 Sid Banerjee (sbanerjee@cornell.edu)

Problem 1: Practice with Dynamic Programming Formulation

A product manager has to order stock daily. Each unit cost is c, there is a fixed cost of K for placing an order. If you order on day t, the items will be available at the beginning of day $t+1$ with probability $1-\theta$, and at the beginning of day $t+2$ with probability θ.

The store has N clients. Each day, a client demands a product independently with probability q. The selling price is $p > c$. Each time a client wants a product and there is no stock, the store experiences a loss of ℓ, representing a penalty for being unable to provide service.

The warehouse has capacity for $L > N$ units. At the beginning of day $t = 0$ there are S units in stock. At the end of period T all the units become obsolete and must be thrown away.

Write down a dynamic program to maximize the utility over T months. You should clearly define the states, actions and the value function, and write the terminal value function and the Bellman equation.

Solution: The state variables are

$$A_t = \text{number of available products at the start of day } t,$$
$$B_t = \text{number backlogged products that will arrive at } t+1.$$

The decision variables are

$$y_t = 1 \text{ if we order products for next day and 0 otherwise,}$$
$$x_t = \text{number of products we order for next day.}$$

The boundary conditions are $A_0 = S$ and $B_0 = 0$. The terminal value function, is

$$V^*_{T+1}(A_{T+1}, B_{T+1}) = 0.$$

At day t, the value function can be written as

$$\begin{aligned} V_t(A_t, B_t, x_t, y_t) = \sum_{n=0}^{N} \binom{N}{n} q^n (1-q)^{N-n} \Big(& p \cdot \min\{A_t, n\} - \ell \cdot \max\{n - A_t, 0\} \\ & + (1-\theta) V^*_{t+1}(\min\{L, A_t - \min\{A_t, n\} + x_t + B_t\}, 0) \\ & + \theta V^*_t(\min\{L, A_t - \min\{A_t, n\} + B_t\}, x_t) - K y_t - c x_t \Big). \end{aligned}$$

The Bellman equation is

$$V^*_t(A_t, B_t) = \max_{0 \le x_t \le L y_t} V_t(A_t, B_t, x_t, y_t).$$

Problem 2: Two-class Capacity Allocation with Upsell

In this problem, we will extend our basic two-class capacity allocation model from class to a setting where a fraction of the rejected discount customers are willing to purchase full-fare tickets.

As before, we have a total capacity C and let $p_\ell < p_h$ denote the price of the discount (low) and full-fare tickets (high) respectively. We again use D_ℓ to denote the random demand for the low fare, and let D_h denote the demand which is *exclusively for the full fare*, i.e., the demand for high fare assuming that all low-fare bookings are accepted. We assume that D_ℓ and D_h are independent discrete random variables.

To model the *upsell opportunity* from rejected discount customers, we assume that a *known* fraction $\alpha \in [0,1]$ of the *rejected* low-fare demand will seek to book high-fare tickets, if the low-fare seats are not available. Thus, given a booking limit b for discount tickets, the total high-fare demand $\bar{D}_h$ is given by:

$$\bar{D}_h = D_h + \alpha\,[D_\ell - b]^+$$

where for any $x \in \mathbb{R}$, $x^+ = \max\{x, 0\}$ denotes its positive part. Note that the booking limit decision influences the total full-fare demand when discount demand exceeds the booking limit.

Part (a)

Let $R(b)$ denote the expected total revenue given that the booking limit for discount customers is set at b. Show that

$$R(b) = p_\ell \mathbb{E}\left[\min\{D_\ell, b\}\right] + p_h \mathbb{E}\left[\min\left\{D_h + \alpha\,[D_\ell - b]^+, C - \min\{D_\ell, b\}\right\}\right]$$

Solution: This follows from observing that given a booking limit b, we sell at most $\min\{D_\ell, b\}$ at price p_ℓ. This leaves $[D_\ell - b]^+$ low-fare customers without a seat, a fraction α of whom are willing to buy a high-fare seat. Thus, the total demand for high-fare seats is $D_h + \alpha\,[D_\ell - b]^+$, while the number of seats remaining is $C - \min\{D_\ell, b\}$.

Part (b)

(OPTIONAL) Show that the *forward difference* $\Delta R(b) = R(b+1) - R(b)$ is given by

$$\Delta R(b) = (p_\ell - \alpha p_h)\,\mathbb{P}[b < D_\ell] - p_h(1-\alpha)\mathbb{P}\left[\{b < D_\ell\} \text{ AND } \{D_h + \alpha\,(D_\ell - b) \geq C - b\}\right]$$

(Note: You should try and replicate the technique we use in Lecture 3 for discrete values).

Solution:

Part (c)

Using the expression for the derivative of the revenue function in part (b), determine the optimal booking limit when $\alpha \geq p_\ell / p_h$. Can you think of an intuition behind this answer?

Solution: Note that if $\alpha \geq p_\ell/p_h$, then both terms in the above expression for $\Delta R(b)$ are ≤ 0 for all b. This means that $\arg\max_{b\geq 0} R(b) = 0$, since $R(b+1) \leq R(b)$, $R(b+2) \leq R(b+1)$ and so on; hence the optimal booking limit is 0.

One way to view this is to consider a setting where the low-fare agents arrive randomly, and thus each agent has a probability α of being willing to buy a high-fare seat. Now observe that under the condition $\alpha p_h \geq p_\ell$, each agent looking for a low-fare seat gives a higher revenue on average if denied the seat – thus, the optimal policy is to sell all seats only at the high fare.

Part (d)

Next, consider the case when $\alpha < p_\ell/p_h$. In general, computing the optimal booking limit b^* is difficult because it requires computing the probability $\mathbb{P}[\{b < D_\ell\} \text{ AND } \{D_h + \alpha\,(D_\ell - b) \geq C - b\}]$. However, suppose we assume that the market is highly competitive and only a very small fractional of the rejected discount customers will actually purchase our full-fare tickets, that is, α is close to zero. In this case, we can approximate the probability as follows:

$$\begin{aligned}\mathbb{P}[(b < D_\ell) \text{ and } (D_f + \alpha\,(D_\ell - b) > C - b)] &\approx \mathbb{P}[\{b < D_\ell\} \text{ AND } \{D_f > C - b\}] \\ &= \mathbb{P}[b < D_\ell] \cdot \mathbb{P}[D_\ell > C - b]\ ,\end{aligned}$$

where the equality follows from our assumption that D_ℓ and D_f are independent random variables. Therefore, we have the following approximation for the derivative:

$$\Delta R(b) \approx (p_\ell - \alpha p_h)\,\mathbb{P}[b < D_\ell] - p_h(1-\alpha)\mathbb{P}[b < D_\ell]\mathbb{P}[D_h > C - b]\ .$$

Using the above, show that the optimal protection level x^* is approximately equal to

$$F_h^{-1}\left(\left(\frac{1}{1-\alpha}\right)\left(1 - \frac{p_\ell}{p_h}\right)\right)\ .$$

Note that when $\alpha = 0$, the above expression gives the original Littlewood's rule.

Solution: As in the lecture, the optimal booking limit $b^* = \min_{b\geq 0}\{\Delta R(b) \leq 0\}$. Now under the proposed approximation, we want to find the smallest b such that:

$$\begin{aligned}&(p_\ell - \alpha p_h)\,\mathbb{P}[b < D_\ell] \leq p_h(1-\alpha)\mathbb{P}[b < D_\ell]\mathbb{P}[D_h > C - b] \\ &\Leftrightarrow (p_\ell - \alpha p_h) \leq p_h(1-\alpha)(1 - F_h(C-b)) \\ &\Leftrightarrow x^* = C - b^* = \min_{x\geq 0}\left\{F_h(x) \geq \left(\frac{1}{1-\alpha}\right)\left(1 - \frac{p_\ell}{p_h}\right)\right\}.\end{aligned}$$

Problem 3: Practice with convex/concave functions

In class we used several properties of concave functions in deriving the protection levels. We now briefly revise some of these, as they will be useful for later topics as well.

We use the following definition: a function $f : \mathbb{R} \to \mathbb{R}$ is convex if for every pair $x, y \in \mathbb{R}$ and every $\lambda \in [0,1]$, we have $f(\lambda x + (1-\lambda)y) \leq \lambda f(x) + (1-\lambda)f(y)$. f is concave if $-f$ is convex.

Part (a)

Given convex functions $f_1(\cdot), f_2(\cdot)$, show that the following are convex: **(1)** $g(x) = f_1(ax+b)$ for any $a, b \in \mathbb{R}$; **(2)** $g(x) = af_1(x) + bf_2(x)$ for $a, b > 0$; **(3)** $h(x) = \max\{f_1(x), f_2(x)\}$

Solution: **(1)**

$g(\lambda x+(1-\lambda)y) = f_1(\lambda(ax+b)+(1-\lambda)(ay+b)) \leq \lambda f_1(ax+b)+(1-\lambda)f_1(ay+b) = \lambda g(x)+(1-\lambda)g(y).$

(2)

$$\begin{aligned}
g(\lambda x + (1-\lambda)y) &= af_1(\lambda x + (1-\lambda)y) + bf_2(\lambda x + (1-\lambda)y) \\
&\leq a[\lambda f_1(x) + (1-\lambda)f_1(y)] + b[\lambda f_2(x) + (1-\lambda)f_2(y)] \\
&= \lambda g(x) + (1-\lambda)g(y).
\end{aligned}$$

Notice that for the second line we need $a, b \geq 0$, otherwise the inequality may be reversed.

(3) One can check that, for any $a, b, c, d \in \mathbb{R}$, it holds $\max\{a+b, c+d\} \leq \max\{a, c\} + \max\{b, d\}$. Using this,

$$\begin{aligned}
h(\lambda x + (1-\lambda)y) &= \max\{f_1(\lambda x + (1-\lambda)y), f_2(\lambda x + (1-\lambda)y)\} \\
&\leq \max\{\lambda f_1(x) + (1-\lambda)f_1(y), \lambda f_2(x) + (1-\lambda)f_2(y))\} \\
&\leq \max\{\lambda f_1(x), \lambda f_2(x)\} + \max\{(1-\lambda)f_1(y), (1-\lambda)f_2(y)\} \\
&= \lambda h(x) + (1-\lambda)h(y).
\end{aligned}$$

Part (b)

(Discrete Jensen's Inequality) For any convex function $f(\cdot)$, n points $\{x_i : i \in \{1, 2, \ldots, n\}\}$ and $\{\theta_1, \theta_2, \ldots, \theta_n\}$ such that $\theta_i \geq 0$ and $\sum_{i=1}^n \theta_i = 1$, prove that

$$f\left(\sum_{i=1}^n \theta_i x_i\right) \leq \sum_{i=1}^n \theta_i f(x_i)$$

Hint: This is true for 2 *points by definition - how can you extend to* n*?*

Solution: We use induction on n. Assume the result is true for $n-1$ points. Assume also that $\theta_n < 1$, otherwise the result is trivial (just one point). We want to express the sum of n points as $\lambda x + (1-\lambda)y$ for some well chosen x, y, λ. Let $x := \sum_{i=1}^{n-1} \frac{\theta_i}{1-\theta_n} x_i$ and $y = x_n$. If we define $\lambda := 1 - \theta_n$, then $\lambda x + (1-\lambda)y = \sum_{i=1}^n \theta_i x_i$. We use convexity for two points to obtain

$$f\left(\sum_{i=1}^n \theta_i x_i\right) = f(\lambda x + (1-\lambda)y) \leq (1-\theta_n) f\left(\sum_{i=1}^{n-1} \frac{\theta_i}{1-\theta_n} x_i\right) + \theta_n f(x_n).$$

Now we can use the induction hypothesis since we have $n-1$ points and $\sum_{i=1}^{n-1} \frac{\theta_i}{1-\theta_n} = 1$. A straightforward computation finishes the proof.

Problem 4: Two ways of computing protection levels

Suppose we have a capacity of 100 seats and 5 demand classes. As before, we assume that demands arrive sequentially, with class 5 (the lowest fare class) arriving first, and class 1 arriving last. The fare and demand distribution for each class are given in the the following table.

Class	Fare	Demand Distribution
1	$300	Geometric(1/10)
2	$200	Geometric(1/20)
3	$160	Geometric(1/25)
4	$140	Geometric(1/30)
5	$120	Geometric(1/40)

(Note that $Y \sim \text{Geometric}(p)$ $(p \in (0,1))$ means that $\mathbb{P}[Y = k] = (1-p)^k p$, for $k = 0, 1, \ldots$.)

Our aim is to compute the optimal protection levels.

Part (a)

Write down the dynamic program for the problem. You should give an expression for $V_1(s)$, and also write the Bellman equation.

Solution: First, note that $V_1(s) = p_1 \cdot \min\{s, D_1\}$; moreover, given the value function $V_k(\cdot)$ for any $k \in \{1, 2, 3, 4\}$, we have the Bellman equation:

$$V_{k+1}(s) = \min_{w \in \{0,1,\ldots,s\}} \mathbb{E}[p_{k+1} \cdot \min\{w, D_{k+1}\} + V_k(s - \min\{w, D_{k+1}\})]$$

In the lectures, we saw that the optimal policy corresponding to the above DP is in the form of a set of protection levels $\{x_k^*\}_{k \in \{1,2,\ldots,n-1\}}$, where x_k^* denotes the minimum number of seats we need to preserve for fare-classes $k, k+1, \ldots, 1$. Suppose we are given the optimal protection-level x_k^*; then the expression for the value-function is given by:

$$V_{k+1}(s) = \mathbb{E}[p_{k+1} \cdot \min\{D_{k+1}, (s - x_k^*)^+\} + V_k(\max\{s - D_{k+1}, x_k^*\})]$$

Part (b)

Write a program (in a language of your choice - ideally Python) to compute the protection levels. You should plot the value functions $V_k(s)$ to check concavity.

Hint: Note that you do not need to look at infinite values of s, since there is a a maximum capacity.

Part (c)

Now try and compute the protection levels using the sampling approach we discussed in class. You should generate a collection of K samples, for different K (say $K \in \{1000, 2000, 3000, \ldots, 10000\}$) and plot how your estimate of the protection levels x_k^* changes with K.