closing price with an additional 1/3 power as an approximation for the density function of the discrete bid levels. In doing so, we can be confident that this additional 1/3 power ensures that the effect of errors in our approximation are likely to be small.

8. CONCLUSIONS AND FUTURE WORK

In this article, we considered a canonical online auction protocol, the ascending price English auction with discrete bid levels. This auction protocol forms the basis of nearly all current online auctions, and our aim was to understand how these discrete bid levels affect the auction properties and also to provide the optimal auction design for this setting.

To this end, we derived a general expression which describes the revenue of the auction in terms of the actual bid levels implemented, the number of bidders participating, and the distribution from which these bidders' valuations are drawn. Using this expression, we showed that, in the case of a uniform valuation distribution, we could derive analytical results that described the distribution of the discrete bid levels. Specifically, we proved that, in the case of two bidders, a fixed bid increment is optimal, while, for greater numbers of bidders, it is optimal to implement a decreasing bid increment so that the interval between bid levels decreases as the auction proceeds.

We then developed a numerical solution that allowed us to calculate the values of the discrete bid levels in the general case, and we compared results for uniform and exponential valuation distributions. We showed that contrary to the standard literature on continuous auctions when discrete bid levels are used, the reserve price of the auction increases as the number of bidders increases. In addition, we then used the same numerical algorithm to consider two extensions to the auction model that increased its applicability and realism. Specifically, we considered the case that the number of bidders participating was not fixed but was described by a probability distribution, and we explicitly included the costs of the auctioneer into the derivation of its expected revenue.

Finally, we considered the limiting case where the number of discrete bid levels became very large and the auction closing price approached that of the equivalent continuous auction. Here we derived an analytical expression for the density of the discrete bid levels and showed that this expression was similar, but not identical, to that describing the expected closing price of the auction. Thus, we suggested that the later distribution (which is easier to estimate from historical auction data) could serve as an estimate for the former distribution, and we noted that the presence of an additional 1/3 power would ensure that the effect of errors in our estimation would likely be small.

Our future work in this area considers two areas. First, we intend to examine and attempt to relax some of the assumptions of the auction model that we consider. In Section 6, we presented some initial work in this area. However, we are particularly interested in the common assumption that all bidders are a priori symmetric (i.e., they all draw their private valuations from the same distribution). This is often a useful simplifying assumption, but it has consequences for the allocation of bid levels. A simple example can illustrate this. Suppose that we have two bidders. In one case, they both draw their values from a uniform distribution on [0,2]; in the other, one draws from a uniform distribution on [0,1] and one from a uniform distribution on [1,2]. If we suppose that we use two bid levels, in the first case, we can calculate their optimal value (in this case, $l_0 = 1.0532$ and $l_1 = 1.5266$) and show that the expected revenue is 0.8142. However, in the second case, where the bidders' valuations are drawn from separate distributions, the expected revenue is greatest if both bid levels are allocated within the range [1,2]. Since there is certainly one bidder with a valuation in this range, an expected revenue of 1 can be guaranteed by simply setting $l_0 = 1$ (the value of l_1 is irrelevant). Thus, clearly the presence of asymmetric bidders has a significant effect on the optimal allocation of bid levels and also on the resulting expected auction revenue. We intend to investigate this further and incorporate it within the framework presented here. A further extension of this model of the bidders would be to consider that, like the auctioneer, they incur a participation cost that is related to the number of discrete bid levels that the auction proceeds through.

Second, we intend to address the question of how an auctioneer should attempt to use the analysis that we have presented to design real-world online auctions. In this respect, online settings are particularly interesting since these typically involve auctioneers who hold many similar repeated auctions. Thus, there is an opportunity for the auctioneer to attempt to use the experience of earlier auctions in order to tailor the discrete bid levels implemented within future auctions. In our preliminary work, we have used Bayesian inference to allow the auctioneer to infer the number and valuation distribution of bidders who participate in a series of repeated auctions [Rogers et al. 2005]. However, building upon the analysis that we presented in the previous section, our aim is to extend this inference technique to the more general setting where rather than deal with the individual bidder's valuation distributions, the auctioneer uses the more easily estimated expected auction closing price to guide his design. Our intention then is to develop efficient methods that automate the design of optimal discrete bid auctions.

APPENDIXES

A. EXPECTED AUCTION REVENUE

Our initial expression for the revenue of the auction was derived by summing over all possible bid levels contributions from the three cases in which the auction closed at bid level l_i, to give:

$$E = \sum_{i=0}^{m} l_i[P(\text{case1}, l_i) + P(\text{case2}, l_i) + P(\text{case3}, l_i)]. \tag{21}$$

In Equations (1), (2) and (3) we presented expressions for these three probabilities. In order to reduce the complexity of the final expression for the revenue of the auction, we must first simplify the combinatorial summations in each of these equations. To do so, we note that by substituting $k - 1$ for k in

Equations (2) and (3) we can then change the indices of the summation in these equations from from $k = 1 \dots n-1$ to $k = 2 \dots n$ to give:

$$P(\text{case1}, l_i) = \sum_{k=2}^{n} \binom{n}{k} F(l_i)^{n-k} [F(l_{i+1}) - F(l_i)]^k \tag{22}$$

$$P(\text{case2}, l_i) = \sum_{k=2}^{n} \binom{n-1}{k-1} \frac{n}{k} F(l_i)^{n-k} [F(l_{i+1}) - F(l_i)]^{k-1} [1 - F(l_{i+1})] \tag{23}$$

$$P(\text{case3}, l_i) = \begin{cases} nF(l_0)^{n-1}[1 - F(l_0)] & i = 0 \\ \sum_{k=2}^{n} \binom{n-1}{k-1} \frac{(k-1)n}{k} F(l_{i-1})^{n-k} [F(l_i) - F(l_{i-1})]^{k-1} [1 - F(l_i)] & i > 0 \end{cases} \tag{24}$$

Then, by noting that $\binom{n-1}{k-1}\frac{n}{k} = \binom{n}{k}$, we can again rewrite Equations (23) and (24) such that the binomial expressions and the indices of the summation in all three equations are identical.

$$P(\text{case1}, l_i) = \sum_{k=2}^{n} \binom{n}{k} F(l_i)^{n-k} [F(l_{i+1}) - F(l_i)]^k \tag{25}$$

$$P(\text{case2}, l_i) = \sum_{k=2}^{n} \binom{n}{k} F(l_i)^{n-k} [F(l_{i+1}) - F(l_i)]^{k-1} [1 - F(l_{i+1})] \tag{26}$$

$$P(\text{case3}, l_i) = \begin{cases} nF(l_0)^{n-1}[1 - F(l_0)] & i = 0 \\ \sum_{k=2}^{n} \binom{n}{k} (k-1) F(l_{i-1})^{n-k} [F(l_i) - F(l_{i-1})]^{k-1} [1 - F(l_i)] & i > 0. \end{cases} \tag{27}$$

We then use the identity $\sum_{k=0}^{n} \binom{n}{k} a^{n-k} b^k = (a+b)^n$ to derive the result that $\sum_{k=2}^{n} \binom{n}{k} a^{n-k} b^k = (a+b)^n - na^{n-1}b - a^n$. Using this result, we can simplify Equations (25) and (26) to give:

$$P(\text{case1}, l_i) = F(l_{i+1})^n - nF(l_i)^{n-1}[F(l_{i+1}) - F(l_i)] - F(l_i)^n \tag{28}$$

$$P(\text{case2}, l_i) = \frac{1 - F(l_{i+1})}{F(l_{i+1}) - F(l_i)} \times [F(l_{i+1})^n - nF(l_i)^{n-1}[F(l_{i+1}) - F(l_i)] - F(l_i)^n] \tag{29}$$

We use a similar procedure for the case of $P(\text{case3}, l_i)$. This case is more complex as we have an additional factor of $k-1$ inside the summation. However, we observe that this factor arises through the differentiation of b^{k-1}, and thus we can derive the identity $\sum_{k=2}^{n} \binom{n}{k} (k-1) a^{n-k} b^k = b^2 \frac{d}{db} [\frac{1}{b} \sum_{k=2}^{n} \binom{n}{k} a^{n-k} b^k]$.

By substituting $(a+b)^n - na^{n-1}b - a^n$ for $\sum_{k=2}^{n}\binom{n}{k}a^{n-k}b^k$ in this expression and differentiating the result, we can show that $\sum_{k=2}^{n}\binom{n}{k}(k-1)a^{n-k}b^k = (a+b)^{n-1}[b(n-1)-a] + a^n$. Using this result in Equation (27) gives:

$$P(\text{case3}, l_i) = \begin{cases} nF(l_0)^{n-1}[1-F(l_0)] & i = 0 \\ \dfrac{1-F(l_i)}{F(l_i)-F(l_{i-1})}[F(l_{i-1})^n - F(l_i)^n + nF(l_i)^{n-1}(F(l_i)-F(l_{i-1}))] & i > 0 \end{cases} \tag{30}$$

We can now substitute Equations (28), (29), and (30) into our expression for the expected revenue of the auction (Equation (21)), to give:

$$\begin{aligned} E = & \sum_{i=0}^{m} l_i \frac{1-F(l_i)}{F(l_{i+1})-F(l_i)}[F(l_{i+1})^n - nF(l_i)^{n-1}(F(l_{i+1})-F(l_i)) - F(l_i)^n] \\ & + \sum_{i=1}^{m} l_i \frac{1-F(l_i)}{F(l_i)-F(l_{i-1})}[F(l_{i-1})^n + nF(l_i)^{n-1}(F(l_i)-F(l_{i-1})) - F(l_i)^n] \\ & + l_0 nF(l_0)^{n-1}(1-F(l_0)) \end{aligned} \tag{31}$$

Clearly, many terms in these expressions cancel each other. The middle terms of each summation are equal and opposite when l_i is between l_1 and l_m. Additionally, the term that is left over from this cancellation (i.e., when $i = 0$) cancels the additional term $P(\text{case3}, l_0)$. This gives the simpler result:

$$\begin{aligned} E = & \sum_{i=0}^{m} l_i \frac{1-F(l_i)}{F(l_{i+1})-F(l_i)}[F(l_{i+1})^n - F(l_i)^n] \\ & + \sum_{i=1}^{m} l_i \frac{1-F(l_i)}{F(l_i)-F(l_{i-1})}[F(l_{i-1})^n - F(l_i)^n]. \end{aligned} \tag{32}$$

Finally, we note that by substituting $i+1$ for i in each term within the second summation and stating with no loss of generality that $F(l_{m+1}) = 1$, we can change the indices of this summation from $i = 1 \ldots m$ to $i = 0 \ldots m$, and combine it with the first summation to give the final result:

$$E = \sum_{i=0}^{m} \frac{F(l_{i+1})^n - F(l_i)^n}{F(l_{i+1})-F(l_i)}[l_i(1-F(l_i)) - l_{i+1}(1-F(l_{i+1}))]. \tag{33}$$

This final expression relates the expected revenue of the auctioneer to the discrete bid levels used in the auction and the cumulative distribution from which the bidders' independent valuations are drawn.

B. PROOF OF OPTIMAL DECREASING BID INCREMENTS

In order to show that the optimal bid levels show a decreasing bid increment when $n > 2$, it is sufficient to show that, in this case:

$$l_i > \frac{l_{i-1} + l_{i+1}}{2}. \tag{34}$$

Thus, using the result from Equation (8), we must show that:

$$\underline{v} + \sqrt[n-1]{\frac{(l_{i+1} - \underline{v})^n - (l_{i-1} - \underline{v})^n}{n(l_{i+1} - l_{i-1})}} > \frac{l_{i-1} + l_{i+1}}{2}. \tag{35}$$

If we define $a = l_{i-1} - \underline{v}$ and $b = l_{i+1} - \underline{v}$, then we must show, for $0 < a < b$, that:

$$\frac{b^n - a^n}{b - a} > n\left(\frac{a + b}{2}\right)^{n-1}. \tag{36}$$

PROOF. If $f(t)$ is a strictly convex function over the interval $[a, b]$, then it follows from Jensen's inequality that:

$$\frac{1}{b - a}\int_a^b f(t)dt > f\left(\frac{a + b}{2}\right). \tag{37}$$

We take $f(t) = nt^{n-1}$. This has $f''(t) > 0$, when $n > 2$, and thus is strictly convex. Substituting into Equation (37) and integrating between the limits gives Equation (36) as required. □

C. CALCULATING AUCTION PROPERTIES

We can calculate the average number of bid levels that any auction proceeds through, t, by simply summing over the probability that the auction will close at any bid level to give:

$$t = \sum_{i=0}^{m}(i + 1)[P(\text{case1}, l_i) + P(\text{case2}, l_i) + P(\text{case3}, l_i)], \tag{38}$$

where $P(\text{case1}, l_i)$, $P(\text{case2}, l_i)$, and $P(\text{case3}, l_i)$ are given in Equations (1), (2) and (3).

We can calculate the probability that the bidder with the highest valuation does win the auction by considering the ways in which this does not happen. In Cases 2 and 3 (see Section 4), the bidder with the highest valuation always wins the auction. However, the bidder with the highest valuation will not win if all the bidders valuations are below l_0 or, in Case 1, when the valuations of the k highest bidders fall between bid levels, but one of the $k - 1$ other bidders is selected as the winner. Thus, we can simply sum the probabilities of these events occurring and subtract this probability from one, to give:

$$P(\textit{Highest Bidder Wins}) = 1 - F(l_0)^n - \sum_{i=0}^{m}\sum_{k=2}^{n}\binom{n}{k}\frac{k-1}{k}F(l_i)^{n-k} \times [F(l_{i+1}) - F(l_i)]^k. \tag{39}$$

D. CONTINUOUS LIMIT DERIVATION

In this appendix, we give some of the details necessary to obtain the results stated in Section 7. We first consider how to find a smooth function $l(t)$, defined over the range $0 \leq t \leq 1$, such that the actual bid levels l_i are equally-spaced samples of this function:

$$l_i = l\left(\frac{i}{m}\right) \qquad \text{for} \quad i = 0, \ldots, m, \tag{40}$$

and the endpoints of the range are given by $l_0 = l(0)$ and $l_m = l(1)$. We take as our starting point the expression for the expected revenue of the auction given in Equation (5) and using the fact that we have defined $F(l_{m+1}) = 1$, we separate out the final term to give:

$$E = \sum_{i=0}^{m-1} \frac{F(l_{i+1})^n - F(l_i)^n}{F(l_{i+1}) - F(l_i)} [l_i(1 - F(l_i)) - l_{i+1}(1 - F(l_{i+1}))] + l_m(1 - F(l_m)^n). \tag{41}$$

Given that we are considering the continuous limit where m is large, we can now express each term within the summand of this expression as a Taylor expansion of our smooth function $l(t)$ in increasing powers of $\frac{1}{m}$. For example, for $F(l_{i+1}) - F(l_i)$ we have:

$$\begin{aligned} F(l_{i+1}) - F(l_i) &= F\left(l\left(\frac{i+1}{m}\right)\right) - F\left(l\left(\frac{i}{m}\right)\right) \\ &\approx \frac{1}{m} F'\left(l\left(\frac{i}{m}\right)\right) l'\left(\frac{i}{m}\right) \\ &\quad + \frac{1}{2m^2}\left(F'\left(l\left(\frac{i}{m}\right)\right) l''\left(\frac{i}{m}\right) + F''\left(l\left(\frac{i}{m}\right)\right) l'\left(\frac{i}{m}\right)^2\right) \\ &\quad + \frac{1}{6m^3}\left(F'\left(l\left(\frac{i}{m}\right)\right) l'''\left(\frac{i}{m}\right) + F'''\left(l\left(\frac{i}{m}\right)\right) l'\left(\frac{i}{m}\right)^3 \right. \\ &\quad \left. + 3F''\left(l\left(\frac{i}{m}\right)\right) l'\left(\frac{i}{m}\right) l''\left(\frac{i}{m}\right)\right). \end{aligned} \tag{42}$$

Since m is large, each term in this Taylor expansion is significantly smaller than the preceding one, and thus, we need only retain the first three terms (i.e., those in $\frac{1}{m}$, $\frac{1}{m^2}$, and $\frac{1}{m^3}$). Repeating this procedure for each of the terms in Equation (41) and then substituting these results back into Equation (41) results in a rather complicated expression for the auctioneer's expected revenue. This expression is difficult to manipulate by hand and is unfortunately too long to present here. However, using Maple to perform the necessary symbolic manipulation, and again keeping just the first three powers of $\frac{1}{m}$, we can simplify it to an expression of the form:

$$\begin{aligned} E &\approx l(1)(1 - F(l(1))^n) + \frac{1}{m} \sum_{i=0}^{m-1} \mathcal{G}_1\left(\frac{i}{m}\right) \\ &\quad + \frac{1}{m^2} \sum_{i=0}^{m-1} \mathcal{G}_2\left(\frac{i}{m}\right) + \frac{1}{m^3} \sum_{i=0}^{m-1} \mathcal{G}_3\left(\frac{i}{m}\right), \end{aligned} \tag{43}$$

where the functions $\mathcal{G}_1(\frac{i}{m})$, $\mathcal{G}_2(\frac{i}{m})$, and $\mathcal{G}_3(\frac{i}{m})$ are functions composed of multiple powers and derivatives of $l(\frac{i}{m})$ and $F(l(\frac{i}{m}))$ that are similar in structure to

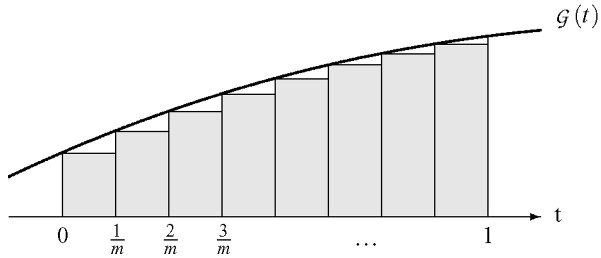

Fig. 10. Graphical interpretation of the Euler-Maclaurin sum formula.

those in Equation (42), and are too long to present here. Now, when m is large, the summations of these functions can be approximated by integrals using the Euler-Maclaurin sum formula, which for our purpose can be conveniently written as:

$$\begin{aligned}\frac{1}{m}\sum_{i=0}^{m-1}\mathcal{G}\left(\frac{i}{m}\right) &= \int_0^1 \mathcal{G}(t)dt - \frac{1}{2m}[\mathcal{G}(1)-\mathcal{G}(0)] + \frac{1}{12m^2}[\mathcal{G}'(1)-\mathcal{G}'(0)] + O\left(\frac{1}{m^4}\right)\\ &= \int_0^1 \left(\mathcal{G}(t) - \frac{1}{2m}\mathcal{G}'(t) + \frac{1}{12m^2}\mathcal{G}''(t)\right)dt + O\left(\frac{1}{m^4}\right) \qquad (44)\end{aligned}$$

Figure 10 shows a graphical interpretation of this result where the shaded rectangles represent the summation $\frac{1}{m}\sum_{i=0}^{m-1}\mathcal{G}(\frac{i}{m})$. This summation can clearly be approximated as the integral of the function $\mathcal{G}(t)$ over the range $0 \le t \le 1$ plus a number of correction terms dependent on the derivatives of $\mathcal{G}(t)$ (for details see, e.g., Eric W. Weisstein.[19] Using these integral approximations within Equation (43) and using Maple to manipulate the expression, allows the expected revenue of the auction to be written as:

$$\begin{aligned}E &\approx n\int_{l(0)}^{l(1)} F(v)^{n-1}[vF'(v)+F(v)-1]dv + l(1)(1-F(l(1))^n)\\ &\quad - \frac{n(n-1)}{12m^2}\int_0^1 F(l(t))^{n-2}[2F'(l(t))^2 + (1-F(l(t)))F''(l(t))]l'(t)^3dt, \qquad (45)\end{aligned}$$

where we have introduced the variable v by noting that $\int_0^1 F(l(t))dt = \int_{l(0)}^{l(1)} F(v)dv$. This change of variable is pertinent since the first line of this expression can now be seen to be equivalent to the revenue of a continuous price auction with n bidders, a starting price $l(0)$, and a maximum possible price $l(1)$ (see, e.g., Riley and Samuelson [1981] and note that since it is typically assumed that $F(l(1)) = 1$, the second term vanishes).

The terms in the second line describe the loss in revenue due to the discrete bid levels, and thus the revenue of the auctioneer is maximized when $l(t)$ is chosen such that the integral here is minimized. To do this, we first note that

[19]Euler-Maclaurin Integration Formulas, from Mathworld—a Wolfram Web resource, `http://-mathworld.wolfram.com/EulerMaclaurinIntegrationFormulas.html`).

this integral can be written in the form:

$$\int_0^1 p(l(t))l'(t)^3 dt. \tag{46}$$

Now, the usual Euler-Lagrange equation for minimization of integrals of this form gives:

$$p'(l(t))l'(t)^3 - \frac{d}{dt}(3p(l(t))l'(t)^2) = 0, \tag{47}$$

or after simplification:

$$3p(l(t))l'(t)l''(t) + p'(l(t))l'(t)^3 = 0. \tag{48}$$

Multiplying by $l'(t)$, the left-hand side is precisely the derivative of $p(l(t))l'(t)^3$, which must therefore be constant. That is, we have:

$$l'(t) = \left(\frac{C}{p(l(t))}\right)^{1/3}. \tag{49}$$

Finally for the correct choice of p, we have:

$$\frac{1}{l'(t)} = \frac{[F(l(t))^{n-2}(2F'(l(t))^2 + (1 - F(l(t)))F''(l(t)))]^{1/3}}{C}, \tag{50}$$

which describes the optimal discrete bid level density in the large m limit.

ACKNOWLEDGMENTS

The authors would like the thank the reviewers, whose suggestions lead to significant improvements in both the content and clarity of the final paper.

REFERENCES

ANTHONY, P. AND JENNINGS, N. R. 2003. Developing a bidding agent for multiple heterogeneous auctions. *ACM Trans. Intern. Tech. 3*, 3, 185–217.

BAJARI, P. AND HORTACSU, A. 2003. The winneŕs curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions. *RAND J. Econom. 34*, 2, 329–355.

BAPNA, R., GOES, P., AND GUPTA, A. 2003. Analysis and design of business-to-consumer online auctions. *Manage. Sci. 49*, 1, 85–101.

BAPNA, R., GOES, P., GUPTA, A., AND KARUGA, G. 2002. Optimal design of the online auction channel: Analytical, empirical and computational insights. *Decision Sci. 33*, 4, 557–577.

BLUMROSEN, L., NISAN, N., AND SEGAL, I. 2007. Auctions with severely bounded communication. *J. AI Resear. 28*, 233–266.

CASSIDY, R. 1967. *Auctions and Auctioneering*. University of California Press.

CHWE, M. S. 1989. The discrete bid first auction. *Econom. Lett. 31*, 303–306.

DAVID, E., ROGERS, A., SCHIFF, J., KRAUS, S., AND JENNINGS, N. R. 2005. Optimal design of English auctions with discrete bid levels. In *Proceedings of ACM Conference on Electronic Commerce (EC'05)*. 98–107.

HAGEMAN, L. A. AND YOUNG, D. M. 1981. *Applied Iterative Methods*. Academic Press.

JIANG, A. X. AND LEYTON-BROWN, K. 2005. Estimating bidders' valuation distributions in online auctions. In *Proceedings of IJCAI-05 Workshop on Game Theoretic and Decision Theoretic Agents*. 98–107.

KRESS, A. AND BOUTILIER, C. 2004. A study of limited-precision, incremental elicitation in auctions. In *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS' 04)*. 1344–1345.

LEVIN, D. AND SMITH, J. L. 1994. Equilibrium in auctions with entry. *Amer. Econom. Rev. 84*, 3, 585–99.

LUCKING-REILEY, D. H. 2000. Auctions on the Internet: What's being auctioned, and how? *J. Indus. Econom. 48*, 3, 227–252.

MYERSON, R. 1981. Optimal auction design. *Mathem. Operat. Resea. 6*, 1, 58–73.

PRESS, W. H., FLANNERY, B. P., TEUKOLSKY, S. A., AND VETTERLING, W. T. 1992. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press.

RILEY, J. G. AND SAMUELSON, W. F. 1981. Optimal auctions. *Amer. Econom. Rev. 71*, 381–392.

ROGERS, A., DAVID, E., SCHIFF, J., KRAUS, S., AND JENNINGS, N. R. 2005. Learning environmental parameters for the design of optimal English auctions with discrete bid levels. In *Proceedings of 7th International Workshop on Agent-Mediated E-Commerce (AMEC-VII)*. 81–94.

ROGERS, A., DAVID, E., SCHIFF, J., KRAUS, S., AND JENNINGS, N. R. 2007. The effects of proxy bidding and minimum bid increments within eBay auctions. *ACM Trans. Web*. http://eprints.ecs.soton.ac.uk/12716/

ROTHKOPF, M. H. AND HARSTAD, R. M. 1994. On the role of discrete bid levels in oral auctions. *Europ. J. Operat. Resea. 74*, 572–581.

YAMEY, B. S. 1972. Why 2,31,000 [pounds] for a Velazquez?: An auction bidding rule. *J. Politic. Economy 80*, 1323–1327.

YU, J. 1999. Discrete approximation of continous allocation machanisms. Ph.D. thesis, Division of Humanities and Social Science, California Institute of Technology.

Received January 2006; revised June 2006; accepted November 2006