Recent Applications in Mechanism Design

Joseph Thompson

April 2016

Abstract

This paper will demonstrate heuristics for solving revenue-maximizing single parameter mechanism design problems. The bidders in each environment will share the same relative valuations for goods, and their utility functions will be linear with respect to what they are allocated. We will apply these heuristics for three auctions: k-Vickrey, sponsored search, and knapsack. Formalizing all of these notions, as well as techniques for solving welfare-maximizing design problems, will be developed along the way.

1 Introduction

Game theory is the study of finding optimal strategies given a collection of players, actions, and payoffs. Mechanism design is a field within game theory, but it focuses on optimal strategies from the rule-maker's perspective rather than the player's. In other words, given a collection of players and actions, mechanism design is the study of finding rules and payoffs that yield optimal outcomes. We assume that players will behave rationally, so the designer is more concerned about the preferences of players rather than their list of actions.

In this paper, we will focus on the applications of mechanism design through various auctions. Our auctions will have one seller, the auctioneer, who will also operate as the mechanism designer. The seller has one of two objectives: either she wants to maximize the total utility of everyone participating in the auction (including herself), or she wants to maximize her own personal utility. Regardless of her objective, the seller has to design an auction such that bidders are incentivized to both join the auction and play truthfully. If a bidder doesn't want to join the auction, then the auctioneer will have no customers and collect zero revenue. If a bidder doesn't play truthfully, then the auctioneer will be allocating under false preferences, where the outcome may be suboptimal. It is therefore crucial to design an auction such that the incentives of bidders align with the seller's expectations.

2 Definitions and Assumptions

Definition 2.1. A **mechanism environment** is a pair of sets (N, G), where N is the set of all bidders and G is the set of all goods. A **mechanism** M is an algorithm that does the following:

1. Queries information from bidders in N.
2. Allocates goods in G to bidders in N.
3. Collects payments from bidders in N.

Intuitively, a mechanism environment is the setting for an auction, and the mechanism is the auction itself. Bidders will meet the seller in an attempt to purchase some or all of the goods that the seller is selling. Given a mechanism environment, let n denote the number of bidders and g denote the number of goods. We assume for now that N and G are finite, so we will be able to refer to elements in these sets with an index. Let $i \in \{1, 2, ..., n\}$ be an index for an arbitrary bidder in N, and let $j \in \{1, 2, ..., g\}$ be an index for an arbitrary good in G. If there are multiple copies of the same good, G becomes a multiset. Note that a particular bidder i can receive more than one good from G. In fact, bidder i can receive any subset of goods $S \subseteq G$. Let $k \in \{1, 2, ..., 2^g\}$ be an index for an arbitrary bundle of goods in $\mathcal{P}(G)$, where $\mathcal{P}(G)$ denotes the power set of G.

Once our environment (N, G) is established, the mechanism performs the three steps mentioned above. Each step in the mechanism provides a challenge for the designer:

1. How will she collect information from the bidders?
2. Assuming all necessary information is collected, for any $j \in \{1, 2, ..., g\}$, to whom should she give good j?
3. Assuming all goods are allocated optimally, for any $i \in \{1, 2, ..., n\}$, what price should she charge bidder i?

The first challenge has been solved for any general mechanism in [4]. Suppose there is a mechanism M that collects all necessary information in whatever means. There is then another mechanism M' that collects the same information by having each bidder privately reveal their portion of the information to the seller. The class of mechanisms M' share a name:

Definition 2.2. Given an environment (N, G), a mechanism M is said to be **direct-revelation** if, for every $i \in \{1, 2, ..., n\}$, M collects bidder i's information by privately querying bidder i only once.

Since every mechanism M has an equivalent direct-revelation mechanism M', we assume that M is direct-revelation. Of course, the information received from bidders may not be truthful. The seller has to still construct allocation and payment schemes such that revealing truthful information is the incentive for every bidder. We now formalize our notion of "information" that each bidder i submits to M and what it means to be "truthful":

Definition 2.3. For any $i \in \{1, 2, ..., n\}$ and $k \in \{1, 2, ..., 2^g\}$, bidder i's **bid** for bundle k is a signal to M that denotes the maximum amount that bidder i will pay in exchange for bundle k. Bidder i's **valuation** for bundle k is the amount that bidder i would need to pay to be indifferent between receiving bundle k and receiving nothing.

Let $B_k^{(i)}$ denote bidder i's bid for bundle k, and let $\mathbf{B}^{(i)} = (B_1^{(i)}, B_2^{(i)}, ..., B_{2^g}^{(i)})$ be a vector of bids from bidder i. Similarly, let $V_k^{(i)}$ denote bidder i's valuation for bundle k, and let $\mathbf{V}^{(i)} = (V_1^{(i)}, V_2^{(i)}, ..., V_{2^g}^{(i)})$ be a vector of valuations from bidder i. The information that M will receive from bidder i is exactly $\mathbf{B}^{(i)}$, and we say that bidder i submitted truthful information (or "bid truthfully") if $\mathbf{B}^{(i)} = \mathbf{V}^{(i)}$.

Note that $\mathbf{B}^{(i)}$ need not equal $\mathbf{V}^{(i)}$. Each bidder has the objective to maximize their own utility, and the optimal bid for any good may be under or over bidder i's valuation. For example, consider a mechanism environment with one item and one bidder. Bundle 1 will correspond to not receiving the good, and bundle 2 will correspond to receiving the good. Suppose a mechanism M will allocate to the bidder and charge his bid $B_2^{(1)}$ if $B_2^{(1)} \geq 0$. Suppose the bidder receives zero utility if he receives bundle 1 and utility $5 - p$, where p denotes the price, when he receives bundle 2. The bidder is indifferent when $p = 5$, meaning the bidder's valuation for this good is $V_2^{(1)} = 5$. He'll want to bid such that he receives the good while paying the smallest amount possible. By the design of M, p equals the bidder's bid $B_2^{(1)}$. Hence, the bidder will set $B_2^{(1)} = 0$ even though his valuation $V_2^{(1)} = 5$.

If a mechanism M is receiving a bid vector $\mathbf{B}^{(i)}$ from every bidder, then M will be optimizing over $n2^g$ variables. That number grows exponentially in g, making our problem computationally expensive. To reduce the number of variables, we limit our space of auction problems to single parameter mechanisms from this point forward:

Definition 2.4. A mechanism environment is **single parameter** if for every $i \in \{1, 2, ..., n\}$ and every $k \in \{1, 2, ..., 2^g\}$, bidder i can express their valuation for bundle k through a private parameter v_i and a publicly known scalar $c_k^{(i)}$. In other words, for any $\mathbf{V}^{(i)} \in \mathbb{R}^{2^g}$, there exists a publicly known vector of scalars $\mathbf{c}^{(i)} = (c_1^{(i)}, c_2^{(i)}, ..., c_{2^g}^{(i)}) \in \mathbb{R}^{2^g}$ such that $\mathbf{V}^{(i)} = v_i\mathbf{c}^{(i)}$.

For each i, v_i is known as bidder i's **valuation parameter**. Let $\mathbf{v} = (v_1, v_2, ..., v_n)$ be a vector of valuation parameters, where v_i is the parameter for bidder i. We assume that for each i, v_i is a random variable that takes on a value in bidder i's **type space** T_i. T_i is the sample space for v_i, and the probability that v_i takes on some value at most $t \in T_i$ is denoted $F_i(t)$, where F_i is the cumulative distribution function of v_i. We assume that each T_i is closed and bounded. Let $\mathbf{T} = T_1 \times T_2 \times ... \times T_n$ be a cross product of type spaces, where T_i is the type space of bidder i.

For each bidder i, our mechanism M will know $\mathbf{c}^{(i)}$, T_i and F_i. M now only needs to query a single parameter b_i from each bidder because $b_i\mathbf{c}^{(i)}$ will equal bidder i's bid $\mathbf{B}^{(i)}$. We call b_i the **bid parameter** for bidder i. Note that $\mathbf{V}_i = \mathbf{B}_i \Leftrightarrow v_i\mathbf{c}^{(i)} = b_i\mathbf{c}^{(i)} \Leftrightarrow v_i = b_i$, so bidder i is truthful if $v_i = b_i$. Let $\mathbf{b} = (b_1, b_2, ..., b_n)$ be a vector of bid parameters, where b_i is the parameter for bidder i.

It will be useful to have notation to discuss parameters of every other bidder besides i. Let $\mathbf{v}_{-i}$, $\mathbf{b}_{-i}$, and $\mathbf{T}_{-i}$ denote a vector of valuation parameters, a vector of bid parameters, and a cross product of type spaces of all bidders after bidder i is removed. Let $(v_i, \mathbf{v}_{-i}) = \mathbf{v}$, $(b_i, \mathbf{b}_{-i}) = \mathbf{b}$, and $(T_i, \mathbf{T}_{-i}) = \mathbf{T}$.

Some single parameter environments allow us to simply our set of goods G quite nicely. If all goods in G are identical, then G is said to be **homogeneous**. For a set of homogeneous goods, we let G denote a set of numbers as opposed to a set of goods. When G is homogeneous, G will denote the sizes of feasible bundles for bidders and j will denote a quantity of the good as opposed to an index. This allows us to get rid of the notion of $\mathcal{P}(G)$ when G is homogeneous. For example, if an auctioneer were selling four identical goods, then $G = \{1, 2, 3, 4\}$. In the case of homogeneous goods, $g \neq |G|$ in general.

Let $A = [\mathbf{V}_1^\intercal, \mathbf{V}_2^\intercal, ..., \mathbf{V}_n^\intercal]$ denote the matrix of valuation vectors for an environment (N, G). Consider a mechanism environment (N, G) such that the rank of A is one. We can construct a seemingly identical environment (N, G') where G' is homogeneous. Let $\{\mathbf{c}\}$ be the basis for A. Each bidder i now shares the *same* vector of scalars $\mathbf{c}$, meaning each bidder's valuation is simply their valuation parameter v_i multiplied by $\mathbf{c}$. Consider the set $\{c_k : k \in \{1, 2, ..., 2^g\}\}$. G' is exactly that set; for a fixed good in G, every other good in G will be represented as a feasible bundle in G' consisting solely of the fixed good.

For example, consider $n = 2$ and G ={apple, orange}. Suppose bidder 1 values receiving nothing at zero, the apple at one, the orange at two, and both goods at four. Suppose bidder 2 values receiving nothing at zero, the apple at three, the orange at six, and both goods at twelve. This mechanism environment is essentially identical to one where four apples are being auctioned with feasible bundles being $\{0, 1, 2, 4\}$. This is because both bidders value the orange twice as much as the apple, and the opportunity to receive both the apple and the orange is valued at four apples. This was possible because the relative valuations for these goods were the same for each bidder. Their valuation vectors are off by a single parameter v_i, and in this example, v_i corresponds to the valuation of one apple for each bidder.

For the remainder of this paper, we assume that the valuation matrix A has rank 1 for a mechanism environment (N, G). Since we also assume that our environment is single parameter, we will furthermore assume that G is homogeneous. These assumptions will allow us to unify v_i, $\mathbf{c}^{(i)}$, and G with marvelous simplicity. For any i, bidder i's valuation parameter v_i will correspond to how much bidder i values 1 unit of the good. For any $j \in G$, j will correspond to both the public scalar and allocation of receiving j units. In other words, for all $i \in \{1, 2, ..., n\}$ and all $j \in G$, bidder i values receiving j units at jv_i. The terms in $\mathbf{c}^{(i)}$ are exactly the elements in G, so we get rid of the notion of $\mathbf{c}^{(i)}$. Now that we've formalized (N, G) and the input b for M, we shall define the outputs of M: allocation and payments.

Definition 2.5. An **allocation rule** is a function $\mathbf{x} : \mathbf{T} \to G^n$, where $x_i(\mathbf{b})$ denotes the amount of the good allocated to bidder i. A **payment rule** is a function $\mathbf{p} : \mathbf{T} \to \mathbb{R}^n$, where $p_i(\mathbf{b})$ denotes the amount bidder i pays to M.

Given a single parameter mechanism environment (N, G) where G is homogeneous, a mechanism M will take an input of bid parameters $\mathbf{b}$ and output an allocation and payment rule $(\mathbf{x}, \mathbf{p})$. The structure of $(\mathbf{x}, \mathbf{p})$ will depend on $\mathbf{b}$ and the objective of M. In general, M will either want to maximize **total revenue**, the sum of all payments collected by M, or **total welfare**, the sum of bidder's utilities and total revenue. M will maximize these objectives in expectation over $\mathbf{v}$. Each bidder i has a **utility function** u_i, and each bidder has the objective to maximize their own utility function. Bidder i's utility will depend on their valuation, their allocation, and their payment. Bidder i's valuation will be fixed, but allocation and payment rules depend on the bid parameters of all bidders. Therefore, bidder i's utility is a function of allocation, payment, and bid parameters. In other words, $u_i = u_i(\mathbf{x}, \mathbf{p}, \mathbf{b}) = u_i(\mathbf{x}, \mathbf{p}, b_i, \mathbf{b}_{-i})$. Assume that each bidder i has a utility function of the form $u_i(\mathbf{x}, \mathbf{p}, b_i, \mathbf{b}_{-i}) = v_i x_i(b_i, \mathbf{b}_{-i}) - h(p_i(b_i, \mathbf{b}_{-i}))$, where h is called the **cost function**. This cost function will be the same among all bidders.

Not only will M be constructing $(\mathbf{x}, \mathbf{p})$ to maximize an objective, but M must also adhere to certain constraints. M must construct allocation and payment rules that will incentive bidders to participate, incentivize bidders to bid truthfully, and not over-allocate. We define each of these constraints more formally:

Definition 2.6. An allocation and payment rule $(\mathbf{x}, \mathbf{p})$ satisfies **Individual Rationality (IR)** if for each $i \in \{1, 2, ..., n\}$, bidder i is at least as happy participating in the auction than otherwise. In other words, IR is satisfied when $\forall i \in \{1, 2, ..., n\}, \forall v_i \in T_i, \exists b_i \in T_i$ such that $u_i(\mathbf{x}, \mathbf{p}, b_i, \mathbf{b}_{-i}) \geq 0 \ \forall \mathbf{b}_{-i} \in \mathbf{T}_{-i}$.

Definition 2.7. An allocation and payment rule $(\mathbf{x}, \mathbf{p})$ satisfies **Incentive Compatibility (IC)** if for each $i \in \{1, 2, ..., n\}$, bidder i maximizes their utility when they bid truthfully. In other words, IC is satisfied when $\forall i \in \{1, 2, ..., n\}, \forall v_i \in T_i, u_i(\mathbf{x}, \mathbf{p}, v_i, \mathbf{b}_{-i}) \geq u_i(\mathbf{x}, \mathbf{p}, b_i, \mathbf{b}_{-i}) \ \forall b_i \in T_i, \forall \mathbf{b}_{-i} \in \mathbf{T}_{-i}$.

Definition 2.8. An allocation and payment rule $\mathbf{x}$ satisfies **Ex-Post Feasibility (XP)** if there are enough goods available to allocate to the bidders according to $\mathbf{x}$. In other words, XP is satisfied when $\sum_{i=1}^{n} x_i(\mathbf{b}) \leq g \ \forall \mathbf{b} \in \mathbf{T}$.

There is an additional constraint that requires $\forall i \in \{1, 2, ..., n\}, \ x_i(b_i, \mathbf{b}_{-i}) \geq 0 \ \forall b \in T_i, \forall \mathbf{b}_{-i} \in \mathbf{T}_{-i}$. This means that no bidder will ever give goods away; they will only receive goods. Note that if IC is satisfied, then bidders are truthful. In that case, we can assume the input $\mathbf{b}$ is actually $\mathbf{v}$. It is possible to relax the IR and IC constraints such that they only need to hold for each bidder i in expectation under $\mathbf{v}_{-i}$ as opposed to for all possible $\mathbf{v}_{-i}$. The relaxed IR constraint is called Bayesian Individual Rationality (BIR), and the relaxed IC constraint is called Bayesian Incentive Compatibility (BIC). In this paper, we will consider most design problems subject to IR, IC, and XP, but [1] shows that M's objective in the Bayesian setting will be at least as good as its objective in our setting.

With all terminology formalized, the following program is the revenue-maximizing design problem for M:

On input $\mathbf{v}$, construct $(\mathbf{x}, \mathbf{p})$ to maximize

$$\mathbb{E}_{\mathbf{v}}\left[\sum_{i=1}^{n} p_i(\mathbf{v})\right] \quad \text{subject to}$$

$$\forall i \in \{1,2,...,n\} \;\; x_i(b_i, \mathbf{b}_{-i}) \geq 0 \;\; \forall b \in T_i, \forall \mathbf{b}_{-i} \in \mathbf{T}_{-i}$$

$$\forall i \in \{1,2,...,n\}, \forall v_i \in T_i, \exists b_i \in T_i \text{ such that } u_i(\mathbf{x}, \mathbf{p}, b_i, \mathbf{b}_{-i}) \geq 0 \;\; \forall \mathbf{b}_{-i} \in \mathbf{T}_{-i} \quad \text{(IR)},$$

$$\forall i \in \{1,2,...,n\}, \forall v_i \in T_i, \;\; u_i(\mathbf{x}, \mathbf{p}, v_i, \mathbf{b}_{-i}) \geq u_i(\mathbf{x}, \mathbf{p}, b_i, \mathbf{b}_{-i}) \;\; \forall b_i \in T_i, \forall \mathbf{b}_{-i} \in \mathbf{T}_{-i} \quad \text{(IC), and}$$

$$\sum_{i=1}^{n} x_i(\mathbf{v}) \leq g \;\; \forall \mathbf{v} \in \mathbf{T} \quad \text{(XP).}$$

And this program is the welfare-maximizing design problem for M:

On input $\mathbf{v}$, construct $(\mathbf{x}, \mathbf{p})$ to maximize

$$\mathbb{E}_{\mathbf{v}}\left[\sum_{i=1}^{n} (u_i(\mathbf{x}, \mathbf{p}, v_i, \mathbf{v}_{-i}) + p_i(\mathbf{v}))\right] \quad \text{subject to}$$

$$\forall i \in \{1,2,...,n\} \;\; x_i(b_i, \mathbf{b}_{-i}) \geq 0 \;\; \forall b \in T_i, \forall \mathbf{b}_{-i} \in \mathbf{T}_{-i}$$

$$\forall i \in \{1,2,...,n\}, \forall v_i \in T_i, \exists b_i \in T_i \text{ such that } u_i(\mathbf{x}, \mathbf{p}, b_i, \mathbf{b}_{-i}) \geq 0 \;\; \forall \mathbf{b}_{-i} \in \mathbf{T}_{-i} \quad \text{(IR)},$$

$$\forall i \in \{1,2,...,n\}, \forall v_i \in T_i, \;\; u_i(\mathbf{x}, \mathbf{p}, v_i, \mathbf{b}_{-i}) \geq u_i(\mathbf{x}, \mathbf{p}, b_i, \mathbf{b}_{-i}) \;\; \forall b_i \in T_i, \forall \mathbf{b}_{-i} \in \mathbf{T}_{-i} \quad \text{(IC), and}$$

$$\sum_{i=1}^{n} x_i(\mathbf{v}) \leq g \;\; \forall \mathbf{v} \in \mathbf{T} \quad \text{(XP).}$$

3 Useful Tools

In this section, we assume the cost function h is the identity function. In other words, $u_i(\mathbf{x}, \mathbf{p}, b_i, \mathbf{b}_{-i}) = v_i x_i(b_i, \mathbf{b}_{-i}) - p_i(b_i, \mathbf{b}_{-i})$ for each bidder i. Under this assumption, total welfare is independent of payments:

$$\begin{aligned}\sum_{i=1}^{n}(u_i(\mathbf{x}, \mathbf{p}, b_i, \mathbf{b}_{-i}) + p_i(\mathbf{b})) &= \sum_{i=1}^{n}(v_i x_i(b_i, \mathbf{b}_{-i}) - p_i(b_i, \mathbf{b}_{-i})) + p_i(b_i, \mathbf{b}_{-i})) \\ &= \sum_{i=1}^{n}(v_i x_i(b_i, \mathbf{b}_{-i})).\end{aligned}$$

We have found a way to determine allocations and payments separately, effectively decoupling $\mathbf{x}$ and $\mathbf{p}$. This allows us to break our welfare-maximizing design problem into two subproblems. First, M will assume that its input $\mathbf{b}$ is equal to $\mathbf{v}$ and construct an allocation rule $\mathbf{x}$ to maximize welfare and satisfy XP. Then, given $\mathbf{x}$, M can construct a payment rule $\mathbf{p}$ to satisfy IR and IC, justifying M's assumption that the input $\mathbf{b} = \mathbf{v}$.

The allocation subproblem is an optimization problem which can usually be solved or approximated with a greedy algorithm. Finding the optimal $\mathbf{x}$ will depend on the mechanism environment, so we assume that M can construct the welfare-maximizing $\mathbf{x}$ in general. Note that $x_i(b_i, \mathbf{b}_{-i})$ is non-decreasing in b_i for any i. This is because M assumes that $b_i = v_i$, and M would not want to take items away from bidder i if bidder i valued the good more than before; that would go against M's welfare-maximizing strategy. If any $x_i(b_i, \mathbf{b}_{-i})$ is non-decreasing in b_i for all i, then we say $\mathbf{x}$ is **monotone**.

Once $\mathbf{x}$ is determined, we can apply **Myerson's Payment Formula** from [3]:

$$p_i(b_i, \mathbf{b}_{-i}) = \int_0^{b_i} z \frac{d}{dz}[x_i(z, \mathbf{b}_{-i})]dz.$$

This formula may seem counter-intuitive at first, but after some rewriting and consideration of bidder i's utility under this payment rule, it will become clear that bidding v_i will be the utility-maximizing strategy. First, we separate the payment formula into two terms using integration by parts:

$$p_i(b_i, \mathbf{b}_{-i}) = b_i x_i(b_i, \mathbf{b}_{-i}) - \int_0^{b_i} x_i(z, \mathbf{b}_{-i})dz.$$

Now we use this payment rule in bidder i's utility function:

$$\begin{aligned} u_i(\mathbf{x}, \mathbf{p}, b_i, \mathbf{b}_{-i}) &= v_i x_i(b_i, \mathbf{b}_{-i}) - [b_i x_i(b_i, \mathbf{b}_{-i}) - \int_0^{b_i} x_i(z, \mathbf{b}_{-i})dz] \\ &= (v_i - b_i)x_i(b_i, \mathbf{b}_{-i}) + \int_0^{b_i} x_i(z, \mathbf{b}_{-i})dz. \end{aligned}$$

Suppose bidder i bid some value $a \neq v_i$. Let us consider the difference in utility between this bidding strategy and telling the truth:

$$\begin{aligned} &u_i(\mathbf{x}, \mathbf{p}, a, \mathbf{b}_{-i}) - u_i(\mathbf{x}, \mathbf{p}, v_i, \mathbf{b}_{-i}) \\ &= \left[(v_i - a)x_i(a, \mathbf{b}_{-i}) + \int_0^a x_i(z, \mathbf{b}_{-i})dz\right] - \left[(v_i - v_i)x_i(v_i, \mathbf{b}_{-i}) + \int_0^{v_i} x_i(z, \mathbf{b}_{-i})dz\right] \\ &= \begin{cases} (v_i - a)x_i(a, \mathbf{b}_{-i}) - \int_a^{v_i} x_i(z, \mathbf{b}_{-i})dz & \text{if } a < v_i \\ \int_{v_i}^a x_i(z, \mathbf{b}_{-i})dz - (a - v_i)x_i(a, \mathbf{b}_{-i}) & \text{if } a > v_i \end{cases} \end{aligned}$$

Because $\mathbf{x}$ is monotone, $u_i(\mathbf{x}, \mathbf{p}, a, \mathbf{b}_{-i}) - u_i(\mathbf{x}, \mathbf{p}, v_i, \mathbf{b}_{-i}) \leq 0$ in either case. In other words, if bidder i bids any other value besides v_i, they risk losing utility. Hence, bidder i should bid v_i if they want to maximize utility. Applying this payment formula satisfies IC. Since, bidders won't ever pay more than their valuation under this formula, Myerson's rule also satisfies IR. M can successfully construct a $(\mathbf{x}, \mathbf{p})$ that solves the welfare-maximizing design problem.

Does this scheme work for the revenue-maximizing design problem? At first, it doesn't seem like we can decouple $\mathbf{x}$ and $\mathbf{p}$ because $\mathbf{x}$ is trying to maximize the sum of payments determined by $\mathbf{p}$. However, we can apply a trick via change of variables, changing valuation parameters into **virtual valuations**. For each bidder i let $\psi_i(v_i)$ denote bidder i's virtual valuation, which is obtained by the following from [4]:

$$\psi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}$$

F_i is the cumulative distribution function of the random variable v_i, and f_i is the probability density function of v_i. If we treated ψ_i as the valuation parameter for each bidder i, [4] shows that maximizing the expected virtual valuation of bidder i is equivalent to maximizing the expected price that bidder i pays. Then, by linearity of expectation:

$$\mathbb{E}_{\mathbf{v}}[\sum_{i=1}^{n} p_i(\mathbf{v})] = \sum_{i=1}^{n} \mathbb{E}_{\mathbf{v}}[p_i(\mathbf{v})] = \sum_{i=1}^{n} \mathbb{E}_{\mathbf{v}}[\psi_i(v_i)x_i(v_i, \mathbf{v}_{-i})] = \mathbb{E}_{\mathbf{v}}[\sum_{i=1}^{n} \psi_i(v_i)x_i(v_i, \mathbf{v}_{-i})]$$

This allows us to decouple $\mathbf{x}$ and $\mathbf{p}$ as before. However, there are two noticeable differences:

1. Since we're maximizing $\psi_i(v_i)$ instead of v_i, $\mathbf{x}$ may no longer be monotone. Monotonicity of $\mathbf{x}$ will depend on the cumulative distribution functions F_i. If a cumulative distribution function F_i preserves the fact that $x_i(b_i, \mathbf{b}_{-i})$ is non-decreasing in b_i for all $\mathbf{b}_{-i}$, we say that F_i is **regular**. For this paper, we assume that all cumulative distribution functions are regular, thereby preserving the monotonicity of $\mathbf{x}$.

2. $\psi_i(v_i)$ may be negative. If we allocate to a bidder with a negative virtual value, it will lower the total virtual surplus. To take care of this, we set up a **reserve price** $\psi_i^{-1}(0)$ for each bidder. If a bidder submits a bid smaller than their reserve price, they will not be allocated anything and will be charged with a price of zero.

Once the bidders with negative virtual values are "tossed out" of the auction, M can construct $\mathbf{x}$ as before: optimize the virtual welfare over all remaining bidders while preserving XP. This $\mathbf{x}$ will be monotone, so once it is constructed, M can apply Myerson's payment formula as before, satisfying IR and IC. Hence, M can successfully construct an $(\mathbf{x}, \mathbf{p})$ that solves the revenue-maximizing design problem as well.

4 Changing the Cost Function

We now attempt to solve these design problems with other cost functions h. Once we move away from h being the identity function, total welfare will begin to depend on payments $\mathbf{p}$:

$$\sum_{i=1}^{n}(u_i(\mathbf{x},\mathbf{p},b_i,\mathbf{b}_{-i})+p_i(\mathbf{b})) = \sum_{i=1}^{n}(v_i x_i(b_i,\mathbf{b}_{-i}) - h(p_i(b_i,\mathbf{b}_{-i}))) + p_i(b_i,\mathbf{b}_{-i}))$$
$$\neq \sum_{i=1}^{n}(v_i x_i(b_i,\mathbf{b}_{-i})).$$

We won't be able to decouple $\mathbf{x}$ from $\mathbf{p}$ as before in the welfare-maximizing design problem, so we will focus our efforts on the revenue-maximizing design problem. For a bidder i, let $q_i(b_i,\mathbf{b}_{-i})$ be bidder i's **perceived payment**. The perceived payment for bidder i is the cost of bidder i's utility when bidder i pays price $p_i(b_i,\mathbf{b}_{-i})$. In other words, $q_i(b_i,\mathbf{b}_{-i}) = h(p_i(b_i,\mathbf{b}_{-i})$. If we treat q_i as if they were the actual payments p_i, bidder i's utility function $u_i(\mathbf{x},\mathbf{p},b_i,\mathbf{b}_{-i}) = v_i x_i(b_i,\mathbf{b}_{-i}) - q_i(b_i,\mathbf{b}_{-i})$ looks exactly like one where we assumed that the cost function was the identity function. That means we can apply the same virtual valuation trick as before, giving us:

$$\mathbb{E}_{\mathbf{v}}[\psi_i(v_i)x_i(v_i,\mathbf{v}_{-i})] = \mathbb{E}_{\mathbf{v}}[q_i(\mathbf{v})] = \mathbb{E}_{\mathbf{v}}[h(p_i(\mathbf{v}))]$$

The result looks hopeful, but maximizing $\mathbb{E}_{\mathbf{v}}[\sum_{i=1}^{n} h(p_i(\mathbf{v}))]$ will not be the same as maximizing $\mathbb{E}_{\mathbf{v}}[\sum_{i=1}^{n} p_i(\mathbf{v})]$ in general. For particular cost functions h, this problem can be solved or bounded with a heuristic. We first consider the case where h is a general linear function. For a fixed $m > 0$, let $h(p_i(b_i,\mathbf{b}_{-i})) = mp_i(b_i,\mathbf{b}_{-i}) + a$, where a is a constant determined by M. In this example, m represents the price elasticity of the bidders and a represents the "admission price" to enter the auction. Linearity of h implies that its inverse h^{-1} is also linear. Because h^{-1} is linear, we can use linearity of expectation to show that the expected total inverse cost of virtual values is equal to the expected total revenue:

$$\mathbb{E}_{\mathbf{v}}[\sum_{i=1}^{n} h^{-1}(\psi_i(v_i)x_i(v_i,\mathbf{v}_{-i}))] = \sum_{i=1}^{n}\mathbb{E}_{\mathbf{v}}[h^{-1}(\psi_i(v_i)x_i(v_i,\mathbf{v}_{-i}))]$$
$$= \sum_{i=1}^{n} h^{-1}(\mathbb{E}_{\mathbf{v}}[\psi_i(v_i)x_i(v_i,\mathbf{v}_{-i})]) = \sum_{i=1}^{n} h^{-1}(\mathbb{E}_{\mathbf{v}}[h(p_i(\mathbf{v}))])$$
$$= \sum_{i=1}^{n}\mathbb{E}_{\mathbf{v}}[h^{-1}(h(p_i(\mathbf{v})))] = \sum_{i=1}^{n}\mathbb{E}_{\mathbf{v}}[p_i(\mathbf{v})] = \mathbb{E}_{\mathbf{v}}[\sum_{i=1}^{n} p_i(\mathbf{v})]$$

Hence, if M chooses $\mathbf{x}$ that maximizes $\mathbb{E}_{\mathbf{v}}[\sum_{i=1}^{n} h^{-1}(\psi_i(v_i)x_i(v_i,\mathbf{v}_{-i}))]$, $\mathbf{x}$ will maximize expected total revenue. This process will be very similar to maximizing $\mathbb{E}_{\mathbf{v}}[\sum_{i=1}^{n}\psi_i(v_i)x_i(v_i,\mathbf{v}_{-i})]$, only this time the reserve price is set to $\psi_i^{-1}(a)$ to ensure that each $h^{-1}(\psi_i(v_i)x_i(v_i,\mathbf{v}_{-i}))$ is non-negative. Since $m > 0$, $\frac{1}{m} > 0$ and h^{-1} is increasing. Maintaining the assumption that all distributions F_i are regular, $\mathbf{x}$ will remain monotone. Therefore, we can apply Myerson's Payment formula as before, satisfying the IC constraint.

For $a > 0$, however, this mechanism does not satisfy IR; some bidders will enter the auction, pay the admission price, and go home with nothing. This will yield $-a$ utility, incentivizing those bidders to not participate in the auction. If the seller wants to charge an admission fee, we can only build a mechanism that satisfies Bayesian Individual Rationality (BIR). Luckily, we can do that by "reverse-engineering" the value a:

1. Have M solve the auction by constructing $(\mathbf{x}, \mathbf{p})$ for with $a = 0$. This will maximize expected revenue while preserving XP and IC.

2. Given $(\mathbf{x}, \mathbf{p})$, calculate each bidder's expected utility assuming $a = 0$.

3. Set a equal to the minimum expected utility found. Now, $a > 0$ while preserving BIR. Changing the value of a does not change the allocation or payment rule.

Sadly, our journey of completely solving revenue-maximizing design problems ends once we assume a more general cost function h. However, [1] provides a heuristic lower bound on maximum total revenue. For the remainder of this paper, we assume that h is non-decreasing, invertible, and $h(0) = 0$. To demonstrate this lower bound, we assume the following heuristic:

$$\forall i \in \{1, 2, ..., n\}, \forall \mathbf{v} \in \mathbf{T}, \;\; h(p_i(v_i, \mathbf{v}_{-i})) \geq \psi_i(v_i) x_i(v_i, \mathbf{v}_{-i})$$

That means $p_i(v_i, \mathbf{v}_{-i}) \geq h^{-1}(\psi_i(v_i) x_i(v_i, \mathbf{v}_{-i}))$. Taking an expected value over the sum will preserve the same relationship:

$$\mathbb{E}_{\mathbf{v}} \left[\sum_{i=1}^{n} p_i(\mathbf{v}) \right] \geq \mathbb{E}_{\mathbf{v}} \left[\sum_{i=1}^{n} h^{-1}(\psi_i(v_i) x_i(v_i, \mathbf{v}_{-i})) \right]$$

Constructing an auction that yields this revenue bound will be equivalent to solving the revenue-maximizing design problem for linear h:

1. Assume that the input $\mathbf{b} = \mathbf{v}$, change valuation parameters v_i into virtual valuations $\psi_i(v_i)$. Remove bidders that have negative virtual values.

2. Of the remaining bidders, construct an allocation rule $\mathbf{x}$ that maximizes $\mathbb{E}_{\mathbf{v}}[\sum_{i=1}^{n} h^{-1}(\psi_i(v_i) x_i(v_i, \mathbf{v}_{-i}))]$ such that XP is satisfied.

3. Given $\mathbf{x}$, apply Myerson's payment formula to generate $\mathbf{p}$, satisfying IR and IC.

5 Applications

We first look at the **k-Vickrey auction**. Consider an environment with n bidders, each wanting exactly one out of k identical goods. In this case, $G = \{0, 1\}$ and $g = k$. The mechanism is described as follows:

$M =$ "On input $\mathbf{v}$:

1. Transform each valuation parameter v_i into a virtual value $\psi_i(v_i)$. Remove all bidders that have a negative virtual value. Let m be the number of remaining bidders. For each removed bidder i, set $x_i(\mathbf{v}) = 0$ and $p_i(\mathbf{v}) = 0$.

2. Sort and re-index the remaining bidders such that $v_1 \geq v_2 \geq ... \geq v_m$.

3. Let $l = \min\{m, k\}$. For each bidder i, set $x_i(\mathbf{v}) = 1$ if $1 \leq i \leq l$ and $x_i(\mathbf{v}) = 0$ otherwise.

4. For each bidder i, set $p_i(\mathbf{v}) = \max\{\psi_i^{-1}(0), v_{l+1}\}$ if $1 \leq i \leq l$ and $p_i(\mathbf{v}) = 0$ otherwise.

5. Return $(\mathbf{x}, \mathbf{p})$."

M selected the l bidders with the highest valuation parameters. Regular distributions will imply that the ordering is preserved when converting valuation parameters into virtual values, and the inverse cost function h^{-1} is non-decreasing. Hence, selecting the l highest bids will also select the l highest $h^{-1}(\psi_i(v_i))$ quantities. The payment function found in step four is the application of Myerson's payment formula. Both of these tricks will be used again in the next two auctions.

Our next mechanism environment is the **sponsored search auction**. Suppose each bidder was a website owner who wanted their site to be featured on the front page of a seller's search engine. Suppose there are k locations where each bidder's site could be advertised on the page, and locations are valued based on their probability that a user clicks on that location when entering the site. For a location j, let α_j denote the probability that j is clicked. For simplicity, assume that $\alpha_1 \geq \alpha_2 \geq ... \geq \alpha_k$. Each bidder may have a different valuation for one click, but the relative values of these locations are constant among all bidders. Hence, the set of goods is homogeneous with $G = \{0, \alpha_1, \alpha_2, ..., \alpha_k\}$ and $g = 1$. The mechanism is described as follows:

$M =$ "On input $\mathbf{v}$:

1. Transform each valuation parameter v_i into a virtual value $\psi_i(v_i)$. Remove all bidders that have a negative virtual value. Let m be the number of remaining bidders. For each removed bidder i, set $x_i(\mathbf{v}) = 0$ and $p_i(\mathbf{v}) = 0$.

2. Sort and re-index the remaining bidders such that $v_1 \geq v_2 \geq ... \geq v_m$.

3. Let $l = \min\{m, k\}$ and set $\alpha_{l+1} = 0$. For each bidder i, set $x_i(\mathbf{v}) = \alpha_i$ if $1 \leq i \leq l$ and $x_i(\mathbf{v}) = 0$ otherwise.

4. For each bidder i, set $p_i(\mathbf{v}) = 0$ if $i \geq l$. For $1 \leq i \leq l$, set payment to be:

$$p_i(\mathbf{v}) = \psi_i^{-1}(0) x_i(\psi_i^{-1}(0), \mathbf{v}_{-i}) + \sum_{z=i}^{l} \left[v_{z+1}(\alpha_z - \alpha_{z+1}) \mathbb{1}_{\{v_{z+1} \geq \psi_i^{-1}(0)\}} \right]$$

5. Return $(\mathbf{x}, \mathbf{p})$."

Finally, we consider the **knapsack auction**. Suppose each bidder was a television advertiser wanting to feature their commercial on the seller's channel during a specific time frame. Suppose the time frame was of length one. The homogeneous good in this case is time, so $G = [0, 1]$ and $g = 1$. In this case, each bidder has two parameters: a valuation parameter v_i corresponding to their valuation of one unit of time and a **public weight parameter** w_i corresponding to the length of their commercial. For simplicity, assume that M can only allocate either 0 or w_i units of time to bidder M. Let W_i be the amount of time given away after allocating to i bidders. For M to construct an $\mathbf{x}$ to maximize revenue subject to XP will be equivalent to solving the knapsack problem. Here are two possible mechanisms proposed by [2] that approximate the solution:

M_1 = "On input $\mathbf{v}$:

1. Transform each valuation parameter v_i into a virtual value $\psi_i(v_i)$. Remove all bidders that have a negative virtual value. Let m be the number of remaining bidders. For each removed bidder i, set $x_i(\mathbf{v}) = 0$ and $p_i(\mathbf{v}) = 0$.
2. Sort and re-index the remaining bidders such that $v_1 \geq v_2 \geq ... \geq v_m$.
3. Set $W_0 = 0$.
4. For $i = 1, 2, ..., m$:
   - If $W_{i-1} + w_i \leq 1$, set $x_i(\mathbf{v}) = w_i$ and set $W_i = W_{i-1} + w_i$.
   - Else, set $x_i(\mathbf{v}) = 0$ and set $W_i = W_{i-1}$.
5. Holding $\mathbf{v}_{-i}$ fixed, let θ_i be the smallest possible valuation parameter bidder i can have such that $x_i(\theta_i, \mathbf{v}_{-i}) = x_i(v_i, \mathbf{v}_{-i})$. For each i, set $p_i(\mathbf{v}) = \theta_i x_i(v_i, \mathbf{v}_{-i})$.
6. Return $(\mathbf{x}, \mathbf{p})$."

The second mechanism M_2 looks almost identical to M_1. The difference is that the underlying optimization algorithm is greedy in $v_i w_i$ as opposed to v_i:

M_2 = "On input $\mathbf{v}$:

1. Transform each valuation parameter v_i into a virtual value $\psi_i(v_i)$. Remove all bidders that have a negative virtual value. Let m be the number of remaining bidders. For each removed bidder i, set $x_i(\mathbf{v}) = 0$ and $p_i(\mathbf{v}) = 0$.
2. Sort and re-index the remaining bidders such that $v_1 w_1 \geq v_2 w_2 \geq ... \geq v_m w_m$.
3. Set $W_0 = 0$.
4. For $i = 1, 2, ..., m$:
   - If $W_{i-1} + w_i \leq 1$, set $x_i(\mathbf{v}) = w_i$ and set $W_i = W_{i-1} + w_i$.
   - Else, set $x_i(\mathbf{v}) = 0$ and set $W_i = W_{i-1}$.

5. Holding $\mathbf{v}_{-i}$ fixed, let θ_i be the smallest possible valuation parameter bidder i can have such that $x_i(\theta_i, \mathbf{v}_{-i}) = x_i(v_i, \mathbf{v}_{-i})$. For each i, set $p_i(\mathbf{v}) = \theta_i x_i(v_i, \mathbf{v}_{-i})$.

6. Return $(\mathbf{x}, \mathbf{p})$."

Our mechanism M will first run both M_1 and M_2, and then M will choose $(\mathbf{x}, \mathbf{p})$ from the sub-mechanism that yielded more revenue. [2] shows that M 2-approximates M''s objective function, and this will suffice for our heuristic lower bound.

6 Conclusion

The design problems in this paper, while intricate at first, can be divided into manageable subproblems and conquered accordingly. The power of Myerson's payment formula reduced our welfare-maximizing design problems into optimization problems, and the use of virtual values allowed us to apply the same reduction to our revenue-maximizing auctions. For a more general cost function, design problems could not be divvied up so easily, but a heuristic lower bound could be derived for revenue-maximizing design problems using previous decoupling methods. With the help of these techniques, we were able to construct mechanisms for k-Vickrey, sponsored search, and knapsack auctions. There are many other mechanism environments that relax the assumptions made in this paper, and the pursuit to divide those design problems into subproblems continues.

References

[1] Amy Greenwald, Takehiro Oyakawa, and Vasilis Syrgkanis. Optimal auctions with convex perceived payments. *CoRR*, abs/1601.07163, 2016. URL `http://arxiv.org/abs/1601.07163`.

[2] Ahuva Mu'Alem and Noam Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. *Games and Economic Behavior*, 64(2):612–631, 2008.

[3] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1): 58–73, 1981.

[4] Tim Roughgarden. Lectures notes on algorithmic game theory. `http://theory.stanford.edu/ tim/f13/f13.pdf`, Accessed: 2015-09-13, 2013.