ORIE 4154 **Homework 5: Due May 10th, 11.59pm (on CMS)**

Spring 2017 Sid Banerjee (sbanerjee@cornell.edu)

Problem 1: Choice models and assortment optimization

Part (a)

Consider a MNL choice model over five products with prices $(p1, \ldots, p5) = (7, 6, 4, 3, 2)$ and preference weights (i.e., MNL parameters) $(v1, \ldots, v5) = (3, 5, 6, 4, 5)$. The preference weight of the no-purchase alternative is $v_0 = 10$. Compute the optimal expected-revenue assortment.

Part (b)

Next, consider a mixed-MNL choice model, wherein we have two consumer types and three products. The probability of observing each consumer type is $(\alpha_1, \alpha_2) = (0.5, 0.5)$. The product prices are $(p1, p2, p3) = (8, 4, 3)$. A consumer of type 1 has preference weights $(v1, v21, v31) = (5, 20, 0)$, and a consumer of type 2 has preference weights $(v12, v2, v32) = (1/5, 10, 10)$; the preference weight of the no-purchase alternative is 1 for both types.

i. First, find the optimal assortments $S_1^*, , S_2^*$ for each individual type, and compute the expected revenue of S_1^* and S_2^* for the mixed-MNL model.

ii. Next, consider the assortment $\{1, 3\}$, and show that this achieves a higher revenue under the mixed-MNL model than the two assortments in the previous part. (In fact, $\{1, 3\}$ is the optimal assortment, but you do not need to show that).

Note: To maximize the expected revenue over the two consumer types, we need to offer product 3; however, this product is not offered when we want to maximize the expected revenue from either one of the consumer types individually.

Problem 2: Pivot rules and procurement auctions

Consider a single-item auction setting, where each bidder i has a private value v_i. In class, we studied the Vickrey (or second-price) auction for such settings, and saw that it has two properties: *i. incentive compatibility* (DSIC), i.e., $u_i(b_i, \mathbf{b}_{-i}) \leq u_i(v_i, \mathbf{b}_{-i})$, and *ii. individual rationality*, i.e., $u_i(v_i, \mathbf{b}_{-i}) \geq 0$. We now will see a more general mechanism that has the DSIC property.

Given bids $\mathbf{b}$, let $i^* = \arg\max_i\{b_i\}$, and consider the mechanism $(\mathbf{x}, \mathbf{p})$, with allocation rule $x_i = \mathbb{1}_{\{i=i^*\}}$ (i.e., award item to highest bidder), and payment rule $p_j = 0$ for all $j \neq i^*$, and $p_{i^*} = f_{i^*}(\mathbf{b}_{-i^*})$, where $f_i(\mathbf{b}_{-i^*})$ is some function which only depends on *the bids of other bidders* (and any other publicly-known constant). The term $f_{i^*}(\mathbf{b}_{-i^*})$ is sometimes referred to as a *pivot rule.*

Part (a)

Argue that $f_i(\mathbf{b}_{-i}) = \max_{j \neq i}\{b_j\}$ always ensures individual rationality in any single-item auction setting. This gives us the Vickrey auction (and more generally, the pivot rule is a special case of the so-called *Clark pivot rule*).

Part (b)

Find the *maximum* pivot rule $f_i(\mathbf{b}_{-i})$ that ensures individual rationality in the following settings:

i. Bidder i's value v_i is known to satisfy $v_i \in [v_{\min}(i), v_{\max}(i)]$, i.e., the maximum and minimum values of each bidder's valuations is *public* knowledge.

ii. There is a publicly-known constant δ such that any two bidders i and j, we have $|v_i - v_j| \geq \delta$ (i.e., any two bidders' values are at least δ apart).

Part (c)

Argue that the pivot rule is dominant-strategy incentive compatible for the following choices of f

i. The function from (b) part i.

ii. The function from (b) part ii.

iii. $f_i(\mathbf{b}_{-i}) = \max_{j \neq i}\{b_j\}$

Part (d)

Finally, consider a *procurement auction*, where you want to buy an item (or enter into a contract for some work) from among a group of n sellers. Each seller has a private cost c_i which is known to lie in a (publicly-known) range $[c_{\min}(i), c_{\max}(i)]$. Consider a mechanism that collects bids from the sellers, chooses the seller i^* with the lowest bid, and then pays i^* an amount $p_{i^*}(b_{i^*}, \mathbf{b}_{-i})$. What is the minimum payment that you can offer such that the mechanism is DSIC and IR.
Hint: One way to view a cost is as a negative value $v_i < 0$ for entering into a contract, and a payment as a negative price $p_i \leq 0$. As in a single-item auction, each seller i has utility $u_i = v_i - p_i$, where $u_i = 0$ for no-trade; the seller's aim is to bid so as to maximize utility $u_i(b_i, \mathbf{b}_{-i})$.

Problem 3: Welfare maximization and externality pricing

Part (a)

Consider an arbitrary single-parameter environment, with feasible set $\mathcal{X}$. Given values v_i, the welfare-maximizing allocation rule is $\mathbf{x}(\mathbf{v}) = \arg\max_{(x_1,\ldots,x_n)\in\mathcal{X}} \sum_{i=1}^n v_i x_i$. Prove that this allocation rule is monotone. You can assume for convenience that all values are distinct (or more generally, that ties are broken in some deterministic and consistent way, for example, lexicographically.)

Part (b)

Next, consider feasible sets $\mathcal{X}$ that contain only $0-1$ vectors, i.e., each bidder either wins or loses. Now, given any monotone allocation rule $\mathbf{x}(\mathbf{b})$, for any bidder i and other bids $\mathbf{b}_{-i}$, argue that the Myerson payment rule can be written as:

$$p(b_i, \mathbf{b}_{-i}) = \begin{cases} 0 & \text{if } x_i(b_i, \mathbf{b}_{-i}) = 0 \\ b_i^*(\mathbf{b}_{-i}) & \text{if } x_i(b_i, \mathbf{b}_{-i}) = 1 \end{cases}$$

where $b_i^*(\mathbf{b}_{-i})$ is bidder i's *critical bid*, i.e., the lowest bid at which i gets a non-0 allocation.

Part (c)

For feasible sets $\mathcal{X}$ containing only $0-1$ vectors, we can identify each feasible allocation with a 'winning set' of bidders. Assume further that the environment is 'downward closed', meaning that subsets of a feasible set are again feasible. Prove that, when $\mathcal{S}^*$ is the set of winning bidders and $i \in \mathcal{S}^*$, then is critical bid equals the difference between (i) the maximum surplus of a feasible set that excludes i (you should assume there is at least one such set); and (ii) the surplus $\sum_{j\in\mathcal{S}^*\setminus\{i\}} v_j$ of the bidders other than i in the chosen outcome $\mathcal{S}^*$. Also, is this difference always nonnegative? *Note: In other words, a winning bidder pays its 'externality' – the welfare loss it imposes on others.*

Part (d)

To see how to use the above result, consider the knapsack auction we discussed in class (for allocating TV advertisements to ad slots). We want to choose ads to fill a slot of length at most 120 seconds. Bidders have private values v_i and public ad-lengths ℓ_i. In particular, consider a setting with 4 bidders $\{a, b, c, d\}$, with private values (v_a, v_b, v_c, v_d) and lengths $(60s, 40s, 40s, 40s)$. Now suppose all bidders truthfully report their bids, and the auctioneer finds that the optimal allocation is to choose ads b, c and d (note that this is feasible). Use part c to compute the Myerson payments for all the bidders.

Problem 4: The Myerson optimal-revenue auction

In this question, we will get some practice with the Myerson optimal-revenue auction.

Part (a)

Consider an auction with k identical goods, with at most one given to each bidder. There are n bidders whose valuations are i.i.d. draws from a regular distribution F. Describe the optimal auction in this case.

Part (b)

Next, consider a single-item auction with independent but non-identical values; in particular, assume bidder is valuation is drawn from its own regular distribution F_i.

i. Give a formula for the winners payment in the optimal revenue auction, in terms of the bidders virtual valuation functions.

ii. Show by example that, in an optimal auction, the highest bidder need not win, even if he has a positive virtual valuation.

Hint: For the last part, a simple setting with two bidders with valuations from different uniform distributions suffices.