

Auction Mechanism for Optimally Trading Off Revenue and Efficiency in Multi-unit Auctions

Anton Likhodedov and Tuomas Sandholm

Carnegie Mellon University
Computer Science Department
5000 Forbes Avenue
Pittsburgh, PA 15213
{likh,sandholm}@cs.cmu.edu

November 8, 2003

Abstract

We study auctioning multiple units of the same good to potential buyers with single unit demand (i.e. every buyer wants only one unit of the good). Depending on the objective of the seller, different selling mechanisms are desirable. The Vickrey auction with a truthful reserve price is optimal when the objective is efficiency - allocating the units to the parties who values them the most. The Myerson auction is optimal when the objective is the seller's expected utility. These two objectives are generally in conflict, and cannot be maximized with one mechanism. In many real-world settings—such as privatization and competing electronic marketplaces—it is not clear that the objective should be either efficiency or seller's expected utility. Typically, one of these objectives should weigh more than the other, but both are important. We account for both objectives by designing a new *deterministic* dominant strategy auction mechanism that maximizes expected social welfare subject to a minimum constraint on the seller's expected utility. This way the seller can maximize social welfare subject to doing well enough for himself.

The results in this paper are derived under the asymmetric independent private values model, which assumes that the distributions of buyers' valuations are common knowledge. We also describe a prior-free mechanism, which does not assume that the distributions are known. When the number of buyers tends to infinity and the number of units on sale is at least two, this auction approaches expected efficiency and expected seller's utility of the auction, designed with distributions known upfront.

1 Introduction

Electronic commerce has spawned the use of increasingly sophisticated auction mechanisms. There are several possible reasons for this. First, in many ecommerce settings the bidders are automated agents that are programmed to act rationally even in complex situations. Also human users of ecommerce systems are typically quite savvy and able to recognize attractive properties of sophisticated mechanisms. This motivates the analysis of unintuitive auction mechanisms, which are capable of meeting complicated design objectives.

One example is the Vickrey auction [Vi61]. In this auction with q_0 units of the same good on sale, q_0 highest bidders win, but only pay the price of the first *unsuccessful* bid. The Vickrey auction maximizes economic *efficiency*, aka. *social welfare* (assuming the reserve price is set to equal the seller's valuation for a unit of good being sold), that is, the units end in the hands of the party who values it the most.

Another example is the Myerson auction [My81], which maximizes the *seller's expected utility* (expected *revenue* in case the seller does not value the object).¹ The unintuitive aspect of the Myerson auction is that it sometimes allocates the goods to bidders other than the q_0 highest bidders.

¹The formal characterization of the seller's optimal multi-unit mechanism was given by E.Maskin and J.Riley in [MR89]. However in the single-unit demand case, which is considered in this paper, the seller's optimal multi-unit mechanism can be derived similarly to the seller's optimal single-unit mechanism from [My81]. Therefore we will refer to the seller's-optimal mechanism as Myerson auction.

Expected social welfare and the seller’s expected utility cannot be maximized with the same auction mechanism in general because these objectives conflict. Furthermore, in many real-world settings it is not clear that the objective should be either of the two.

For example, most privatization auctions are motivated by the belief that private companies can make more efficient use of an asset than the government can. It seems thus reasonable to allocate the asset to the party who can make the most effective use of it, that is, to use efficiency as the auction objective. At the same time, the government would like to raise as much money from the sale as possible (maximize the seller’s expected utility) because the asset is owned by the tax payers who prefer to pay for government expenditures out of the auction revenue rather than taxes.

As another example, consider electronic auction houses that compete with each other. To attract sellers, an auction house would use an auction mechanism that maximizes the seller’s expected utility. On the other hand, this would not be desirable from the perspective of attracting buyers. Clearly, an auction house needs both buyers and sellers to operate. Therefore, including some element of social welfare measurement in the objective may be desirable.

We account for both objectives by designing a new deterministic auction mechanism that maximizes expected social welfare subject to a minimum constraint on the seller’s expected utility. This way the seller can maximize social welfare subject to doing well enough for himself. We show that this auction mechanism belongs to a family of mechanisms, maximizing a linear combination of the seller’s expected utility and expected social welfare, controlled by one free parameter λ .² We then present an algorithm for determining the optimal value for this parameter and constructing an auction with desired characteristics. This approach is different from the traditional mechanism design - the mechanism is not completely specified upfront - rather we compute the particular mechanism for every instance of the problem (that is, given the constraint on the revenue and the distributions of buyers’ valuations). In a way this is similar in spirit to automated mechanism design [CS02].

We also present a family of much simpler randomized mechanisms that *under some assumptions*, which we state later in the paper achieve the same expected revenue and efficiency. However, in Section 4 we argue, that randomization is inappropriate for the settings, which motivate the present work. This discourages the use of the randomized mechanisms even in the cases, when they yield similar expected revenue and social welfare.

Although most of the results in this paper are derived under the independent private values model, which assumes that the distributions of valuations (types) of buyers are common knowledge, it is possible to relax this assumption. In Section 5 we describe a prior-free mechanism, which utilizes the idea of sampling from [BV03]. When the number of buyers tends to infinity this mechanism approaches expected efficiency and expected seller’s utility of the auction, designed with distributions known upfront. This way we meet the principle, known as Robert Wilson doctrine.³

2 Framework and notation

We focus on settings with one seller, multiple buyers, and multiple units of the same good on sale. For convenience buyers are indexed with numbers from 1 to n and the set of all buyers is denoted by $N = \{1 \dots n\}$. Index 0 always refers to the seller. We analyze the case of the single unit demand: each buyer wants to buy just one unit of the good.⁴

The seller’s valuation for each item is v_0 . The number of items on sale is denoted by q_0 . We also make the following usual assumptions about valuations, known as the *independent private values (IPV) model*:

1. The valuation of buyer i , v_i , is a realization of the random variable X_i with the cumulative distribution F_i and the density $f_i : [a_i, b_i] \rightarrow \mathbf{R}$. Each density function f_i is continuous and positive on $[a_i, b_i]$, and zero elsewhere.
2. All densities f_i are common knowledge.⁵
3. All random variables X_i are independent of each other.

²Note that we *do not* assume in advance that the tradeoff is attained with a linear combination of seller’s expected utility and expected social welfare - instead the linearity of the tradeoff is proved explicitly

³Robert Wilson doctrine of mechanism design states that the mechanism should be independent of the prior distribution of the bidder’s valuations

⁴The case when buyers want to buy more than one good can be handled by the same model as long as buyers’ demand functions are constant (i.e. a buyer’s valuation for each additional unit of the good is the same as her valuation for the first unit).

⁵We relax this assumption in Section 5.

We use the following notation. V denotes the set of all possible combinations of valuations of buyers:

$$V = \times_{i \in N} [a_i, b_i] \quad \text{and} \quad V_{-i} = \times_{j \in \{N \setminus i\}} [a_j, b_j]$$

We also need a special notation for vectors of valuations:

$$\begin{aligned} v &= (v_1, \dots, v_n) \\ v_{-i} &= (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n) \\ (v_{-i}, w_i) &= (v_1, \dots, v_{i-1}, w_i, v_{i+1}, \dots, v_n) \end{aligned}$$

The joint distribution of the valuations is denoted by $f(v)$. By the independence assumption we have

$$f(v) = \prod_{j=1}^n f_j(v_j) \quad \text{and} \quad f_{-i}(v) = \prod_{j \neq i}^n f_j(v_j)$$

Utilities of buyers are assumed to be quasi linear: The utility of buyer i is $u_i = p_i \cdot v_i - t_i$, where p_i is the probability that she gets an item, v_i is her valuation, and t_i is the amount that she has to pay. The utility of the seller is $u_0 = (q_0 - \sum_{i=1}^n p_i)v_0 + \sum_{i=1}^n t_i$ (the first term in the expected number of items kept by the seller).

We try to design a *mechanism* in order to meet some objective, described later, when each buyer plays the game so as to maximize his own expected utility. By the *revelation principle* we can now, without loss in the objective, restrict our attention to mechanisms where each buyer *truthfully* bids for a unit of good on sale in a sealed-bid format.

Definition 2.1 (sealed-bid mechanism) *Each buyer i submits a bid b_i for a unit of the good on sale. Upon obtaining the bids, the seller computes the allocation and the payment of each buyer. The allocation is the probability vector $p(b) = (p_1(b_1), \dots, p_n(b_n))$, where $p_i(b_i)$ is the probability that buyer i gets a unit of good, when bidding b_i . These probabilities do not have to sum to q_0 : the seller may keep some units. The payments are specified by the vector $t(b) = (t_1(b_1), \dots, t_n(b_n))$, where $t_i(b_i)$ is the payment of buyer i . The allocation rule $p(b)$ and the payment rule $t(b)$ are common knowledge.*

The bid of buyer i depends on his valuation v_i for a unit of good on sale. If each buyer is motivated to submit a bid that equals that buyers valuation, the mechanism is *incentive compatible*. As is standard in the literature on optimal auction design, throughout this paper we focus on *Bayes-Nash* incentive compatibility, that is, each buyer expects to get highest utility by bidding truthfully rather than insincerely—given that the other buyers bid truthfully. However, the mechanisms, which we derive in this paper turn to be also (*ex post*) *incentive compatible*.

When all participant are truthful, the expected utility of buyer i can be expressed as follows:

$$U_i(p, t, v_i) = E_{v_{-i}} [p_i(v)v_i - t_i(v)] = \int_{V_{-i}} (p_i(v)v_i - t_i(v)) f_{-i}(v_{-i}) dv_{-i} \quad (2.1)$$

When a buyer chooses to bid differently from his true type and given that other bidders bid vector is v_{-i} , her utility is

$$p_i(v_{-i}, w_i)v_i - t_i(v_{-i}, w_i)$$

Now, the buyers' incentive compatibility constraints can be stated formally:

Definition 2.2 (incentive compatibility (IC))

$$U_i(p, t, v_i) \geq E_{v_{-i}} [p_i(v_{-i}, w_i)v_i - t_i(v_{-i}, w_i)] \forall v_i, w_i \in [a_i, b_i], \quad \forall i \in N$$

The expression on the right side of the inequality is the expected utility of bidding w_i when the true valuation is v_i .

Another important property of auction mechanisms is individual rationality. An auction mechanism is *ex ante* individually rational if each buyer is no worse off participating than not—on an expected utility basis:

Definition 2.3 (individual rationality (IR))

$$U_i(p, t, v_i) \geq 0 \quad \forall v_i \in [a_i, b_i], \quad \forall i \in N$$

Same as with incentive compatibility constraint, the optimal mechanism turns to be (ex post) individually rational.

Different (individually rational, incentive compatible) auctions are usually evaluated either according to the expected utility of the seller or *efficiency* (aka *social welfare* in this setting where parties have quasi linear utility functions):

Definition 2.4 *The expected utility of the seller is*

$$U_0(p, t) = E_v \left[(q_0 - \sum_{i=1}^n p_i(v)) \cdot v_0 + \sum_{i=1}^n t_i(v) \right] = \int_V \left((q_0 - \sum_{i=1}^n p_i(v)) \cdot v_0 + \sum_{i=1}^n t_i(v) \right) f(v) dv \quad (2.2)$$

In the important special case where the seller's valuation for a unit of the good on sale is zero, then the seller's expected utility is the seller's expected revenue.

Definition 2.5 *Given the allocation rule $p(v)$, the expected social welfare is*

$$SW(p) = E_v \left[\sum_{i=1}^n p_i(v) \cdot v_i + (q_0 - \sum_{i=1}^n p_i(v)) \cdot v_0 \right] = \int_V \left(\sum_{i=1}^n p_i(v) \cdot v_i + (q_0 - \sum_{i=1}^n p_i(v)) \cdot v_0 \right) f(v) dv \quad (2.3)$$

When the objective of the auctioneer is efficiency (i.e., allocating the items to the parties who values them the most), the optimal mechanism is given by the Vickrey auction [Vi61]. The Myerson auction is optimal when the objective is the seller's expected utility (see [My81] for a single-unit case and [MR89] for a multi-unit case). These two objectives are generally in conflict, and cannot be maximized with one mechanism. In many real-world settings—such as privatization and competing electronic marketplaces—it is not clear that the objective should be either efficiency or expected utility of the seller.

One way to account for importance of both objectives is to set up a constrained optimization problem: optimize one of the objectives, subject to a constraint on the other. In this paper we derive a new auction mechanism that maximizes expected social welfare *subject to a minimum constraint on the seller's expected utility*. We now give the formal statement of this auction design problem:

Problem 2.1 *Maximize $SW(p)$ subject to the following constraints:*

1. *Constraint on the seller's expected utility:*

$$U_0(p, t) \geq R_0 \quad (2.4)$$

2. *The usual probability normalization (PN) constraints:*

$$\begin{cases} 1 \geq p_i(v) \geq 0, & \forall i \in N, \quad \forall v \in V \\ \sum_{i=1}^n p_i(v) \leq q_0 \end{cases} \quad (2.5)$$

3. *Incentive compatibility (IC) constraints:*

$$\begin{aligned} U_i(p, t, v_i) &\geq E_{v_{-i}} [p_i(v_{-i}, w_i) v_i - t_i(v_{-i}, w_i)] \\ \forall v_i, w_i &\in [a_i, b_i], \quad \forall i \in N \end{aligned} \quad (2.6)$$

4. *Individual rationality (IR) constraints:*

$$U_i(p, t, v_i) \geq 0 \quad \forall v_i \in [a_i, b_i], \quad \forall i \in N \quad (2.7)$$

We call a solution (mechanism) to this problem a welfare maximizing R_0 -seller's expected utility (WM- R_0) auction. Constraints IC, IR and PN are referred to as feasibility constraints [My81].

The main difference between this problem and Myerson's seller's expected utility maximization is the choice of the objective function. In Problem 2.1 the objective is efficiency and the seller's expected utility appears as an additional constraint, while in [My81] and [MR89] the objective is the seller's expected utility.

The problem is easier in the symmetric case where the valuations of different buyers come from the same probability distribution ($f_i = f_j, \forall i, j \in N$). In that setting, both social welfare and seller's expected utility are maximized by second-price auctions with reserve prices.⁶ The two auctions differ by the value of the reserve price: in the welfare

⁶The second-price auction with a reserve price is defined by the following allocation and payment rules:

maximizing auction the reserve price equals the seller's valuation for a unit of good. In the auction that maximizes seller's expected utility, the reserve price is generally greater than that. However, the set of buyers, receiving items in the seller's utility maximizing auction is a subset of the set of buyers, receiving items in the welfare maximizing auction.

In the asymmetric case where the densities f_i are not the same, the mechanism that maximizes welfare and the one that maximizes seller's expected utility differ fundamentally. Myerson showed that the latter sometimes allocates the items on sale to bidders other than the highest bidders (we discuss Myerson's auction in Section 3.2). With respect to our problem, the Myerson auction introduces the following complication: mechanisms optimizing seller's expected utility and those optimizing welfare might yield different allocations.

3 Designing the optimal mechanism

It can be shown that due to the particular form of the functionals $SW(p)$ and $U_0(p, t)$ the problem has a solution with one free parameter, whose optimal value can be easily found numerically.

Before deriving the mechanism in detail, we present the high-level ideas of the derivation, in order:

1. We show that the optimal payment rule t is the same as in the Myerson auction. We also demonstrate that the problem of designing the optimal mechanism (p, t) can be reduced to an optimization problem in p only. The optimal allocation rule is computed by maximizing $SW(p)$ subject to the constraint

$$\hat{U}_0(p) \geq R_0 \tag{3.1}$$

where $p(t)$ are valid probability distributions such that the mechanism (p, t) is incentive compatible and individually rational. Here $\hat{U}_0(p)$ is a linear functional that depends only on the allocation rule $p(t)$, and not on the payment rule t . (The expression is given in Sec. 3.1.)

2. Constraint 3.1 is either inactive (i.e. the unconstrained global maximum satisfies it) or is satisfied with equality at the maximum.⁷
 - (a) If Constraint 3.1 is inactive, the optimal auction is the standard Vickrey auction with the reserve price set to equal the seller's valuation for a unit of the good on sale.
 - (b) If Constraint 3.1 is active, we do the following:
 - i. We solve the problem using Lagrangian relaxation. The optimal allocation rule $p^{opt}(v)$ is found as the maximum of

$$\hat{L}(p, \lambda) = SW(p) + \lambda \cdot (\hat{U}_0(p) - R_0) \tag{3.2}$$

with respect to (p, λ) on the convex set of feasible allocation rules $p(v)$. In Sec. 3.2 we argue that this indeed yields a solution to the original optimization problem. We derive the allocation rule $p^\lambda(v)$ that, for given λ , maximizes $\hat{L}(p, \lambda)$.

- ii. This way the problem reduces to finding a value λ so that p^λ maximizes the objective $SW(p)$ subject to Constraint 3.1. Since Constraint 3.1 is active, the maximum can be found by solving the following integral equation:

$$\hat{U}_0(p^\lambda) = R_0$$

We prove that $\hat{U}_0(p^\lambda)$ is non-decreasing and continuous in λ . This allows us to find the optimal λ with a simple numerical algorithm, as we explain in Section 3.3.

The following subsections present the derivation of the optimal mechanism in detail.

-
- Each of the q_0 highest bidders gets one unit of the good, provided that his bid exceeds the reserve price.
 - Every bidder getting a unit of good, pays the maximum of the first unsuccessful bid and the reserve price. The other bidders pay 0.

Since we only consider feasible auctions, the bid of the buyer is equal to her private valuation for the good. The reserve price is a threshold value, such that all the bids below it are ignored.

⁷The feasibility constraints PN , IC , IR and Constraint 3.1 are linear in p , so the feasible region is convex. Since the objective $SW(p)$ is a linear functional, the maximum is attained on the boundary of the region, and if a constraint is active, it must be satisfied with equality in extremum. While this is analogous to linear programming, the problem is more complex because the optimization is over functions rather than variables.

3.1 New formulation of the optimization problem

In this subsection we reduce the problem of designing the optimal mechanism (p, t) to an optimization problem in p only. The derivation relies on the following lemma from [My81] and [MR89], which we state without proof.

Lemma 3.1 (Myerson and Maskin, Riley) *The expected utility of the seller in any feasible multi-unit auction is given by*

$$U_0(p, t) = \sum_{i=1}^n E_v \left[(v_i - v_0 - \frac{1 - F_i(v_i)}{f_i(v_i)}) p_i(v) \right] + q_0 \cdot v_0 - \sum_{i=1}^n U_i(p, t, a_i) \quad (3.3)$$

[MR89] proves a more general result about multi-unit auctions with general (not necessarily unit) demand. Restricting the buyers' demand functions to single-unit demand yields Lemma 3.1. The special case of the above lemma (with one item on sale) is stated as Lemma 3 in [My81].

Lemma 3.1 and the revenue equivalence theorem (see [My81] and [KP98] for generalized version) imply that the expected utility of the seller in a feasible WM- R_0 auction is given by

$$\hat{U}_0(p^{opt}) = \sum_{i=1}^n E_v \left[(v_i - v_0 - \frac{1 - F_i(v_i)}{f_i(v_i)}) p_i(v) \right] + q_0 v_0 \quad (3.4)$$

where p^{opt} is a solution to Problem 2.1. The payment of buyer i in a feasible WM- R_0 auction is given by

$$t_i(v) = p_i(v) v_i - \int_{a_i}^{v_i} p_i(v_{-i}, w_i) ds_i \quad (3.5)$$

We call $\hat{U}_0(p)$ the *pseudo-utility* of the seller. Trivially, $\hat{U}_0(p)$ and $U_0(p, t)$ are not equal for arbitrary t and p . However, by revenue equivalence

$$\hat{U}_0(p^{opt}) = U_0(p^{opt}, t)$$

Note that the expression for the seller's utility in the constrained optimum that is given by (3.4) does not involve t .

We now can restate the original optimization problem in terms of the allocation rule p only.

Problem 3.1 *Maximize $SW(p)$ subject to*

1. *Pseudo-utility (PU) constraint*

$$\hat{U}_0(p) = \sum_{i=1}^n E_v \left[(v_i - v_0 - \frac{1 - F_i(v_i)}{f_i(v_i)}) p_i(v) \right] + q_0 v_0 \geq R_0 \quad (3.6)$$

2. *Monotonicity condition - the expected probability of buyer i getting an item -*

$$E_{v_{-i}} p_i(v_i, v_{-i}) = \int_{V_{-i}} p_i(v) f_i(v_{-i}) dv_{-i} \quad (3.7)$$

is non-decreasing in v_i .

3. *Probability normalization constraints (PN).*

Again, correctness of this approach can be justified by revenue equivalence theorem.

We now solve the auction design problem using a form of Lagrangian relaxation. We find a solution to Problem 3.1 by computing the saddle point of the following functional:

$$\hat{L}(p, \lambda) = SW(p) + \lambda \cdot (\hat{U}_0(p) - R_0) \quad (3.8)$$

in (p, λ) , where p is restricted to the convex set of valid feasible allocation rules.

This approach of moving just one of the constraints into the objective is not the standard way of using Lagrangian relaxation, but is nevertheless valid for the following reason. The full Lagrangian corresponding to Problem 3.1 (and containing the terms corresponding to the feasibility constraints and the PU constraint (3.6)) is exactly the same as the Lagrangian corresponding to maximizing the objective (3.8) subject to the feasibility constraints only. And by the Kuhn-Tucker Theorem, the saddle point of the Lagrangian is a solution to Problem 3.1.

Problem 3.1 can be solved with the following algorithm:

1. For all λ , find the allocation rule $p^\lambda(v)$ that maximizes (3.8).
2. Find the optimal λ .⁸

The next two subsections explain the implementation of the two steps of this algorithm, respectively.

3.2 Finding the optimal allocation rule p^λ

We now present the mechanism for maximizing $\hat{L}(p, \lambda)$ with respect to p , for given λ . We first transform \hat{L} into a more convenient form:

$$\begin{aligned} \hat{L}(p, \lambda) &= E_v \left[\sum_{i=1}^n p_i(v) \cdot v_i + (q_0 - \sum_{i=1}^n p_i(v)) \cdot v_0 \right] + \lambda \left(\sum_{i=1}^n E_v \left[(v_i - v_0 - \frac{1 - F_i(v_i)}{f_i(v_i)}) p_i(v) \right] + \right. \\ &\left. q_0 v_0 - R_0 \right) = \int_V \sum_{i=1}^n \left((c_i^\lambda(v) - (1 + \lambda)v_0) p_i(v) \right) f(v) dv + (1 + \lambda)q_0 v_0 - \lambda R_0 \end{aligned} \quad (3.9)$$

We call the quantities c_i^λ *virtual valuations*.

Definition 3.1 For buyer i with valuation v_i drawn from distribution F_i , the virtual valuation v_i is

$$c_i^\lambda(v) = c_i^\lambda(v_i) = (1 + \lambda)v_i - \lambda \cdot \frac{1 - F_i(v_i)}{f_i(v_i)}$$

We now deviate from the main flow of this subsection and briefly compare our mechanism to that of Myerson, who introduced the use of virtual valuations in his analysis of expected revenue maximizing auctions. One of the differences between his auction and ours is in the form of virtual valuations. In the Myerson auction, a buyer's virtual valuation $c_i(v_i)$ is the difference between the buyer's real valuation v_i and the hazard rate:

$$c_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} \quad (3.10)$$

The Myerson auction, operating on those virtual valuations $c_i(v_i)$ rather than on real bids, is biased in favor of disadvantaged buyers [Kr02, p.73], thus creating an artificial competition between “weak” and “strong” buyers.⁹ Such a mechanism allows the auctioneer to set a high sell price for a strong buyer while motivating him to stay truthful even if he is sure that his valuation exceeds any possible valuation of any other buyer. This approach provides the auctioneer with higher expected revenue. In our case, the virtual valuations depend on an additional parameter λ which controls the tradeoff between expected social welfare and seller's expected utility. For ease of comparison, our virtual valuations $c_i^\lambda(v_i)$ can be written as

$$c_i^\lambda(v_i) = (1 + \lambda) \left(v_i - \frac{\lambda}{1 + \lambda} \cdot \frac{1 - F_i(v_i)}{f_i(v_i)} \right) \quad (3.11)$$

We now return to the derivation of the optimal allocation rule. Despite the differences in virtual valuations and functionals being optimized, for any given λ , the optimal allocation rule $p^\lambda(v)$ can be derived similarly to the one in [My81] and [MR89], without substantial changes in the argument.

⁸When $\lambda = 0$, the PU constraint (3.6) is inactive and the optimum is attained at the Vickrey auction with no reserve price.

⁹The terms “weak” and “strong” refer to buyers' valuation *distributions*. Distributions of “strong” buyers are concentrated around higher values.

In order to describe the optimal allocation rule we need a few more definitions. Let $\hat{c}_i^\lambda(v_i)$ be the closest *non-decreasing* continuous approximation for $v_i(v_i)$. Formally,

$$\hat{c}_i^\lambda(v_i) = \frac{d}{dq} G_i^\lambda(q), \quad \text{where } q = F_i(v_i) \quad (3.12)$$

$G_i^\lambda(q)$ is the lower convex hull of the function $H_i : [0, 1] \rightarrow \mathbb{R}$, defined as

$$H_i^\lambda(q) = \int_0^q c_i^\lambda(F_i^{-1}(r)) dr \quad (3.13)$$

That is, $G_i^\lambda(q)$ is the highest convex function on $[0, 1]$,¹⁰ such that

$$G_i^\lambda(q) \leq H_i^\lambda(q)$$

Redefined this way, virtual valuations are "ironed" over certain portions of the domain. The ironing procedure transforms the virtual valuations into non-decreasing functions, which is necessary to preserve incentive compatibility (see [My81] for the discussion).

Theorem 3.1 *For any λ , $\hat{L}(p, \lambda)$ is maximized when the allocation rule is given by*

$$p_i^\lambda(v) = \begin{cases} 1, & \text{if } \hat{c}_i^\lambda(v_i) \text{ is among } q_0 \text{ highest virtual valuations} \\ & \text{and } \hat{c}_i^\lambda(v_i) > (1 + \lambda)v_0 \\ 0, & \text{otherwise.} \end{cases} \quad (3.14)$$

Ties are broken by randomizing.

The proof follows that in [MR89] or [My81] (except that it is for different virtual valuations).

Theorem 3.1 proves that for all λ , the optimal mechanism is integral, that is, the probability of a buyer getting an item is always 0/1 (except for the case of ties). The following provides some intuition behind the proof: integral allocation mechanisms can be thought of as "corners" of the feasible region, given by constraints. Since the objective is a linear functional and all the constraints are also linear in p , analogy with the linear programming problem suggests that the optimum should be in one of those "corners". It should be noted, that although the feasibility constraints (2.6, 2.7) are on expected utility of the buyer, Mechanism 3.14 is (ex post) incentive compatible and individually rational. The latter can be verified by substituting the payment rule (3.5) in the expression of buyer's utility.

The payment rule (3.5) means that the winning buyer pays the minimum valuation, required to win the item. Therefore Mechanism 3.14 can be viewed as a form of Vickrey auction, run on virtual, rather than real valuations. When $\lambda = 0$, Mechanism (3.14) yields the standard Vickrey auction.

3.3 Finding the optimal value for Lagrange multiplier λ

Computing the optimal mechanism requires choosing a λ_{opt} so that the allocation rule $p^{\lambda_{opt}}$ maximizes expected social welfare over all allocation rules p^λ , which satisfy the PU constraint (3.6). (There is no need to account for feasibility constraints, since all p^λ satisfy them.) As we argued in the beginning Section 3, the optimal allocation rule satisfies (3.6) with equality and λ_{opt} is the solution to the integral equation

$$\tilde{U}_0(\lambda_{opt}) = R_0 \quad (3.15)$$

where \tilde{U}_0 is defined as

$$\tilde{U}_0(\lambda_{opt}) = \hat{U}_0(\hat{p}^{\lambda_{opt}})$$

Although $\tilde{U}_0(\lambda)$ can only be evaluated numerically, the following theorem states that \tilde{U}_0 is nicely behaved. The importance of this fact is that it makes Equation (3.15) easy to solve numerically (we present an algorithm for doing so in the next subsection).

Theorem 3.2 *$\tilde{U}_0(\lambda)$ is continuous and non-decreasing in λ .*

The proof of the theorem essentially relies on the fact that Mechanism 3.14 is (ex post) incentive compatible (the full proof is in the Appendix).

¹⁰See [My81] and [Ro96, p.36] for details.

3.4 Algorithm for computing the optimal allocation rule

The optimal allocation rule can be computed with the following algorithm, where the optimal λ_{opt} is the root of $\tilde{U}_0(\lambda) - R_0$. It is easy to find because \tilde{U}_0 is continuous and increasing.

Algorithm 1 (Computing the optimal allocation rule)

1. Check whether the Vickrey auction with reserve price equal to the seller's valuation satisfies the PU constraint (3.6). If it does, output it as a solution. Otherwise, go to 2.
2. Set λ_{min} to zero and λ_{max} to some positive number, such that $p^{\lambda_{max}}$, given by Mechanism (3.14), does not satisfy the PU constraint. The PU constraint is checked by evaluating $\tilde{U}_0(\lambda_{max})$ numerically.
3. Repeat the following step until $\lambda_{max} - \lambda_{min}$ converges to zero:
Set $\lambda_{new} = \frac{\lambda_{min} + \lambda_{max}}{2}$. Construct the allocation rule $p^{\lambda_{new}}$ and check the PU constraint. If it is satisfied, set $\lambda_{min} = \lambda_{new}$, otherwise set $\lambda_{max} = \lambda_{new}$.
4. Set the allocation rule to $p^{\lambda_{opt}}$, where λ_{opt} is given by Step 3. Set the payment rule t according to 3.5:

$$t_i(v) = p_i^{\lambda_{opt}}(v)v_i - \int_{a_i}^{v_i} p_i^{\lambda_{opt}}(v_{-i}, w_i) ds_i$$

\tilde{U}_0 is an n -dimensional integral over the allocation probabilities p^λ , but nevertheless it can be evaluated numerically. It can be estimated by sampling v and computing the Monte-Carlo sum, which converges to the true value of the integral. Sampling v is easy because the valuations v_i are drawn independently: v can be obtained by sampling individual valuations v_i from the distributions f_i .¹¹

4 Randomized mechanisms

When virtual valuations of the buyers c_i^λ given by (3.11) are non-decreasing and no ironing is needed,¹² there exist simple randomized mechanisms where the Vickrey rules are used w.p. t and the Myerson rules w.p. $1 - t$ that yield the same expected social welfare and seller's expected revenue as our mechanism.

Rather than computing the optimal λ using the algorithm above and running Mechanism (3.14), the seller can use the distributions f_i to evaluate the expected revenue of the Myerson auction and the Vickrey auction in advance, and use these revenue values to analytically determine t . When no ironing is applied this indeed yields the same expected revenue and social welfare as Mechanism 3.14.

However, randomization is often undesirable. In many settings an auction is only run once. For instance, each privatization auction usually has different participants and/or a different object (company) for sale. Similarly, in Internet auctions, the set of buyers generally differs for every auction, as may the object for sale. Now, say that in a given setting, the auction designer is unsatisfied with the seller's expected utility in the Vickrey auction, and with the expected social welfare of the Myerson auction. Still, the designer can be satisfied with the seller's expected utility and expected social welfare in our deterministic mechanism. So, the deterministic mechanism is satisfactory, but the randomized mechanism would run an unsatisfactory auction for sure. (For randomization to really make sense, the designer would have to be able to repeat the random drawing multiple times, i.e., to repeat the same auction in the same setting.)

5 Prior-free mechanisms

The results of this paper are derived under the independent private values model, which assumes that the distributions of valuations (types) are common knowledge. However, in many practical application this assumption does not hold and the seller has no information about these distributions (for instance when the good on sale is new to the market).

¹¹If one can solve for the inverse of F_i , it is possible to sample directly from f_i . Otherwise a technique such as importance sampling [GC95, p. 305] should be used.

¹²this is called a *regular* case in economics literature

This raises the question of the possibility of prior-free mechanisms with desired properties, which is also motivated by the *Robert Wilson doctrine*. In the special case of symmetric valuations (the valuations of all buyers come from the same distribution F) and with more than one unit of the good on sale (i.e. $q_0 > 1$) it is possible to design a mechanism which does *not* use any prior knowledge about distributions, but still approaches the characteristics of WM- R_0 auction, derived in this paper, when the number of buyers is large.

Such a prior-free mechanism can be constructed, applying the idea of sampling, developed in [BV03]: since the valuations of buyers are independent draws from the same *unknown* distribution F , they can be used to estimate F , using the following technique:

$$F_n(t) = \frac{|j \in N, v_j < t|}{n - 1}$$

Then the estimates F_n can be used in place of the true distribution F in Mechanism (3.14) (the derivative of F , f can also be estimated from the sample, as discussed in [BV03]). The one issue remaining is incentive compatibility: when the reported valuation of buyer i is used to estimate F and f there exist an incentive to lie (since through F and f , buyer i gets a chance to manipulate the estimates of the virtual valuations for herself and the buyers, she is competing with. A solution to this problem is to apply the idea of dual-sampling auction from [FGHK02]:

1. Split the set of buyers randomly into two groups of size $\frac{n}{2}$ (if n is odd, place the remaining buyer into the first group).
2. Assign half of the units on sale ($\frac{q_0}{2}$ units) to group 1 and the other half to group 2.
3. Use the valuations of the buyers from group 1 to estimate F and f for the buyers from group 2 and vice versa.
4. Independently run the auctions among the buyers from group 1 and group 2.

Note that this way, buyers from different groups do not compete against each other and the report of any buyer does not affect the estimates of F and f used in her group. This property obviously makes the auction incentive-compatible. By the argument, using the law of large numbers, the expected efficiency of this mechanism approaches that of feasible WM- R_0 auction and expected revenue approaches R_0 when n (the number of buyers) goes to infinity (the proof closely follows that in Theorem 4 from [BV03]).¹³

The same idea can be trivially extended for the non-symmetric case (buyers' valuations are draws from different distributions), when there is sufficient number of buyers coming from each distribution.

6 Conclusions

We demonstrated that the auction, maximizing expected efficiency subject to a constraint on the seller's expected utility belongs to the family of dominant strategy mechanisms, parameterized with one free parameter and presented an algorithm for optimally choosing the value of the parameter. By running this auction the seller can expect to do well enough for himself, and maximize social welfare subject to that.

No matter how the tradeoff between efficiency and seller's utility is struck, the optimal auction has essentially the same form. Furthermore, except for the case of ties in virtual valuations the optimal auction does not rely on randomization, that is, given distinct bids the allocation probabilities are 0/1 (however, the winning bidder is not always the highest bidder).

It is important that the presented mechanism does not use randomization over different auction mechanisms (i.e. when the result is achieved by running different auctions with positive probability). The latter is inappropriate for the auction design problems, which motivate our work.

We derived our results under the asymmetric independent private values model. For the special case of symmetric distributions we also presented the approximation mechanism which does *not* use any prior knowledge about distributions. This mechanism approaches expected efficiency and expected seller's utility of the auction, designed with distributions known upfront, when the number of buyers, coming from every distribution tends to infinity.

¹³It is well known that in symmetric setting with fixed number of goods and the number of bidders going to infinity, the Vickrey auction approaches the optimal revenue (for instance, [MT00] prove this statement for general combinatorial auctions). However the statement does not hold when the number of goods on sale is also increasing. This is exactly the case when the approximation technique of [BV03] is helpful.

References

- [BV03] Baliga S. and Vohra R. *Market Research and Market Design*, Advances in Theoretical Economics: Vol. 3: No. 1, 2003.
- [CS02] Conitzer, V., and Sandholm, V. *Complexity of Mechanism Design*, UAI 2002.
- [FGHK02] Fiat A., Goldberg A., Hartline J. and Karlin A., *Competitive Generalized Auctions*, STOC 2002.
- [GC95] Gelman A.B., Carlin J.S., Stern H.S., Rubin D.B. *Bayesian data analysis*, Chapman and Hall/CRC, 1995.
- [Kr02] Krishna V. *Auction theory*, Academic Press, 2002.
- [KP98] Krishna V., Perry M. *Efficient mechanism design*, Working Paper, 1998.
- [MR89] E. S. Maskin and J. G. Riley. *Optimal multi-unit auctions*. The Economics of Missing Markets, Information, and Games. Oxford University Press, Clarendon Press, 1989.
- [MT00] Monderer D. and Tennenholtz M. *Asymptotically Optimal Multi-Object Auctions for Risk-Averse Agents*, manuscript, 2000.
- [My81] Myerson, R. *Optimal auction design*, In Mathematics of operational research. Vol. 6, pp. 58-73, 1981.
- [Ro96] Rockafellar, R.V. *Convex Analysis*, Princeton University Press, 1996.
- [Vi61] Vickrey W. *Counterspeculation, auctions and competitive sealed tenders*. Journal of Finance, 16 pp.8-37, 1961.

Appendix

Proof of Theorem 3.2. Monotonicity of $\tilde{U}_0(\lambda) = \hat{U}_0(p^\lambda)$ can be proved as follows: Take arbitrary $\lambda_1 < \lambda_2$. Then, by definition of $p^{\lambda_1}, p^{\lambda_2}$

$$\begin{cases} \hat{L}(p^{\lambda_1}, \lambda_1) \geq SW(p^{\lambda_2}) + \lambda_1 \cdot \hat{U}_0(p^{\lambda_2}) \\ \hat{L}(p^{\lambda_2}, \lambda_2) \geq SW(p^{\lambda_1}) + \lambda_2 \cdot \hat{U}_0(p^{\lambda_1}) \end{cases} \quad (6.1)$$

Denote $\Delta p = p^{\lambda_2} - p^{\lambda_1}$ and $\Delta \lambda = \lambda_2 - \lambda_1$. Since SW, \hat{U}_0 and \hat{L} are linear functionals, (6.1) implies

$$\begin{cases} \hat{L}(-\Delta p, \lambda_1) = -\hat{L}(\Delta p, \lambda_1) \geq 0 \\ \hat{L}(\Delta p, \lambda_2) \geq 0 \end{cases}$$

Therefore

$$\hat{L}(\Delta p, \lambda_2) - \hat{L}(\Delta p, \lambda_1) = (\lambda_2 - \lambda_1) \hat{U}_0(\Delta p) \geq 0$$

and

$$\hat{U}_0(\Delta p) = \hat{U}_0(p^{\lambda_2}) - \hat{U}_0(p^{\lambda_1}) \geq 0$$

Thus, $\tilde{U}_0(\lambda) = \hat{U}_0(p^\lambda)$ is non-decreasing in λ .

We now prove the continuity of $\tilde{U}_0(\lambda)$. The idea of the proof uses the (ex-post) incentive compatibility of the Mechanism 3.14 and shows that if $\tilde{U}_0(\lambda)$ is not continuous, the mechanism cannot be ex post incentive compatible.

\tilde{U}_0 is given by

$$\tilde{U}_0(\lambda) = \sum_{i=1}^n \left(\int_V (c_i(v_i) - v_0) p_i^\lambda(v) f(v) dv \right) + q_0 \cdot v_0$$

where c_i are the Myerson virtual valuations, defined in (3.10). For arbitrary λ_1, λ_2 , we have

$$|\tilde{U}_0(\lambda_1) - \tilde{U}_0(\lambda_2)| \leq c_{max} \int_V \sum_{i=1}^n |p_i^{\lambda_1}(v) - p_i^{\lambda_2}(v)| f(v) dv \leq 2 \cdot c_{max} \cdot \int_{V_{\lambda_1, \lambda_2}} f(v) dv \quad (6.2)$$

where

$$c_{max} = \max_{i \in N, v_i \in [a_i, b_i]} |c_i(v_i) - v_0|$$

and

$$V_{\lambda_1, \lambda_2} = \{v \in V \mid p^{\lambda_1}(v) \neq p^{\lambda_2}(v)\}$$

is the set of valuation vectors where allocations $p^{\lambda_1}(v)$ and $p^{\lambda_2}(v)$ differ.

We now show that $V_{\lambda_1, \lambda_2} = O(\Delta\lambda)$. By (3.11), for all i and v_i ,

$$|v_i^{\lambda_1}(v_i) - v_i^{\lambda_2}(v_i)| = |c_i(v_i)\Delta\lambda| \leq |c_{max}\Delta\lambda|$$

Also, using the definition of \hat{c}_i it is easy to show that for all i and v_i

$$\begin{cases} \hat{c}_i^{\lambda_1}(v_i) - |c_{max}\Delta\lambda| \leq \hat{c}_i^{\lambda_2}(v_i) \\ \hat{c}_i^{\lambda_2}(v_i) - |c_{max}\Delta\lambda| \leq \hat{c}_i^{\lambda_1}(v_i) \end{cases}$$

Therefore,

$$|\hat{c}_i^{\lambda_1}(v) - \hat{c}_i^{\lambda_2}(v)| \leq |c_{max}\Delta\lambda| \quad (6.3)$$

Mechanisms, using the payment rule (3.5) and with allocation rule, such that all $p_i(v_{-i}, s)$ are non-decreasing in s are *ex-post* incentive compatible (it can be verified by substituting the payment rule (3.5) into the definition of buyer's utility). Therefore, Mechanisms (3.14) are *ex-post* incentive-compatible. We show now that *ex-post* incentive compatibility together with (6.3) implies $V_{\lambda_1, \lambda_2} = O(\Delta\lambda)$ and therefore continuity of \tilde{U}_0 .

We first demonstrate that utility of buyer i -

$$u_i(p^\lambda, t, v) = p_i^\lambda(v) \cdot v_i - t_i$$

is continuous in parameter λ for all λ and for all v . Assume the contrary: there exists v , λ_1 and ϵ_{λ_1} , such that for all $\delta > 0$, there exist some λ_2 , satisfying

$$\begin{cases} |\lambda_2 - \lambda_1| < \delta \\ |u_i(p^{\lambda_1}, t, v) - u_i(p^{\lambda_2}, t, v)| > \epsilon_{\lambda_1} \end{cases}$$

We now show that such a mechanism is not *ex-post* incentive compatible. W.l.o.g. assume that the utility of buyer i under the allocation p^{λ_2} is higher than under the allocation p^{λ_1} (if this is not the case, interchange λ_1 and λ_2).

Non-continuity of u_i means that arbitrary small changes in λ yield substantial (at least ϵ_{λ_1}) increase in utility of buyer i . Mechanism (3.14) allocates the items to buyers with highest virtual valuations and by (6.3) the virtual valuations of *all* buyers change by at most $|c_{max}\Delta\lambda|$. Since $\hat{c}_i(v_i)$ is non-decreasing in v_i , there exists a type \hat{v}_i , such that when types of other buyers are given by v_{-i} , buyer i benefits from overbidding (i.e. mechanism is not *ex-post* IC).

More formally, by (6.3) $|\hat{c}_j^{\lambda_1}(v_j) - \hat{c}_j^{\lambda_2}(v_j)| \leq |c_{max}\Delta\lambda|$ for all j . Therefore, if $\hat{c}_i^{\lambda_1}(v_i)$ were at most $2 \cdot |c_{max}\Delta\lambda|$ higher, buyer i would get the same probability of winning an item as under allocation p^{λ_2} . Since $\hat{c}_i^{\lambda_1}(v_i)$ is a non-decreasing function of v_i , consider two cases:

1. $\hat{c}_i^{\lambda_1}(v_i)$ is increasing at v_i (i.e. $\hat{c}_i^{\lambda_1}(v_i) = c_i^{\lambda_1}(v_i)$). The derivative of $\hat{c}_i^{\lambda_1}$ is well-defined, positive and continuous in v_i (i.e. it is positive in some neighborhood of v_i). Therefore it is possible to choose $\Delta\lambda = |\lambda_2 - \lambda_1|$ small enough, so that there exist t'_i , such that $0 < t'_i - v_i < \tau_{v_i}^{\lambda_1} \cdot \Delta\lambda$, for some constant $\tau_{v_i}^{\lambda_1}$ and $\hat{c}_i^{\lambda_1}(t'_i) > c_i^{\lambda_1}(v_i) + 2 \cdot |c_{max}\Delta\lambda|$.

If the true type of buyer i is v_i , then reporting t'_i yields an increase in utility of at least ϵ_{λ_1} (due to the increase of the probability of winning), while the payment of the bidder increases by at most $2\tau_{v_i}^{\lambda_1} \cdot \Delta\lambda$ (this can be verified by substituting v_i and \hat{v}_i into the payment rule (3.5)). Therefore if $\Delta\lambda$ is sufficiently small, bidder i benefits from overbidding.

2. $\hat{c}_i^{\lambda_1}(v_i)$ is constant at v_i (i.e. v_i is on the flat (ironed) portion of the virtual valuation). Take \hat{v}_i to be the highest type, such that $\hat{c}_i^{\lambda_1}(\hat{v}_i) = \hat{c}_i^{\lambda_1}(v_i)$ (i.e. \hat{v}_i is at the end of the flat portion of $\hat{c}_i^{\lambda_1}$). It is easy to show that \hat{v}_i has the same probability of winning an item as v_i and the same utility. Applying then the same argument as in the case of increasing $\hat{c}_i^{\lambda_1}(v_i)$ allows to show that \hat{v}_i benefits from overbidding.

The argument yields a contradiction: (Mechanism (3.14) is not ex-post incentive compatible), which is due to our assumption about non-continuity of u_i . When the payment rule is set according to (3.5) the utility of buyer i is given by

$$u_i(p^\lambda, t, (v_{-i}, v_i)) = \int_{a_i}^{v_i} p_i(v_{-i}, w_i) ds_i$$

Therefore, since u_i is continuous in λ , and by (3.14) p_i takes values only in $\{0, \frac{1}{n}, \frac{1}{n-1}, \dots, 1\}$ and is non-decreasing in v_i for all i we must have

$$\{v_i \in [a_i, b_i] \mid p_i^{\lambda_1}(v_{-i}, v_i) \neq p_i^{\lambda_2}(v_{-i}, v_i)\} = O(\Delta\lambda), \quad \forall i, v_{-i}$$

It follows, that

$$V_{\lambda_1, \lambda_2} = \{v \in V \mid p^{\lambda_1}(v) \neq p^{\lambda_2}(v)\} = O(\Delta\lambda)$$

Therefore, by 6.2, $\tilde{U}_0(\lambda)$ is continuous in λ . \square