Problem 1: Random variables, common distributions and the monopoly price

In this problem, we will revise some basic concepts in probability, and use these to better understand the *monopoly price* (alternatively referred to as the optimal posted price or Myerson price). Recall from class that we want to study a setting where we want to sell a single item to a single buyer. The buyer has a random *reservation value* $V \geq 0$ for the item, drawn from a distribution with CDF $F(\cdot)$ (we denote this as $V \sim F$). If we charge a price p, then the buyer gets a *utility* $U = V - p$ from buying the item. We assume that the buyer purchases the item *if and only if* $U \geq 0$.

Part (a)

Fix the price of the item to be p, and let X be an *indicator random variable* for the sale (i.e., $X = 1$ if the buyer purchases the item, else 0); what is the probability distribution of X? Also, let $R(p)$ be the revenue we obtain from the sale; what is the expected value and variance of $R(p)$?

Part (b)

Suppose we have n items and n buyers, and offer the first item to the first buyer at price p, the second to the second buyer at price $2p$, and in general, offer the k^{th} item to the k^{th} buyer at price $k \cdot p$. Let R_k be the revenue obtained from the k^{th} item, and $R = \sum_{k=1}^{n} R_k$ be the net revenue.

Assuming each buyer i has an i.i.d reservation value $V_i \sim F(\cdot)$; what is $\mathbb{E}[R]$ and $\mathrm{Var}(R)$?

Part (c)

Next, assume all buyers have the *same* reservation value V drawn from a UNIFORM$[0, (n+1)p]$ distribution. Now what is the expected value and variance of R?

Part (d)

Returning to the one item/one buyer setting, let $R(p)$ be the revenue we obtain if the posted price is p. Find the optimal price $p^* = \arg\max_{p \geq 0} R(p)$, and check if the function $R(p)$ concave, when:

1. V is uniformly distributed in $[0, m]$.
2. V is distributed as EXPONENTIAL(λ).

Part (e)

Let q denote the probability that we make a sale; we define the *inverse demand function* $p(q)$ to be the maximum price p at which the sale probability is q [1]. For a general (continuous) CDF $F(\cdot)$, write an expression for the inverse demand function in terms of F and q.

[1] We should be careful here, as such a p may not exist, for example, if the distribution is discrete; more generally, we can define $p(q)$ to be the maximum price such that the sale probability is at least q, i.e., $p(q) = \max_{p \geq 0}\{(1 - F(p)) \geq q\}$. For continuous distributions, however, the above definition is fine.

Part (f)

Next, note that we can write the revenue as a function of the sale probability q as $R(q) = q \cdot p(q)$. Write down $R(q)$ for the two distributions in part (d), and show that the function is concave in both cases.

Part (g)

For a general CDF F, show that

$$dR(q)/dq = p(q) - \frac{1 - F(p)}{f(p)}$$

Thus, conclude that the revenue curve $R(q)$ is concave if and only if $p - \frac{1-F(p)}{f(p)}$ is non-decreasing. Such a distribution is said to be *regular*.

Part (h)

For a distribution with CDF F, the *hazard-rate* $\rho(p)$ is defined as

$$\rho(p) = \frac{f(p)}{1 - F(p)}$$

Argue that if a distribution has non-decreasing hazard rate, then its revenue curve $R(q)$ is concave. Such a distribution is said to be a *monotone hazard-rate* (or MHR) distribution.

Note: To see why $\rho(\cdot)$ is called the hazard-rate (and also, to remember the definition), consider F to be the CDF corresponding to the lifetime of a lightbulb before it fuses – then $\rho(t)dt$ is then the probability it will fuse in time $[t, t + dt]$ given that it has survived till time t.

Part (i)

(OPTIONAL) Give an example of a regular distribution that is not MHR.

Hint: Do not think of well-known distributions (these are usually MHR). Instead, recall F can be any non-decreasing bounded continuous function, scaled to lie in $[0, 1]$.

Problem 2: Linear Programming, duality, and the invisible hand of the market

In this problem, we will revise some basic concepts in linear programming, and show how these can be used to demonstrate the power of markets in solving resource allocation problems.

The basic problem we want to consider is that of allocating m items (denoted $\mathcal{I} = \{1, 2, \ldots, m\}$) among n buyers (denoted as $\mathcal{B} = \{1, 2, \ldots, n\}$). One desirable way of doing so is to allocate items to buyers who 'value' them the most. We now show how this notion can be formalized, and how this optimization can be achieved via simple pricing policies. Throughout this problem, we assume that each buyer i has value $v_{ij} \geq 0$ for each item j: let $V = \{v_{ij}\}$ denote the $n \times m$ matrix of buyer values, and assume that all entries of V are distinct. We also assume that each item j has a price p_j, which a buyer must pay to purchase the item. Each item has only one copy, and hence can be sold to either a single buyer, or not sold at all; each buyer can buy as many items as possible.

Part (a)

First, we consider the case of *additive buyers*. We assume that a buyer $i \in \mathcal{B}$ will only buy an item $j \in \mathcal{I}$ if its resulting *utility* $v_{ij} - p_j \geq 0$; moreover, if buyer i purchases a subset of items $\mathcal{I}_i \subseteq \mathcal{I}$, then her net utility is $U_i = \sum_{j \in \mathcal{I}_i} (v_{ij} - p_j)$ [2] . If no item is allocated to buyer i, then $\mathcal{I}_i = \emptyset$ and $U_i = 0$. We define the utility U_s of the seller to be the total amount of money she earns from item sales, and define the *social welfare* $W = U_s + \sum_{i \in \mathcal{B}} U_i$ to be the sum of everyone's utilities.

Let x_{ij} be an indicator that buyer i purchases item j, i.e., $x_{ij} = 1$ if i purchases j, else it is 0. Given buyer valuations V, prices $\{p_j\}$ and indicators $\{x_{ij}\}$, write down the expression for the social welfare. Using this, characterize the allocation of items that maximizes the social welfare.

Part (b)

Now suppose we relax the indicator variables x_{ij} to take values in $[0, 1]$ (in other words, we assume that each item j can be fractionally allocated to a buyer i). Write down a linear program that finds an allocation to maximize the social welfare. Moreover, characterize all the extreme points of the above LP.

Part (c)

Your above LP should have m constraints, one for each item; let π_m be a dual variable associated with each of these constraints. Write down the dual linear program. Moreover, suppose $\{x^*_{ij}\}$ and $\{\pi^*_j\}$ are optimal solutions to the primal and dual programs – write down the complementary slackness conditions.

Part (d)

Given any optimal dual solution π^*, suppose we set the price for each item j as $p_j = \pi^*_j$. Argue that under these prices, there is an allocation of items to agents that obeys the following: (1) if

[2] For example, you visit NYC, and buy entry tickets for multiple museums; (assuming you have enough time) you can now visit them all!

buyer i is allocated item j, then $v_{ij} - p_j \geq 0$, (2) for any item j not allocated to buyer i, we have $v_{ij} - p_j \leq 0$, and (3) the social welfare is maximized.

Part (e)

Next, we consider the case of *unit-demand buyers*. We assume that if buyer $i \in \mathcal{B}$ purchases a subset of items $\mathcal{I}_i \subseteq \mathcal{I}$, then her net utility is $U_i = \max_{j\in\mathcal{I}_i}(v_{ij}) - \sum_{j\in\mathcal{I}_i} p_j$, i.e., she pays for all purchased items, but only gets utility from the highest valued item [3]. As before, the seller's utility U_s is still the total amount of money she earns, and the social welfare is $W = U_s + \sum_{k\in\mathcal{B}} U_i$.

Argue that in any welfare maximizing allocation policy in this setting, each buyer is allocated at most one item. The resulting optimization problem is known as the *maximum-weighted matching* problem (and is a maximization version of the assignment problem that you might have seen in some previous course).

Part (f)

As in part (b), let $x_{ij} \in [0, 1]$ be a fractional allocation of item j to buyer i. Write an LP to choose a fractional allocation in order to maximize welfare.

Note: Recall (or learn...) that the assignment problem LP also has integer corner points – thus, the above relaxation actually gives a valid allocation.

Part (f)

Write down the dual for the above LP, and also write the complementary slackness conditions.

Part (g)

(OPTIONAL) Suppose you are given an optimal primal solution x^*, and associated dual solution μ^*. As in part (d), show that there is a way to use these to set item prices p_j, under which there is an allocation x_{ij} that satisfies: (1) each buyer is allocated at most one item, and each item is allocated to at most one buyer, (2) if buyer i is allocated item j, then $v_{ij} - p_j \geq 0$, (3) for any item j not allocated to buyer i, we have $v_{ij} - p_j \leq 0$, and (4) the social welfare is maximized.

Note: To see why the above result is remarkable, observe that the prices π^* (which are known as Walrasian *or* envy-free *prices) magically coordinate the buyers, such that they each individually pick their favorite item, and by doing so, collectively solve a combinatorial optimization problem!*

[3] You are in NYC again, and buy tickets for multiple broadway shows with identical showtimes (because you were undecided...); you must pay for each ticket, but can only watch one (and are not allowed to sell the other tickets!). This problem should convince you that this is a harder setting, and so you should stick to the museums :)

Problem 3: "You can't always charge what they want (to pay)"

In this question, we will see how the interaction between buyer behavior and seller constraints can sometimes lead to non-intuitive pricing policies, wherein you may not want to sell to everyone who is willing to .

Suppose we own a cabin in the mountains, which we want to rent out on CleanAirBnB. Interested customers arrive sequentially looking to reserve the cabin; they arrive *deterministically* every 2 days, and each customer wants to rent it for a period of 3 nights. Formally, consider discrete time-slots (days) indexed as $t = 0, 1, 2, \ldots$. A single customer arrives at the beginning of every even slot (i.e., at time-slots $0, 2, 4, \ldots$), and if a customer rents the cabin starting at time-slot t, then it becomes free at the beginning of time-slot $t + 3$.

Each customer has value $v = 5$ for staying in the cabin; however, they also have a cost of $c = 1$ per day they need to wait before getting the cabin. For a customer arriving at time t, let d_t denote the number of time slots after which the cabin becomes available, and suppose she is charged a fee of p_t to reserve the cabin from time-slot $t + d_t$ to $t + d_t + 3$: we assume she agrees to reserve the cabin if and only if her *utility* $u_t = v - c \cdot d_t - p_t = 5 - d_t - p_t \geq 0$, else she goes elsewhere looking for cabins to rent. Our aim is to design the pricing policy p_t so as to maximize our revenue.

Part (a)

First, suppose we do not charge customers for reserving the cabin. Assuming $d_0 = 0$ (i.e., the cabin starts as empty), plot how the sequence d_t changes with t.

Hint: You do not need to plot d_t for an infinite number values of t! Observe that whenever d_t becomes equal to some value you have already seen before (i.e., d_s for some $s < t$), then the subsequent evolution is the same as before. Thus, the d_t sequence is periodic in t.

Part (b)

Next, for any customer facing a minimum delay of d days before getting to stay in the cabin, what is the maximum amount $p(d)$ that we can charge her so as to ensure that she makes a reservation?

Part (c)

Suppose we now use the pricing policy $p(d)$ from part (b); what is the resulting sequence d_t (assuming $d_0 = 0$)? Let R_t denote the total earnings up to the start of time slot t (with $R_0 = 0$); compute the *long-term average earning* $R = \lim_{t\to\infty} \frac{R_t}{t}$.

Hint: For any bounded eventually-periodic sequence (i.e., one which is periodic after some initial transient behavior), the long-term average value is the same as the average over any period.

Part (d)

Can you change the pricing policy to get a better long-term average earning?