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 $R(p)$?

Solution: Note that X is a Bernoulli(q) r.v. with parameter $q = \mathbb{P}(X = 1) = \mathbb{P}(V \geq p) = 1 - F(p^-)$ (where $F(p^-) = \lim_{\varepsilon \to 0} F(p - \varepsilon)$ – this is the correct way to write it in case of discrete random variables). Moreover, if $X \sim$Bernoulli(q), we have $\mathbb{E}[X] = q$, $\text{Var}(X) = q(1 - q)$ (you should remember these formulas).

Now for us, the revenue is $R(p) = pX$, hence $\mathbb{E}(R(p)) = p\mathbb{E}(X) = p1 - F(p^-)$; similarly $\text{Var}(R(p)) = p^2\text{Var}(X) = p^2(1 - F(p^-))F(p^-)$.

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 $\text{Var}(R)$?

Solution: Henceforth, let $\pi(p) \triangleq 1 - F(p^-)$. Using part **(a)**, $\mathbb{E}(R_k) = kp\pi(kp)$ and $\text{Var}(R_k) = k^2p^2\pi(pk)(1 - \pi(pk))$. Now we have

$$\mathbb{E}(R) = \sum_k \mathbb{E}(R_k) = p\sum_k k\pi(kp).$$

(By linearity of expectation)

$$\text{Var}(R) = \sum_k \text{Var}(R_k) = p^2\sum_k k^2\pi(pk)(1 - \pi(pk)).$$

(Since V_k, and hence R_k are i.i.d., we can add the variances)

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?

Solution: Now the R_k are correlated, but we can still use linearity of expectation! Note that $\mathbb{E}[R_k]$ is the same as in part **(b)**, but in addition, we now have specified $\pi(kp) = 1 - \frac{kp}{(n+1)p}$. Hence

$$\begin{aligned}\mathbb{E}(R) &= p\sum_{k=1}^{n} k\left(1 - \frac{k}{(n+1)}\right) \\ &= p\left(\frac{n(n+1)}{2} - \frac{1}{n+1}\cdot\frac{n(n+1)(2n+1)}{6}\right) = \frac{pn(n+2)}{6}\end{aligned}$$

To compute the variance, we can not use the sum of individual variances, since they are correlated. In such cases, we have to write out in detail what the random variable R is, and then use $\mathrm{Var}(R) = \mathbb{E}[R^2] - (\mathbb{E}[R])^2$.

To find the pmf for R, observe that if $V \in [kp, (k+1)p)$, then the first k customers purchase the item yielding a revenue of

$$R = \sum_{i=1}^{k} ip = \frac{pk(k+1)}{2}, \qquad \text{when } V \in [kp, (k+1)p).$$

In the case $V \in [0, p)$ no customer buys. Since V is uniform, $\mathbb{P}(V \in [kp, (k+1)p)) = \frac{p}{(n+1)p} = \frac{1}{n+1}$ for each $k \in \{1, 2, \ldots, n\}$. Putting things together, we have that R takes one of the $n+1$ values in $\{0, (p\cdot 1\cdot 2)/2, (p\cdot 2\cdot 3)/2, \ldots, (p\cdot n\cdot(n+1))/2\}$ uniformly (i.e., with probability $1/(n+1)$). Finally,

$$\begin{aligned}\mathrm{Var}(R) &= \mathbb{E}(R^2) - \mathbb{E}(R)^2 \\ &= \sum_{k=1}^{n}\Big(p\frac{k(k+1)}{2}\Big)^2\frac{1}{n+1} - \frac{p^2n^2(n+2)^2}{36} = \frac{p^2}{4(n+1)}\sum_{k=1}^{n}(k^2+k)^2 - \frac{p^2n^2(n+2)^2}{36}.\end{aligned}$$

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).

Solution: From part **(a)**, the expected revenue is $R(p) = p(1 - F(p))$. Also recall that for continuous F, $dF(p)/dp = f(p)$, the pdf of the value distribution.

Irrespective of whether $R(p)$ is concave, the optimal price is either an extreme point (i.e., on the boundary of the domain), or it satisfies

$$R'(p) = 0 \Longleftrightarrow 1 - F(p) - pf(p) = 0.$$

1. In this case $F'(p) = \frac{1}{m}$, so $R'(p) = 1 - \frac{2p}{m}$, hence R is concave. The critical point is $p^* = m/2$; one can check that the extreme points yield revenue zero, so this is the optimal solution.

2. Here $f(p) = \lambda e^{-\lambda p}$, and thus $R'(p) = e^{-\lambda p} - p\lambda e^{-\lambda p} = (1 - \lambda p)e^{-\lambda p}$. The only critical point is $p^* = 1/\lambda$; you can check that this is the only maximum (for example, by plotting the function). However, $R''(p) = (\lambda^2 p - 2\lambda)e^{-\lambda} p$ which is not always negative – hence $R(p)$ is neither concave nor convex.

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.

Solution: We have the equation $q = \mathbb{P}(V \geq p) = 1 - F(p)$, so $p(q) = F^{-1}(1 - q)$.

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.

Solution:

1. The inverse of $F(p) = p/m$ is $F^{-1}(q) = mq$, so $R(q) = qm(1-q)$, which is a concave parabola.
2. The inverse is $F^{-1}(q) = -\frac{1}{\lambda}\ln(1 - q)$, hence $R(q) = -\frac{1}{\lambda} q \ln(q)$. Taking derivatives we get $R''(q) = -1/(\lambda q) \leq 0$; hence this is concave.

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*.

Solution: We know $F(p) = 1 - q$ and hence $f(p)(dp/dq) = -1 \Rightarrow dp/dq = -1/f(p)$. Moreover by the chain rule, we have

$$\frac{d}{dq}R(q) = \frac{d}{dq}(qp(q)) = p(q) + q \cdot dp/dq = p(q) - \frac{q}{f(p)} = p(q) - \frac{1 - F(p(q))}{f(p(q))}.$$

To conclude, recall that differentiable functions are concave iff they have non-increasing derivatives. As $p(q)$ is non-increasing in q, the desired condition follows.

[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 (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.

Solution: Note that, if $\rho(p)$ is non-decreasing, then $-\frac{1}{\rho(p)}$ is non-decreasing and thus $p - \frac{1-F(p)}{f(p)}$ is non-decreasing. This however is the exact condition you derived in part **(g)**.

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 $[0, 1]$.

Solution: Consider the distribution $f(x) = 1/x^2$ for $x \in [1, \infty)$. Note that this is a valid distribution, since $f(x) \geq 0$ and $\int_1^\infty dx/x^2 = 1$; moreover, $F(x) = 1 - 1/x$. Now the hazard rate is given by $\rho(x) = f(x)/(1 - F(x)) = 1/x$ which is decreasing. On the other hand, $x - 1/\rho(x) = 0$ which is non-decreasing – hence the corresponding revenue curve is concave.

(In fact, the revenue curve is given by $R(q) = qF^{-1}(1 - q) = 1$! As a result, this is sometimes called the *equal-revenue* distribution, as we get the same revenue for any price $p \in [1, \infty)$.)

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.

Solution: Note that the seller's utility can be written as $U_s = \sum_{j \in \mathcal{I}} p_j \sum_{i \in \mathcal{B}} x_{ij}$, and the sum of the buyer utilities is $\sum_{i \in \mathcal{B}} \sum_{j \in \mathcal{I}} (v_{ij} - p_j) x_{ij}$. We can add these to get

$$W = \sum_{i \in \mathcal{B}} \sum_{j \in \mathcal{I}} v_{ij} x_{ij}$$

Note that the prices have dropped out of this equation! This is a somewhat remarkable property of welfare - the way to think about it is that the disutility of the purchase price to a buyer is offset by the utility given to the seller – hence it disappears when measuring welfare.

The above expression also tells us that welfare is maximized by allocating *each item j to the buyer i with the highest valuation v_{ij} for that item.* However, once we fix prices, this may not necessarily be a feasible solution (for example, if the price of an item is so high that it is larger than all utilities, then no buyer would want it, even though it may be 'socially better' for the buyer to make the purchase). The rest of the question tries to find prices which make this feasible.

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

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.

Solution: The LP relaxation is given by

$$\begin{aligned} \max_x \quad & \sum_{i\in\mathcal{B}}\sum_{j\in\mathcal{I}} (v_{ij}x_{ij}) & & \\ \text{s.t.} \quad & \sum_i x_{ij} \leq 1 & & \forall\ j \\ & x_{ij} \geq 0 & & \forall\ i,j \end{aligned}$$

The extreme points of the LP will all be integral. On way to see this is to note that the constraints are a subset of the constraints of the *assignment problem*, which you should know from before has integer extreme points.

To see a proof from first principles, suppose that a solution has fractional x_{ij} for some j. Pick any two of the buyers who receive a fraction of the item (Assume WLOG that they are buyers 1 and 2). But it must be that either $v_{1j} > v_{2j}$ or $v_{1j} < v_{2j}$ since the valuations are distinct. Again, suppose WLOG that buyer 1 has a larger valuation. Then we can shift buyer 2's share of the item over to buyer 1, changing our solution by $x_{2j}(v_{1j} - v_{2j}) > 0$. This argument can then be repeated to show that no one receives a fractional allocation.
Note also that we do not need $x_{ij} \leq 1\ \forall\ i,j$ since this is implied by the first constraint.

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.

Solution: The dual LP is

$$\begin{aligned} \min_\pi \quad & \sum_{j\in\mathcal{I}} \pi_j & \\ \text{s.t.} \quad & \pi_j \geq v_{ij} & \forall\ i,j \\ & \pi_j \geq 0 & \forall\ i,j \end{aligned}$$

Given optimal primal and dual solutions (x^*, π^*), the complimentary slackness conditions are

1. $(1 - \sum_{i\in\mathcal{B}} x^*_{ij}) \cdot \pi^*_j = 0\ \forall\ j$
2. $(\pi^*_j - v_{ij}) \cdot x^*_{ij} = 0\ \forall\ i,j$

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 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.

Solution: Set the prices $p_j = \pi^*_j$ and solve the primal LP to get (integer) allocations x^*_{ij}. From the second complementary slackness condition, we see that (1) is true, as x^*_{ij} is non-zero (i.e., i buys item j) iff $v_{ij} \geq p_j$. Moreover, (2) follows from the constraint in the dual, and the fact that the prices π_j are dual feasible. Finally, (3) holds due to the allocation being the optimal solution to the primal LP, whose feasibility in the assignment is guaranteed by (1).

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).

Solution: Again, note that the disutility from item price to the consumers is offset by the utility to the seller. Then the resulting utility is: $\sum_{i \in \mathcal{B}} max_{j \in \mathcal{I}_i} v_{ij}$. Now suppose we are given a welfare maximizing allocation, wherein a buyer i is allocated more than one item. Since she gets value only from one of those items (the one with the highest v_{ij}), hence if we remove the other items from her allocation, her utility remains the same. Hence, it is sufficient to ensure that each buyer receives at most one item.

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.

[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 :)

Solution: From the previous section, we know that the optimal solution will not allocate more than 1 item to each consumer. Therefore, we can add that as a constraint without affecting the optimal solution. The LP relaxation is given by

$$\begin{aligned}
\max_x \quad & \sum_{i\in\mathcal{B}}\sum_{j\in\mathcal{I}} v_{ij}x_{ij} & & \\
\text{s.t.} \quad & \sum_i x_{ij} & \leq 1 \quad & \forall\, j \\
& \sum_j x_{ij} & \leq 1 \quad & \forall\, i \\
& x_{ij} & \geq 0 \quad & \forall\, i,j
\end{aligned}$$

Part (g)

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

Solution: The dual LP will be

$$\begin{aligned}
\min_{\mu,\lambda} \quad & \sum_{i\in\mathcal{B}} \lambda_i + \sum_{j\in\mathcal{I}} \mu_j & & \\
\text{s.t.} \quad & \lambda_i + \mu_j & \geq v_{ij} \quad & \forall\, i,j \\
& \mu_j, \lambda_i & \geq 0 \quad & \forall\, i,j
\end{aligned}$$

with an optimal primal/dual solution $(x^*_{ij}, \mu^*_j, \lambda^*_i)$ complimentary slackness conditions

$$x^*_{ij}\left(\mu^*_j + \lambda^*_i - v_{ij}\right) = 0$$

Part (h)

(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.

Solution: Part (1) follows from the feasibility of $x^*_{ij}, \mu^*_j; r(4)$ follows directly from the optimality of x^*_{ij}. To see the remaining properties, rewrite the complementary slackness condition as $x^*_{ij}\left(\lambda^*_i - (v_{ij} - \mu^*_j)\right) = 0$. Now note that given prices $p_j = \mu^*_j$, an agent i's utility from buying item j is given by $u_i = v_{ij} - \mu^*_j$. Now we will argue that the dual variable λ^*_i can be interpreted to be agent i's utility. To see this, note that for given prices π^*_j, an agent i gets the highest utility from picking the item that has the highest $v_{ij} - \mu^*_j$. Complementary slackness gives us that $\lambda^*_i = v_{ij} - \pi^*_j$ when $x^*_{ij} > 0$, else if $x^*_{ij} = 0$, then $\lambda^*_i \geq v_{ij} - \mu^*_j$.

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

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.

Solution: The plot is given in Figure 1; note that the pattern from $t = 8$ to $t = 14$ repeats.

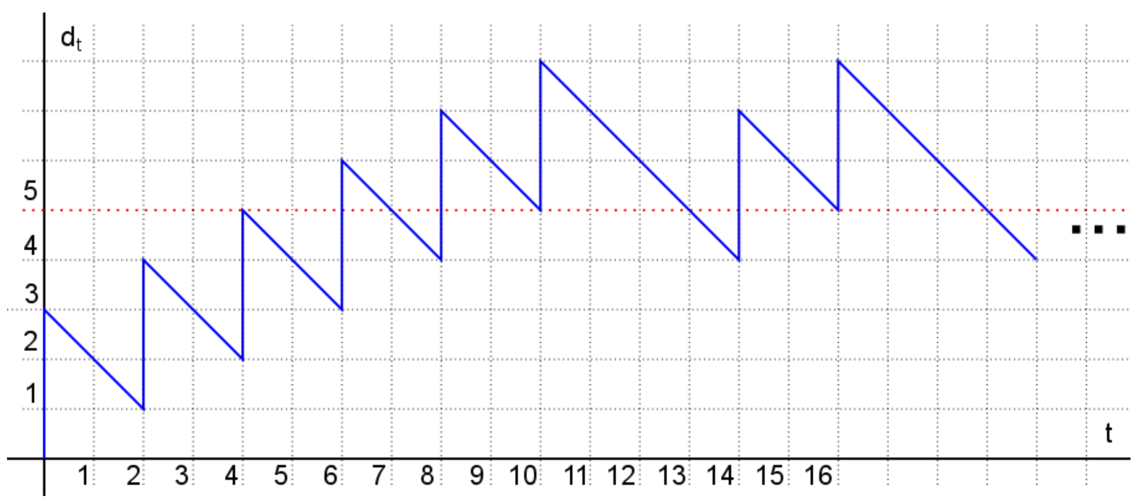

Figure 1: Plotting d_t as a function of t with no prices

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?

Solution: We can charge a customer as much as would ensure that her utility is positive, i.e., $p(d)$ must ensure that $5 - d - p(d) \geq 0$, and hence $p(d) \leq 5 - d$.

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* $R = \lim_{t\to\infty} \frac{R_t}{t}$.

Solution: Since we are charging just enough to ensure that customers rent if $d \leq 5$, we have the same delay sequence as in part (a). In Figure 2, we have annotated the plot with the amount that each reserving customer pays. Note that over each repeating period (i.e., the period between 8 t0 14), we earn only 1 unit – hence the long-term average rate $R = \lim_{t\to\infty} \frac{R_t}{t} = 1/6$.

To see this in more detail, let $T = 8 + 6k$ for some large integer k. Then the total reward earned between $t = 0$ and $t = T$ is $5 + 4 + 3 + 2 + 1 \cdot k = 14 + k$, and thus $\frac{R_T}{T} = \frac{14+k}{8+6k} \to \frac{1}{6}$ as $k \to \infty$.

Part (d)

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

Solution: Consider a policy where we always charge $p = 5$ irrespective of the delay. The delay curve for this is drawn in red in figure 2. Note that now we make $R = \frac{5}{4}$, which is larger than $\frac{1}{6}$!

(In fact, you can do slightly better by charging $p = 5$ when delay is 0, and $p = 4$ for all other delays – try verifying that now you make $R = 9/6 > 5/4$.)

Note: This question was based on a model originally developed by Naor in the 70s. This work has led to a large area of research on pricing in queues – if you are interested (or looking for potential projects), see the book by Hassan and Haviv.

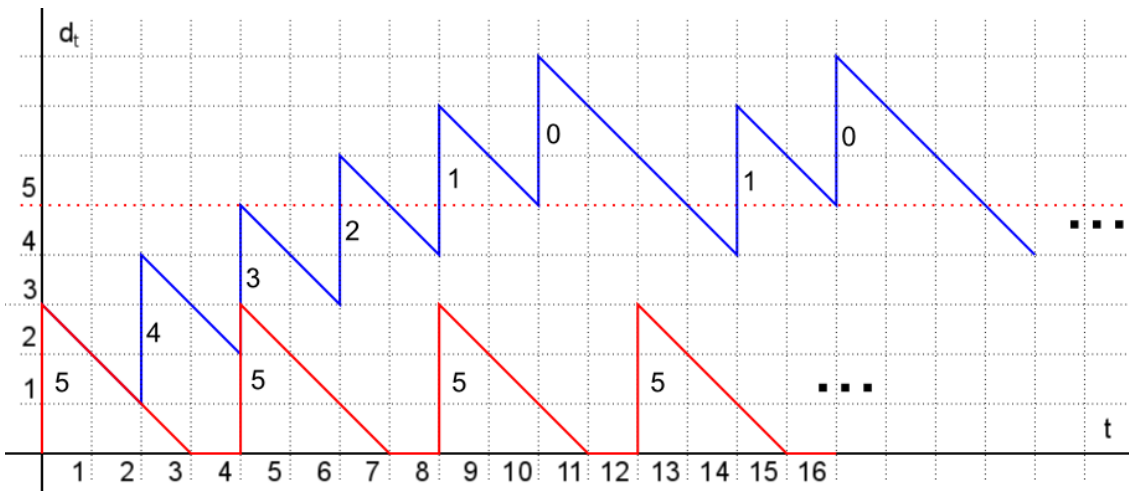

Figure 2: Plotting d_t as a function of t with price $p(d) = 5 - d$ (in blue) and $p(d) = 5$ (in red)