Homework 3: Myerson's Lemma

CSCI1440/2440

2025-09-30

Due Date: Tuesday, September 30, 2025. 11:59 PM.

We encourage you to work in groups of size two. Each group need only submit one solution. Your submission must be typeset using LaTeX. Please submit via Gradescope with you and your partner's Banner ID's and which course you are taking.

For 1000-level credit, you need only solve the first three problems. For 2000-level credit, you should solve all four problems.

---

1 *All-Pay Auction Equilibrium*

In an **all-pay auction**, the good is awarded to the highest bidder, but rather than only the winner paying, *all* bidders i must pay their bid: i.e., $u_i = v_i x_i - p_i$.

Using the envelope theorem, derive (necessary conditions on) the symmetric equilibrium of an all-pay auction in which all the bidders' values are drawn i.i.d. from the same bounded distribution F.

2 *The Revelation Principle and Revenue Equivalence*

Rather than insisting that incentive compatibility and individual rationality hold *always*, suppose we relax these requirements and ask only that these properties hold *in expectation*.

Define the **interim allocation** and **interim payment** functions, respectively, as follows:

$$\hat{x}_i(v_i) = \mathop{\mathbb{E}}_{\mathbf{v}_{-i} \sim F_{-i}} \left[x_i(v_i, \mathbf{v}_{-i}) \right], \quad \forall i \in [n], \forall v_i \in T_i, \tag{1}$$

$$\hat{p}_i(v_i) = \mathop{\mathbb{E}}_{\mathbf{v}_{-i} \sim F_{-i}} \left[p_i(v_i, \mathbf{v}_{-i}) \right], \quad \forall i \in [n], \forall v_i \in T_i. \tag{2}$$

Further, define **Bayesian incentive compatibility** (BIC) to mean that bidding truthfully is, in expectation, utility maximizing:

$$v_i \hat{x}_i(v_i) - \hat{p}_i(v_i) \geq v_i \hat{x}_i(t_i) - \hat{p}_i(t_i), \quad \forall i \in [n], \forall v_i, t_i \in T_i. \tag{3}$$

Likewise, define **interim individual rationaltiy** (IIR) to mean that bidding truthfully, in expectation, leads to non-negative utility:

$$v_i \hat{x}_i(v_i) - \hat{p}_i(v_i) \geq 0, \quad \forall i \in [n], \forall v_i, t_i \in T_i. \tag{4}$$

Myerson's lemma also holds in the interim case, so a mechanism satisfies BIC and IIR iff

1. Interim allocations are monotone non-decreasing:

$$\hat{x}_i(v_i) \geq \hat{x}_i(t_i), \quad \forall i \in [n], \forall v_i \geq t_i \in T. \tag{5}$$

2. Payments take the following form:

$$\hat{p}_i(v_i) = v_i \hat{x}_i(v_i) - \int_0^{v_i} \hat{x}_i(z)\, \mathrm{d}z, \quad \forall i \in [n], \forall v_i \geq t_i \in T. \tag{6}$$

Let's design a BIC auction for a single-parameter environment with n bidders, each of whom draws her values from a uniform distribution on $[0, 1]$.

1. Calculate the interim allocation function $\hat{x}_i(v_i)$. Show your work.
2. Calculate the interim payment formula $\hat{p}_i(v_i)$. Show your work.
3. Recall the symmetric equilibrium strategy in a first-price auction, namely $(n-1/n)v$. Apply the revelation principle to the first-price auction. Interpret the result.
4. Repeat this same exercise for third-price and/or all-pay auction(s). Interpret the result(s).
5. Based on your observations, give a high-level "proof" of the revenue equivalence theorem for this auction design setting.

3 *Sponsored Search Extension*

In this problem, we generalize our model of sponsored search to include an additional *quality* parameter $\beta_i > 0$ that characterizes each bidder i. With this additional parameter, we can view α_j as the probability a user views an ad, and β_i as the conditional probability that a user then clicks, given that she is already viewing the ad. Note that α_j, the view probability, depends only on the slot j, not on the advertiser occupying that slot, while β_i, the conditional click probablity, explicitly depends on the advertiser i.

In this model, given bids $\mathbf{v}$, bidder i's utility is given by:

$$u_i(\mathbf{v}) = \beta_i v_i x(\mathbf{v}) - p(\mathbf{v})$$

So if bidder i is allocated slot j, her utility is:

$$u_i(\mathbf{v}) = \beta_i v_i \alpha_j - p(\mathbf{v})$$

Like click probabilities, you should assume qualities are public, not private, information.

1. Define total welfare for this model of sponsored search, and then describe an allocation rule that maximizes total welfare, given the bidders' reports. Justify your answer.
2. Argue that your allocation rule is monotonic, and use Myerson's characterization lemma to produce a payment rule that yields a DSIC mechanism for this sponsored search setting.

4 *The Knapsack Auction*

The *knapsack problem* is a famous NP-hard[1] problem in combinatorial optimization. The problem can be stated as follows:

[1] There are no known polynomial-time solutions.

> There is a knapsack, which can hold a maximum weight of $W \geq 0$. There are n items; each item i has weight $w_i \leq W$ and value $v_i \geq 0$. The goal is to find a subset of items of maximal total value with total weight no more than W.

Written as an integer linear program,

$$\max_x \sum_{i=1}^{n} x_i v_i$$

subject to

$$\sum_{i=1}^{n} x_i w_i \leq W$$
$$x_i \in \{0,1\}, \quad \forall i \in [n]$$

The key difference between optimization and mechanism design problems is that in mechanism design problems the constants (e.g., v_i and w_i) are not assumed to be known to the center / optimizer; on the contrary, they must be elicted, after which the optimization problem can then be solved as usual.

With this understanding in mind, we can frame the knapsack problem as a mechanism design problem as follows. Each bidder has an item that she would like to put in the knapsack. Each item is characterized by two parameters—a public weight w_i and a private value v_i. An auction takes place, in which bidders report their values. The auctioneer then puts some of the items in the knapsack, and the bidders whose items are selected pay for this privilege. One real-world application of a knapsack auction is the selling of commercial snippets in a 5-minute ad break (e.g., during the Superbowl).[2]

[2] Here, the weight of a commercial is its time in seconds.

Since the problem is NP-hard, we are unlikely to find a polynomial-time welfare-maximizing solution. Instead, we will produce a polynomial-time, DSIC mechanism that is a 2-approximation of the optimal welfare. In particular, for any set possible set of values and weights, we aim to always achieve at least 50% of the optimal welfare.

We propose the following greedy allocation scheme: Sort the bidders' items in decreasing order by their ratios v_i/w_i, and then allocate items in that order until there is no room left in the knapsack.

1. Show that the greedy allocation scheme is not a 2-approximation by producing a counterexample where it fails to achieve 50% of the optimal welfare.

Alice proposes a small improvement to the greedy allocation scheme. Her improved allocation scheme compares the welfare achieved by the greedy allocation scheme to the welfare achieved by simply putting the single item of highest value into the knapsack.[3] She then uses whichever of the two approaches achieves greater welfare. It can be shown that this scheme yields a 2-approximation of optimal welfare. We will use it to create a mechanism that satisfies individual rationality and incentive compatibility.

[3] Note that weakly greater welfare could be achieved by greedily filling the knapsack with items in decreasing order of value until no more items fit. We do not consider this scheme, because it is unnecessary to achieve a 2-approximation; however, it is an obvious heuristic that anyone solving this problem in the real world would of course implement.

2. Argue that Alice's allocation scheme is monotone.

3. Now use Myerson's payment formula to produce payments such that the resulting mechanism is DSIC and IR.