

Sequences of Take-It-or-Leave-It Offers: Near-Optimal Auctions Without Full Valuation Revelation*

Tuomas Sandholm
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
sandholm@cs.cmu.edu

Andrew Gilpin
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
gilpin@cs.cmu.edu

ABSTRACT

We introduce *take-it-or-leave-it auctions (TLAs)* as an allocation mechanism that allows buyers to retain much of their private valuation information, yet generates close-to-optimal expected utility for the seller. We show that if each buyer receives at most one offer, each buyer's dominant strategy is to act truthfully. In more general TLAs, the buyers' optimal strategies are more intricate, and we derive the perfect Bayesian equilibrium for the game. We develop algorithms for finding the equilibrium and also for optimizing the offers so as to maximize the seller's expected utility. In several example settings we show that the seller's expected utility already is close to optimal for a small number of offers. As the number of buyers increases, the seller's expected utility increases, and becomes increasingly (but not monotonically) more competitive with Myerson's expected utility maximizing auction.

Categories and Subject Descriptors

I.2 [Artificial Intelligence]: General

General Terms

Economics, Algorithms

Keywords

auctions, optimal auction design, equilibrium computation, perfect Bayesian equilibrium, multiagent systems

1. INTRODUCTION

Auctions have emerged as a key mechanism for task and resource allocation in multiagent systems. We focus on the problem of allocating an indivisible good to one of a group of interested buyers, while maximizing the seller's expected utility.¹ In this setting, the *Myerson auction* yields the high-

*This material is based upon work supported by the National Science Foundation under CAREER Award IRI-9703122, Grant IIS-9800994, ITR IIS-0081246, ITR IIS-0121678, and ITR IIS-0427858, and a Sloan Fellowship.

¹All of our results apply to reverse auctions as well, in which a single buyer, facing many sellers, wants to minimize her procurement cost.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.
Copyright 2006 ACM 1-59593-303-4/06/0005 ...\$5.00.

est expected utility to the seller, compared to *any* other allocation mechanism [13]. Despite being the optimal allocation mechanism, the Myerson auction has several undesirable features. First, buyers must fully reveal their true valuations for the good.² Buyers may be unwilling to reveal private information because such information is proprietary or its revelation would have a negative strategic long-term impact. For example, after such an auction the buyers will be at a disadvantage in future negotiations with the seller (who may be able to extract more surplus in the future using information about the buyers' valuations). Similar long-term adverse effects have been observed in other truthful full-revelation auction mechanisms [17]. Second, buyers may not understand the complex, unintuitive rules of the Myerson auction, and therefore choose to not participate. Third, for a layman, it is unintuitive to submit true valuations because bid shading is rational in many commonly used auction mechanisms. Buyers are thus likely to shade bids even in Myerson auctions, with unpredictable and sub-optimal results. For these reasons and potentially others, to our knowledge, the theoretically elegant Myerson auction is not being used anywhere.

In the most commonly used auction mechanisms—the *English auction* (a.k.a. ascending open-cry auction) and the *Dutch auction* (a.k.a. descending auction)—bidders do not reveal their valuations completely [9]. Unfortunately, in these auctions, the seller's expected utility is not optimal and in fact can be arbitrarily far from optimal. The *Vickrey auction* [20] (a.k.a. second-price sealed-bid auction) suffers from both full valuation revelation and seller's expected utility being arbitrarily far from optimal.

In this paper we present a mechanism that achieves close-to-optimal expected utility for the seller while allowing each buyer to retain much of her private valuation information. Our *take-it-or-leave-it auction (TLA)* mechanism consists of a sequence of take-it-or-leave-it offers made by the seller to the buyers, each buyer in turn. Before the auction is held, the seller announces the order in which the buyers receive offers and the amount of each offer. During a fixed time after receiving an offer, the buyer has to accept (“take”) the offer, at which point the auction ends, or reject (“leave”) the offer, in which case the seller may proceed to make other offers. After rejecting an offer, the buyer never has the option of

²Recently, cryptographic techniques have also been proposed to address this drawback [14].

accepting that offer.³ We show how rational buyers should act in such an auction and how the seller can optimize the sequence of offers given her knowledge about the buyers. In this sense, our approach is *automated mechanism design* [7].

Sequences of take-it-or-leave-it offers have been studied in the context of optimal auctions previously [3]. However, in that paper there is an assumption that the seller can decide the accuracy to which the buyers learn their own valuations. One way they can achieve this is by withholding information. In our model, the seller has no control of the buyers' information.

In many applications the seller may be aware of some statistical information about the buyers. By taking this information into account, the seller can design mechanisms so as to maximize her expected utility. In our setting, this is done by making the assumption that every buyer's valuation is drawn from some known probability distribution. This assumption violates the *Wilson doctrine* which states that mechanisms should be designed without the use of any priors. Some examples of auctions that follow the Wilson doctrine include Vickrey, English, Dutch, and first-price sealed bid auctions. These types of auctions are not optimal, and in fact can be arbitrarily far from optimal. Further, it is generally not possible to design utility maximizing auctions (or even approximately optimal auctions) without using priors. The Myerson auction, on the other hand, does not adhere to the Wilson doctrine, and as a result is able to extract more surplus from the buyers. Just as is done in the Myerson auction, we use priors in our mechanism.

2. THE SETTING AND BACKGROUND

We study the usual auction setting in which a seller has an indivisible good that she can allocate to any one (or none) of a set of buyers $B = \{1, \dots, n\}$. Each buyer i has a valuation v_i for the good that is drawn from a distribution f_i with cumulative density F_i . The buyers' valuations are drawn independently.⁴ The distributions f_i are common knowledge, but the valuation v_i is only known by the buyer i herself. The setting is *symmetric* if each buyer draws her valuation from the same distribution. Otherwise the setting is *asymmetric*. Finally, we make the standard auction theory assumption that there is no aftermarket.

By designing the rules of the auction, the seller tries to maximize her expected utility. The seller has a valuation v_0 (potentially zero) for the good. By the standard quasilinearity assumption about the seller, the seller's utility is v_0 if she keeps the item, and if she allocates the item to a buyer, her utility is the revenue that she receives from the buyers.

To motivate the buyers to participate in the auction, the rules of the auction have to guarantee that each buyer receives no less utility by participating than by not participating. Such *ex post* participation constraints imply that 1) a buyer who does not get the good cannot be charged, and 2) the buyer who receives the good cannot be charged more than her valuation.

³At first glance, it appears that the Dutch auction and TLAs are equivalent. This is not true. The Dutch auction is more powerful in the sense that the seller can make an infinite number of offers, and less powerful in the sense that the seller has to make the offers symmetrically to the buyers.

⁴If the valuations can have arbitrary correlations, then even designing an approximately optimal auction is hard computationally [16].

The optimal mechanism for our setting is the Myerson auction [13]. It is the benchmark against which we will compare our TLA mechanism.⁵

DEFINITION 1. (Myerson auction [13])

- Every buyer i reveals a valuation \hat{v}_i .
- For each buyer i , compute that buyer's virtual valuation: $\psi_i = \hat{v}_i - \frac{1-F_i(\hat{v}_i)}{f_i(\hat{v}_i)}$.
- Select the buyer i^* with the highest virtual valuation ψ_{i^*} .
- If $\psi_{i^*} > v_0$, buyer i^* receives the item and pays the smallest valuation, \hat{v}_{i^*} , that would still have made her win.
- If $\psi_{i^*} \leq v_0$, the seller keeps the good and no payments are made.

In the symmetric setting, the Myerson auction is a second-price auction with an appropriately tailored reserve price. However, this is not the case in the asymmetric setting.

3. TAKE-IT-OR-LEAVE-IT AUCTION

In this section we define our idea of using a sequence of take-it-or-leave-it offers as an auction mechanism.

DEFINITION 2. (Take-it-or-leave-it auction (TLA)) Let $\mathcal{A} = \langle b_1, a_1 \rangle, \langle b_2, a_2 \rangle, \dots, \langle b_k, a_k \rangle$ be a sequence of k offers. Let an offer \mathcal{A}_j be a tuple $\langle b_j, a_j \rangle$, where a_j is the amount of the offer and $b_j \in B$ is the buyer to whom the offer is made. The entire sequence of offers \mathcal{A} is revealed by the seller to the buyers before the auction. The auction proceeds as follows:

- For j from 1 to k
 - Let $\langle b, a \rangle = \mathcal{A}_j$. The seller offers the good to buyer b for amount a .
 - If buyer b accepts, she gets the good and pays a to the seller. The auction ends.
 - If buyer b rejects offer a , the auction continues.
- If the auction has reached this point, the seller keeps the good and collects no payments.

It is worth mentioning that every instance of a TLA is *ex ante* individually rational (the strongest such form) since a buyer can always reject any offer. The notion of incentive compatibility (*i.e.* truth-telling) does not apply here since it's not a direct revelation mechanism.

The following auction mechanism turns out to be an interesting special case:

DEFINITION 3. (Single-offer TLA) A TLA is a single-offer TLA if each buyer gets at most one offer, *i.e.*, $b_q \neq b_r$ for all $q, r \in \{1, 2, \dots, k\}$, $q \neq r$.

When we need to differentiate between a single-offer TLA and a TLA in which a buyer may receive multiple offers, we will refer to the latter as a *multiple-offer TLA*.

Note that in a TLA it is entirely possible for the seller to allocate the good to a buyer who does not have the highest valuation, thus making the mechanism *inefficient*. This is also the case in the Myerson auction. It is well-known that in the asymmetric setting there is no mechanism that maximizes both seller's expected utility and efficiency.

⁵The proof of optimality of the Myerson auction assumes that the setting is *regular*, *i.e.*, the hazard rate $f_i/(1-F_i)$ is increasing for each buyer i [13]. This assumption is satisfied in all cases where we benchmark against the Myerson auction (in Section 7), so we are benchmarking against the real optimum.

3.1 Strategies and utility

A formal notation for strategies in a TLA is useful. Strategies for TLAs are slightly different than for other auctions since a single buyer may need to make many decisions during the course of an auction.

DEFINITION 4. (Strategy for a TLA) *A buyer's strategy for a TLA is a function from the buyer's valuation and the number of the offer (made to that buyer), to one of two possible actions: accept or reject:*

$$s_i : \mathbb{R} \times \{j \mid b_j = i, j \in \{1, 2, \dots, k\}\} \rightarrow \{\text{accept}, \text{reject}\}.$$

A strategy profile is $s = (s_1, \dots, s_n)$. A strategy profile excluding s_i is $s_{-i} = \{s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n\}$, and a strategy profile s with s_i replaced with s'_i is (s_{-i}, s'_i) .

We make the standard assumption of quasilinearity, that is, a buyer's utility u_i is zero if she does not get the good (and does not pay anything), and $v_i - a$ if she gets the good and pays amount a . Given a strategy profile s and the valuation v_i , the expected utility $u_i(v_i, s)$ of buyer i in a TLA \mathcal{A} is

$$u_i(v_i, s) = \sum_{\{j \mid b_j = i\}} \left(\prod_{l=1}^{j-1} R_l \right) u_i^j(v_i, s_i)$$

$$\text{where } u_i^j(v_i, s_i) = \begin{cases} v_i - a_j & \text{if } s_i(v_i, j) = \text{accept} \\ 0 & \text{if } s_i(v_i, j) = \text{reject}. \end{cases}$$

Here, R_l is the probability that offer l is rejected by buyer b_l if that offer is made. These probabilities depend on the buyers' strategy profile s , as discussed next.⁶

4. EQUILIBRIUM ANALYSIS

We use definitions from game theory to help us understand how rational buyers will behave in a TLA. We use two of the most prevalent solution concepts for sequential games: dominant strategies and perfect Bayesian equilibrium.

4.1 Single-offer TLAs

A strategy s_i^* for a buyer with valuation v_i is called a *dominant strategy* if for all other strategies s'_i and all strategies s_{-i} , we have $u_i(v_i, (s_i^*, s_{-i})) \geq u_i(v_i, (s'_i, s_{-i}))$. As the following proposition illustrates, the strategies employed by rational buyers in a single-offer TLA are straightforward and truthful. They are dominant strategies, so a buyer is best off playing that way regardless of what the other buyers' preferences are and even if others play irrationally. Also, a buyer does not have to know much about her own valuation (the determination of which can be a complex optimization problem itself, see e.g. [18, 19, 15, 10]), merely whether it is greater or less than the offer, making a single-offer TLA easy to play.

PROPOSITION 1. *In a single-offer TLA, a buyer's dominant strategy is*

$$s_i^*(v_i, j) = \begin{cases} \text{accept} & \text{if } v_i > a_j \\ \text{reject} & \text{if } v_i \leq a_j. \end{cases}$$

The strategy profile $s^ = (s_1^*, s_2^*, \dots, s_n^*)$ is thus a dominant strategy equilibrium.*

⁶When buyer i computes the utility, $R_l \in \{0, 1\}$ when $b_l = i$, because the buyer knows her own strategy.

4.2 Multiple-offer TLAs

If the sequence of take-it-or-leave-it offers includes potentially multiple offers per buyer, the buyers' equilibrium strategies become significantly more intricate. The reason is that *a rational buyer will be reluctant to accept an offer at—or somewhat below—her valuation, because there is a chance that the auction will make her a significantly lower offer later on.* The analysis is further complicated by the fact that a buyer's passing on an offer gives a signal to the other buyers and the seller about the buyer's valuation.

In the equilibrium analysis that follows, we study *threshold strategies*. In any finite-action game where types are drawn from independent atomless distributions and each player's utility function does not directly depend on other's types, a pure strategy equilibrium always exists [12]. (This existence result can be extended to the setting with interdependent valuations in several auction-like settings [2].) Clearly, a TLA is such a game. Furthermore, since each buyer receives strictly decreasing offers in an optimal TLA (see Proposition 4), we can restrict our attention to a strategy in which a player can simply decide at the beginning of the game the first offer she will accept. This reasoning partitions the type space into regions, one for each offer. We say these regions are separated by *thresholds*. If the player's type falls within a region, then the first offer the player will accept is the offer corresponding to the region. Thus, without loss of generality we can focus on threshold-based pure strategies in our analysis of an equilibrium of a multiple-offer TLA.

Perfect Bayesian equilibrium (PBE) [8] is the most common solution concept for sequential games with observed actions and private valuations (types). In such a game, with independently drawn valuations, each buyer i has a strategy s_i and beliefs (for each step j of the game, the buyer has, for each other buyer k , a cumulative probability density function F_k^j about the type of k). A strategy profile s^* is a PBE if the following conditions hold: 1) Each buyer updates her beliefs using Bayes rule whenever possible; 2) Whenever it is buyer i 's turn to move, s_i prescribes an action that maximizes i 's expected utility from then on, given i 's beliefs; and the three technical conditions: 3a) At every step of the game, buyer i 's beliefs about other buyers are independent of buyer i 's valuation, 3b) the beliefs about buyer i can only be changed when i acts, and 3c) at every step of the game, for every triple of buyers i, k , and l ($i \neq l, k \neq l$), i and k have the same beliefs about l .

We illustrate the PBE analysis through a 5-offer 2-buyer auction and present the general method after that. For readability, we refer to buyer 1 in the feminine and buyer 2 in the masculine. Suppose that the seller's chosen TLA prescribes that the first offer goes to buyer 1 and the remaining four offers alternate between the buyers.

Buyer 1, when facing the first offer, must decide whether or not she should accept it. If her valuation, v_1 , is less than the first offer, a_1 , then clearly she must reject. However, the converse is not as simple. *Even though she might stand to gain positive utility by accepting the first offer, she might expect to gain even more utility by rejecting the first offer and accepting a later offer.* For each offer made to a buyer, we would like to compute a *threshold* value above which the buyer will accept and below which she will reject. (If the buyer's valuation is equal to the threshold, she is indifferent between rejecting and accepting.) We will refer to buyer i 's

threshold at offer j as t_i^j .

At the last offer made to each buyer, both buyers have dominant strategies. They simply accept if the offer is below their valuation and reject otherwise. This is analogous to the strategies employed by buyers in a single-offer TLA.

Now, consider the third offer (the second offer made to buyer 1). Buyer 1 knows what her utility is if she accepts this offer. She wishes to determine what her *expected* utility is if she rejects the third offer (in a gamble to face the fifth offer). To compute this she needs to compute the probability that buyer 2 rejects the fourth offer. At this point in the game, buyer 1 knows that buyer 2 has rejected the second offer. Assuming buyer 2 is playing in equilibrium, his valuation must be less than t_2^2 , his threshold when facing the second offer. Buyer 1 must take this into account when computing the probability that buyer 2 will reject the fourth offer. *That is, buyer 1 updates her belief about the distribution from which buyer 2's valuation is drawn.* Buyer 1 is able to perform these computations using her knowledge about buyer 2's prior distribution, which is common knowledge. After buyer 2 has rejected the second offer, by Bayes rule, the right tail of his valuation distribution (probability density function) above t_2^2 can be cut, and the remaining distribution needs to be rescaled to make its integral one. Thus the new cumulative distribution is $F_2^4(x) = F_2(x)/F_2(t_2^2)$. (In general, we will write F_i^j to denote the world's belief about buyer i 's distribution when facing offer j . If a buyer i is facing her first offer j , then $F_i^j = F_i$.) With this new notation, we can write equations describing the thresholds for buyers 2 and 1 when facing offers 2 and 3, respectively:

$$\begin{aligned} t_2^2 - a_2 &= F_1^3(t_1^3)(t_2^2 - a_4) \\ t_1^3 - a_3 &= F_2^4(a_4)(t_1^3 - a_5). \end{aligned}$$

We evaluate F_1^3 at t_1^3 since we know that buyer 1 will only reject if her valuation is below t_1^3 .

Determining the strategy that buyer 1 should use at the first offer is slightly more complicated since buyer 1 could reject the first offer in a hope to accept the third offer or the fifth offer. The following equation describes buyer 1's threshold at the first offer:

$$t_1^1 - a_1 = F_2^2(t_2^2) \cdot \max \{t_1^1 - a_3, F_2^4(a_4)(t_1^1 - a_5)\}.$$

4.2.1 A simplifying observation

At this point we have completely *characterized* the strategies that the buyers will use, but we have said nothing about how to compute these strategies. Computing the strategies is complicated by the max operation in the above equation. The determination of the strategies could be split into cases, with each case corresponding to a particular term within the max operation being greatest. However, as we move beyond 5 offers, this becomes complex. As the number of offers increases by one, the number of max operations in the system increases by one, thus doubling the number of cases to analyze. Therefore, the number of cases is exponential in the number of offers.⁷

This computation would be easier if we could assume that the thresholds only depend on the next offer. In the above

⁷Actually, the number of cases is further increased by the fact that as an offer is added, the number of terms in the existing max operations of that buyer's threshold equations increases by one (a max operation with 2 terms is added in the equation where there was no max operation before).

equation, this would mean that $t_1^1 - a_3 \geq F_2^4(a_4)(t_1^1 - a_5)$. We observe that if this inequality does not hold, the third offer would in effect be wasted because it certainly would be rejected. Since this paper is concerned with optimal TLAs, we can assume that such offers will never be made.⁸

4.2.2 Equilibrium for multiple buyers and offers

Algorithm 1 below is a general procedure for generating the system of equations that determines buyers' thresholds. Before describing the algorithm we need to introduce some additional notations. First, we need an expression for evaluating updated distributions:

$$F_i^j(x) = \begin{cases} F_i(x) & \text{if } i \notin \{1, \dots, b_{j-1}\} \\ \frac{F_i(x)}{F_i(t_i^{j'})} & \text{if } i \in \{1, \dots, b_{j-1}\} \end{cases}$$

where j' is the most recent offer made to buyer i . Second, we need an expression for the probability that a buyer will reject a given offer j :

$$R_j = \begin{cases} F_{b_j}^j(a_j) & \text{if } b_j \notin \{b_{j+1}, \dots, b_k\} \\ F_{b_j}^j(t_{b_j}^j) & \text{if } b_j \in \{b_{j+1}, \dots, b_k\}. \end{cases}$$

The above value R_j depends on whether or not the buyer faces another offer or not. If the buyer does face another offer, then the expression involves a threshold value. Otherwise, the expression only involves the offer. The algorithm below makes a pass through the offers. If an offer is not the last one offered to a buyer, then an equation is output containing an expression for the threshold.

ALGORITHM 1.

// Output system of threshold equations

- For j from 1 to k , If buyer b_j faces another offer j' after offer j
 - Then output the equation

$$t_{b_j}^j - a_j = R_{j+1}R_{j+2} \cdots R_{j'-1}(t_{b_j}^{j'} - a_{j'})$$

The product of R_j 's that is output for each equation is the expression describing the probability that the auction will go from offer $j+1$ to offer j' without a buyer accepting an offer.

With the above discussion in mind, we can now present our main equilibrium result.

THEOREM 2. *Suppose we have a multiple-offer TLA $\mathcal{A} = \langle b_1, a_1 \rangle, \langle b_2, a_2 \rangle, \dots, \langle b_k, a_k \rangle$ that is optimal in the sense of maximizing the seller's utility, subject to the constraint that only k offers are made. Consider the following strategy for buyer i with valuation v_i facing offer a_j :*

$$s_i^*(v_i, a_j) = \begin{cases} \text{accept} & \text{if } v_i > a_j, i \notin \{b_{j+1}, \dots, b_k\} \\ \text{reject} & \text{if } v_i \leq a_j, i \notin \{b_{j+1}, \dots, b_k\} \\ \text{accept} & \text{if } v_i > t_i^j, i \in \{b_{j+1}, \dots, b_k\} \\ \text{reject} & \text{if } v_i \leq t_i^j, i \in \{b_{j+1}, \dots, b_k\} \end{cases}$$

where the t_i^j values solve the system of equations generated by Algorithm 1. The strategy profile $s^* = (s_1^*, s_2^*, \dots, s_n^*)$ is a perfect Bayesian equilibrium.

⁸This assumes that the order in which buyers receive offers is not fixed. If this order is fixed (for example, due to external constraints), then making a surely rejected offer can be beneficial in order to *in effect* change the order without actually changing it.

PROOF. It is easy to see that the technical Conditions 3a-3c of PBE are satisfied. As described above, the buyers update their beliefs using Bayes rule whenever possible, so Condition 1 is satisfied. What remains to be shown is that Condition 2 is satisfied, i.e., that it is never profitable to deviate from the above strategy (in any of the four scenarios). For the first two scenarios, the proof is analogous to the proof of Proposition 1. For the other two scenarios, in which the buyers face another offer later in the auction, we need to make use of the thresholds. Consider buyer i with valuation v_i , facing offer a_j . Let $a_{j'}$ be the next offer that buyer i faces. Suppose $v_i > t_i^j$, but buyer i rejects. The highest expected utility she can expect to get later in the auction is $R_{j+1} \cdots R_{j'-1} (v_i - a_{j'})$. But, by the construction of the thresholds (and the fact that the thresholds only depend on the next offer), this is less than $v_i - a_j$. So buyer i is better off accepting the offer. The scenario where $v_i \leq t_i^j$ is proven similarly. \square

5. OPTIMIZING THE OFFERS

So far we have described how we would expect buyers to behave in a TLA. Now we turn to the problem faced by the seller: how to maximize expected utility? The seller can choose the order in which the offers are made as well as the offer values. Before describing the optimization problem, we formalize the seller's objective.

DEFINITION 5. (Seller's expected utility) *Given an instance \mathcal{A} of a TLA, the seller's expected utility from step j onward is defined recursively as*

$$\pi_j = \begin{cases} (1 - R_j) a_j + R_j v_0 & \text{if } j = k \\ (1 - R_j) a_j + R_j \pi_{j+1} & \text{if } j < k. \end{cases}$$

The seller's expected utility is given by π_1 , and we denote it by $\pi(\mathcal{A})$.

We can now formally define the design problem faced by the seller.

DEFINITION 6. (TLA design problem) *Given a limit k_{max} on the number of offers made, the seller's TLA design problem is to compute $\mathcal{A}^* = \operatorname{argmax}_{\{\mathcal{A} \mid |\mathcal{A}| \leq k_{max}\}} \pi(\mathcal{A})$.*

Note that in Definition 6 the k_{max} parameter is exogenous. It is a constraint imposed on the optimization problem that comes from some external source, for example a limit on the amount of time or resources the seller is able to commit to the auction.

We now discuss some characteristic properties of optimal TLAs.

PROPOSITION 3. *For any given TLA design problem, an optimal TLA exists in which no buyer receives consecutive offers.*

PROOF. If a buyer receives consecutive offers, she will never accept any of them except possibly the lowest. Thus, the other ones of those offers can be removed from \mathcal{A} . \square

Offers might not decrease over time in optimal TLAs:

EXAMPLE 1. *Consider the 2-buyer, 3-offer setting where buyer 1's valuation is uniformly distributed on the interval $[0, 1]$ and buyer 2's valuation is uniformly distributed on the interval $[1, 4]$. The only optimal sequence of offers is $\mathcal{A} = \langle 2, 2.125 \rangle, \langle 1, 0.5275 \rangle, \langle 2, 2.0 \rangle$.*

However, decreasing offers can be made to each buyer:

PROPOSITION 4. *For any given TLA design problem, an optimal TLA exists in which each buyer individually receives strictly decreasing offers.*

PROOF. If, in equilibrium, all the rejection probabilities R_j of the other buyers that get offers in between two offers of buyer i are zero, that is equivalent to buyer i receiving consecutive offers, which is unnecessary by Proposition 3. We can therefore restrict attention, without loss in the seller's expected utility, to TLAs where this never occurs.

We now show that a buyer will never accept an offer that is equal to or higher than her previous offer. The main idea is that since the buyer knows the sequence of offers at the start of the auction, she knows she can get more utility by accepting an earlier, lower offer, rather than accepting the later, higher offer. Suppose that in an optimal TLA, buyer i is offered a_j at some point in the auction, and the next offer to buyer i is $a_{j'}$, where $a_j \leq a_{j'}$. Without loss of generality, say buyer i 's valuation v_i is at least $a_{j'}$. (If it is less than, then buyer i will clearly never accept the offer.) Then the expected utility from accepting offer a_j is $v_i - a_j$, while the expected utility from rejecting offer a_j and later accepting offer $a_{j'}$ is strictly less than $v_i - a_j$ since the utility must be multiplied by the probability that the auction will continue to offer $a_{j'}$ without another buyer accepting an offer and thus ending the auction. Let r denote the probability that no other buyer accepts in between. Since $a_j \leq a_{j'}$ and $r < 1$ (by the argument in the paragraph above), we have $v_i - a_j > r(v_i - a_{j'})$, so buyer i will never accept offer $a_{j'}$. Thus, the wasted offer $a_{j'}$ can be removed from the auction and the auction remains optimal. \square

6. COMPUTATIONAL RESULTS

In this section we develop the computational methodology for solving for an optimal TLA (which involves as a sub-problem solving for how the buyers will behave). Designing a TLA involves finding an order (the b_j values) in which the buyers receive the offers, and determining the value of each offer (the a_j values).

6.1 Multiple-offer TLAs

In computing the optimal multiple-offer TLA, we assume that the order of the offers is fixed (e.g., one can include an outer loop to try all orders) and then optimize the TLA for that order as follows.

We run Algorithm 1 to generate the equations for the thresholds t_i^j . Because the shorthand values R_j are expanded in each of those equations, each equation has on the right hand side the offer levels a_1, a_2, \dots, a_k , the specific functional forms of the prior distributions on the types F_1, F_2, \dots, F_n , and the thresholds t_i^j . To compute the optimal offers, we run a non-linear optimizer with these equations as constraints and the expanded recursion for $\pi(\mathcal{A})$ from Definition 6 as the objective, which is now just a function of the offer levels a_1, a_2, \dots, a_k . The output of the optimizer is the values of the optimal offer levels $a_1^*, a_2^*, \dots, a_k^*$.

The computational complexity of the above procedure is difficult to analyze. Indeed, non-linear optimization problems as a class have resisted rigorous analysis. However, algorithms for finding local optima behave well in practice.

6.1.1 A tractable case

Although the question of complexity remains open for the general case, we have developed a linear-time algorithm for the special case of two buyers with uniform distributions on the interval $[0, 1]$. This algorithm relies heavily on results obtained recently [5, 6]. The authors examine auctions in which the amount of information a buyer can transmit to the seller is severely limited. One of the main results is that when a buyer is limited to choosing from among k possible messages, the buyer's valuation range is partitioned into k continuous intervals (the intervals are, in general, asymmetric among the bidders). The buyer signals which interval her valuation happens to fall in. It turns out that the thresholds that separate the intervals are the same as the equilibrium thresholds in an optimal TLA. (Note that these are the equilibrium thresholds, not the offer levels.) This is easily seen by considering a TLA where each buyer receives k offers, and thus has k threshold values. The TLA could be converted to a direct mechanism where the buyer simply states the first offer that she is willing to accept. Then the problem of computing the thresholds in a TLA becomes the same as computing the thresholds in a communication-bounded auction. The following theorem from [5] (stated using our threshold notation) motivates Algorithm 2.

THEOREM 5. ([5]) *When there are two buyers with valuations drawn uniformly from the interval $[0, 1]$ the thresholds in an optimal $2k$ -offer TLA, $k \geq 2$, are:*

$$t_i^j = \begin{cases} t + \frac{(2k-j-1) \cdot (1-t)}{2k-3} & \text{if } j \in \{1, 2, \dots, 2k-1\} \\ \frac{1}{2} & \text{if } j = 2k \end{cases}$$

where $t = \frac{-2\alpha + \sqrt{1+3\alpha}}{2(1-\alpha)}$ and $\alpha = \frac{1}{(2k-3)^2}$. The first offer in the TLA goes to buyer 1, and the offers alternate after that.

Algorithm 2 computes the offer levels for a multiple-offer TLA with two buyers having uniform distributions on the interval $[0, 1]$.

ALGORITHM 2.

// Let $k_{max} \geq 4$ be an even number denoting the total number of offers made.

1. (Initialize.) $k \leftarrow \frac{k_{max}}{2} + 1$, $\alpha \leftarrow \frac{1}{(2k-3)^2}$, $t \leftarrow \frac{-2\alpha + \sqrt{1+3\alpha}}{2(1-\alpha)}$, $t_2^{k_{max}} \leftarrow \frac{1}{2}$
2. (Compute thresholds.) For j from 1 to $k_{max} - 1$
 - (a) If j is odd, then $i \leftarrow 1$, else $i \leftarrow 2$
 - (b) $t_i^j \leftarrow t + \frac{(k_{max}-j-1)(1-t)}{2k-3}$
3. (Compute last offers.) $a_{k_{max}} \leftarrow t_2^{k_{max}}$, $R_{k_{max}} \leftarrow \frac{a_{k_{max}}}{t_2^{k_{max}-2}}$, $a_{k_{max}-1} \leftarrow t_1^{k_{max}-1}$, $R_{k_{max}-1} \leftarrow \frac{a_{k_{max}-1}}{t_1^{k_{max}-3}}$
4. (Compute offers.) For j from $k_{max} - 2$ down to 1
 - (a) If j is odd, then $i \leftarrow 1$, else $i \leftarrow 2$
 - (b) $a_j \leftarrow t_i^j - R_{j+1}(t_i^j - a_{j+2})$
 - (c) If $j > 2$, then $R_j \leftarrow \frac{t_i^j}{t_i^{j-2}}$, else $R_j \leftarrow t_i^j$

THEOREM 6. *Algorithm 2 computes an optimal TLA for the setting with two buyers having uniform distributions on the interval $[0, 1]$.*

PROOF. Immediate from Theorems 2 and 5. \square

6.2 Single-offer TLAs

Now we are considering single-offer TLAs, so we have $k_{max} \leq n$ by definition. Since making at least some offer to every buyer cannot hurt the seller's expected utility, we always have an optimal single-offer TLA with $k = k_{max}$. So, we set $k = k_{max}$.

For convenience, we will think of the optimization problem as two separate problems. The first problem is to determine the order in which the buyers receive the offers. The second problem is to compute the actual offer values.

First, consider symmetric settings, i.e., $F_i = F_j$ for all buyers i, j . In this case the order does not matter, so we set them arbitrarily. The offer values can be determined by first making an offer to the last bidder equal to the inverse virtual valuation of the reserve price and computing the expected profit from this offer as the new reserve price. Formally:

ALGORITHM 3.

// Compute offers for single-offer TLA

1. (Initialize.) $\pi \leftarrow v_0$
2. (Compute offers.) For i from k_{max} down to 1
 - (a) $a_i \leftarrow \operatorname{argmax}_a (1 - F_i(a))a + F_i(a)\pi$
 - (b) $\pi \leftarrow (1 - F_i(a_i))a_i + F_i(a_i)\pi$

PROPOSITION 7. *Algorithm 3 computes an optimal single-offer TLA in the symmetric settings.*

PROOF. Consider the last offer made in a TLA. The result of the last offer does not affect the result of any of the other offers. Thus, the seller should make an optimal offer, which is computed in step 2(a) during the first iteration. Now consider the next to last offer. Again, the result of this offer does not affect any of the previous offers. Also, the expected utility the seller earns if this offer is rejected has already been computed and is stored in π . Thus, the offer computed in step 2(a) is the optimal offer. Continuing with this backward induction, it is easy to see that Algorithm 3 correctly computes the optimal offer levels. \square

If we assume that F_i can be evaluated in constant time and that the "argmax" computation can also be done in constant time, then Algorithm 3 runs in time linear in the number of offers k_{max} . For a wide range of common distributions, including uniform and exponential distributions, these assumptions hold.

Now, consider asymmetric settings. Clearly the order of offers now can make a difference to the seller's expected utility. For general asymmetric preferences, we do not know of an efficient algorithm for ordering the offers. (For example, always greedily making an offer to the remaining buyer from whom the seller can myopically expect the greatest utility is not generally optimal.) One could try all possible orderings, running Algorithm 3 to evaluate each ordering. However, for a large number of buyers this would be infeasible since the number of orderings is $k_{max}!$. Fortunately, we have been able to show that in some commonly used special cases one can order the buyers in polynomial time. Theorem 8 illustrates just one of many results. We denote a uniform distribution over the interval $[a, b]$ as $U[a, b]$.

THEOREM 8. *With $v_i \sim U[0, w_i]$, an optimal single-offer TLA can be determined in $O(n \log n)$ time by sorting the buyers in decreasing order of w_i and passing the sequence to Algorithm 3.*

PROOF. Omitted due to space. \square

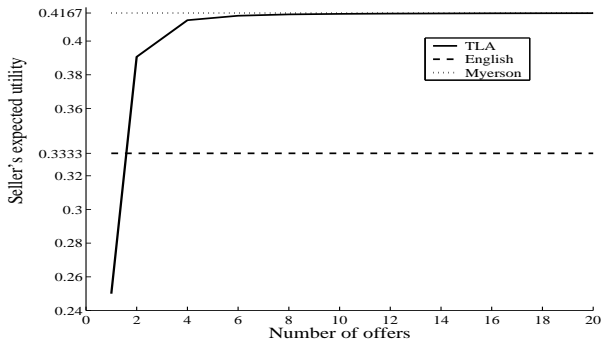


Figure 1: Seller’s expected utility as the number of offers increases. The two buyers’ valuations have uniform distributions on the interval $[0, 1]$.

7. ECONOMIC PERFORMANCE OF TLAS

In this section we present experiments and theoretical results to give an indication of how one can expect our TLAs, optimized as explained above, to perform in practice. (Throughout these experiments, we set the seller’s valuation v_0 to zero.)

First we studied the performance of our optimal single-offer TLA as the number of buyers increases. Naturally, as the number of buyers increases, the expected utility of the single-offer TLA increases. Table 1 shows how well an optimal single-offer TLA fares relative to Myerson’s utility-maximizing auction as the number of buyers increases (and each buyer is made exactly one offer).⁹ The single-offer TLA is close to optimal. Interestingly, the relative performance does not improve monotonically, but does improve when the number of buyers is large.¹⁰ The advantage of our single-offer TLA over the English auction is greatest when the number of buyers is not very large. Also, it tends to be greatest in the asymmetric settings. (Naturally, our multiple-offer TLAs would yield even greater expected utility for the seller.) In the case of 1 buyer, a TLA performs exactly as the Myerson auction, and the English auction does not generate any expected utility for the seller.

Figure 1 shows how the seller’s expected utility in an optimal multiple-offer TLA increases as the number of offers increases. As can be seen, the TLA performs extremely well after a small number of offers. We also studied how well the seller’s expected utility in an optimal multiple-offer TLA increases as the number of offers increases. The following result, which follows immediately from a result in [5] and the correspondence between communication-bounded auctions and TLAs, shows that we can expect a TLA to perform well for a wide variety of distributions.

THEOREM 9. ([6]) *The revenue loss in an optimal k -offer TLA with symmetric buyers having regular distributions¹¹ is $O(\frac{1}{k^2})$.*

For some distributions, including the uniform distribution,

⁹We denote the exponential distribution having mean $1/\mu$ as $\text{Exp}[\mu]$ (its cumulative density is $F(x) = 1 - e^{-\mu x}$).

¹⁰In the third row, the optimal single-offer TLA deterministically allocates the good to the highest buyer, thus making the allocation *efficient* in this case, where the Myerson auction is inefficient.

¹¹Recall that a distribution F_i is said to be regular if $f_i/(1 - F_i)$ is increasing.

the above bound is tight. This result is asymptotic, but in practice TLAs yield close to optimal expected utility for the seller, as exemplified in Table 2 where we considered several 2-buyer settings, showing how close a TLA gets to the expected utility generated by Myerson’s utility-maximizing auction as the number of offers increases. Even with a small number of offers, the expected utility of a TLA is close to optimal, and increases monotonically with the number of offers.

Total number of offers:	1	2	3	English
$v_1 \sim U[0, 1], v_2 \sim U[0, 1]$.600	.938	.979	.800
$v_1 \sim U[0.5, 2], v_2 \sim U[0, 1]$.857	.977	.992	.625
$v_1 \sim U[1, 4], v_2 \sim U[0, 1]$.932	.994	.997	.350
$v_1 \sim \text{Exp}[2], v_2 \sim \text{Exp}[2]$.551	.932	.977	.749
$v_1 \sim \text{Exp}[2], v_2 \sim \text{Exp}[4]$.726	.968	.986	.658

Table 2: Seller’s expected utility of an optimal 2-buyer TLA and English auction divided by the expected utility of the Myerson auction. The number of offers for the TLA is varied from 1 to 3. In each case, it is optimal to give the first offer to buyer 1

In principle, it is possible to achieve optimal expected revenue using an ascending auction, albeit one with asymmetric prices. We call such an auction an *asymmetric ascending auction*. (An asymmetric descending auction can be constructed similarly.) In an ascending auction, there is a single price, visible by all bidders, that rises. Bidders drop out when they no longer would like to buy the item at the current price. As discussed previously, this auction does not generate optimal expected revenue for the seller and in fact can be arbitrarily far from optimal. Alternatively, in an asymmetric ascending auction, there is a single price (the *virtual price*) that rises. But instead of the auctioneer communicating this price to the bidders, she instead communicates asymmetric prices to each bidder, using the inverse of the virtual valuation function evaluated at the current virtual price to compute prices for each bidder. It is straightforward to demonstrate that this method generates optimal expected revenue for the seller.¹²

One advantage of TLAs over asymmetric ascending auctions is that much less communication is required. This is especially important in auctions in which there is some cost to communication, for example when the bidders and the auctioneer communicate via a network. (In general, there is no upper bound on the amount of communication required for an ascending auction.) In fact, a TLA is, in some sense, an approximation of an asymmetric descending auction when there is some limit on the communication resources available to the bidders.

Finally, we mention that a single-offer TLA implements the Myerson auction in the case of a single buyer.

8. MECHANISMS SIMILAR TO TLAS

We are aware of one fielded mechanism that is similar to a TLA. Amazon.com has recently introduced the “Gold Box” on their web site [1]. Each day, registered users are given up to ten take-it-or-leave-it offers for ten different goods. The offers expire after one hour, and after an offer is rejected, it is no longer available to the user. The offer levels appear to

¹²As in the proof of the optimality of the Myerson auction, the bidders’ distributions must be regular for the result to hold.

Num buyers:	1	2	4	6	8	10	20	100
$v_i \sim U[0, 1]$	1 / 0	.938 / .800	.981 / .978	.894 / .997	.900 / 1	.907 / 1	.937 / 1	.992 / 1
$v_i \sim U[0, i]$	1 / 0	.980 / .646	.959 / .844	.953 / .898	.952 / .922	.952 / .936	.958 / .963	.979 / .987
$v_i \sim U[i, i + 1]$	1 / 0	.979 / .735	.990 / .866	.993 / .910	.995 / .933	.996 / .946	.998 / .973	1 / .995
$v_i \sim \text{Exp}[2]$	1 / 0	.932 / .749	.994 / .961	.847 / .991	.837 / .997	.833 / .999	.836 / 1	.872 / 1
$v_i \sim \text{Exp}[i]$	1 / 0	.509 / .657	.917 / .813	.900 / .824	.918 / .822	.927 / .819	.942 / .813	.948 / .810

Table 1: The left (right) number in each cell is the ratio of the expected utility of an optimal single-offer TLA (English auction) to the expected utility of the Myerson auction.

be below the normal retail price, but we do not know what pricing strategy is in use. The techniques developed in this paper could be used to compute optimal values for these prices so as to maximize Amazon’s expected utility (under certain reasonable assumptions).

A piece of closely related research concerns online auctions [11] and online exchanges [4]. In those papers the authors examine how bids should be treated in an online setting where the auctioneer (or exchange administrator) has to make accept/reject decisions about expiring bids before all potential future bids are received. The authors have developed online algorithms and analyzed how well their algorithms do compared to an omniscient algorithm that knows all future bids in advance. Many of those algorithms end up choosing bid acceptance thresholds, so in that sense they use (implicit) take-it-or-leave-it offers.

This setting differs from ours in that there is not only uncertainty about future decisions that will be made, but there is uncertainty about the types of decisions that will be faced. In our setting, the sequence of offers is known in advance by all parties. In addition, in the online model there are multiple units of the item available for exchange, while we have focused only on the single unit case.

Finally, most retail is of the form of making take-it-or-leave-it offers. In the typical retail setting the seller posts price for each item she wishes to sell. The seller generally does not have specific information about a given buyer, and so all buyers face the same prices (i.e. the prices are *non-discriminatory*). This distinguishes the typical retail setting from single-offer TLAs where the prices are *discriminatory*.

9. CONCLUSIONS AND FUTURE WORK

We introduced TLAs as a selling (or buying) mechanism that has low valuation revelation as do the most commonly used auction mechanisms (English and Dutch auctions), yet, unlike them, generates close-to-optimal expected utility for the seller. We showed that in single-offer TLAs, each buyer’s dominant strategy is to act truthfully. In multiple-offer TLAs, the buyers’ optimal strategies are more intricate, and we derived the perfect Bayesian equilibrium for the game. We developed algorithms for finding the equilibrium and for optimizing the offers in single-offer and multiple-offer TLAs so as to maximize the seller’s expected utility (subject to an exogenous constraint on the number of offers), and proved that optimal TLAs have many desirable features. We applied a result due to Blumrosen and Nisan to TLAs to show that the seller’s expected utility is close to optimal already for a small number of offers. We also showed in several examples that as the number of buyers increases, the seller’s expected utility increases and becomes increasingly (but not monotonically) more competitive with Myerson’s expected utility maximizing auction. Myerson’s uses full valuation revelation and is arguably impractical because the rules are unintuitive, unlike ours in which the mechanism is simply

a sequence of take-it-or-leave-it offers. Even our single-offer TLA tends to yield significantly greater utility to the seller than the English auction when the setting is asymmetric or the number of buyers is not very large.

Future research includes developing fast offer-ordering algorithms for asymmetric single-offer TLAs for arbitrary valuation distributions, as well as fast special-purpose algorithms for PBE finding, offer ordering, and offer value optimization in multiple-offer TLAs. We also plan to extend this work to auctions of multiple identical units and to auctions of multiple distinguishable items.

10. REFERENCES

- [1] Amazon.com. Gold box offers. <http://www.amazon.com/gold-box>, 2003.
- [2] S. Athey. Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. *Econometrica*, 69(4):861–889, July 2001.
- [3] D. Bergemann and M. Pesendorfer. Information structures in optimal auctions. Discussion Paper 1323, Yale U., 2001.
- [4] A. Blum, T. Sandholm, and M. Zinkevich. Online algorithms for market clearing. In *SODA-02*, 2002.
- [5] L. Blumrosen and N. Nisan. Auctions with severely bounded communication. In *FOCS-02*, 2002.
- [6] L. Blumrosen, N. Nisan, and I. Segal. Multi-player and multi-round auctions with severely bounded communication. In *ESA*, pages 102–113, 2003.
- [7] V. Conitzer and T. Sandholm. Complexity of mechanism design. In *UAI-02*, pp. 103–110, 2002.
- [8] D. Fudenberg and J. Tirole. Perfect Bayesian equilibrium and sequential equilibrium. *J. of Economic Theory*, 53:236–260, 1991.
- [9] V. Krishna. *Auction Theory*. Academic Press, 2002.
- [10] K. Larson and T. Sandholm. Costly valuation computation in auctions. In *TARK VIII*, 2001.
- [11] R. Lavi and N. Nisan. Competitive analysis of incentive compatible on-line auctions. In *ACM-EC*, 2000.
- [12] P. Milgrom and R. Weber. Distributional strategies for games with incomplete information. *Mathematics of Operations Research*, 10:619–632, 1985.
- [13] R. Myerson. Optimal auction design. *Mathematics of Operation Research*, 6:58–73, 1981.
- [14] M. Naor, B. Pinkas, and R. Sumner. Privacy preserving auctions and mechanism design. In *ACM-EC*, 1999.
- [15] D. Parkes. Optimal auction design for agents with hard valuation problems. In *Agent-Mediated Electronic Commerce Workshop at IJCAI-99*, Stockholm, 1999.
- [16] A. Ronen and A. Saberi. Optimal auctions are hard. In *FOCS*, pages 396–405, 2002.
- [17] M. Rothkopf, T. Teisberg, and E. Kahn. Why are Vickrey auctions rare? *J. of Political Economy*, 98(1):94–109, 1990.
- [18] T. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. In *AAAI*, pages 256–262, Washington, D.C., July 1993.
- [19] T. Sandholm. Issues in computational Vickrey auctions. *International Journal of Electronic Commerce*, 4(3):107–129, 2000. Special Issue on Applying Intelligent Agents for Electronic Commerce.
- [20] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16, 1961.