Revenue Management Review

Tu Ni

AY2019/2020 Semester 1

1 Introduction

- Structural Decision
- Pricing Decision
- Quantity Allocation Decision

2 Single-Product Price Optimization

2.1 Basic Pricing Model

Consider the simplest case when we want to sell a product with the unit cost c, and we have to set a price p, generating the expected demand $d(p)$. If there is no capacity limit, we can build a model to maximize the profit

$$max_p \ f(p) = pd(p) - cd(p)$$

The demand function $d(p)$ is characterized as market share times customer choice probability

$$d(p) = N\mathbb{P}(W \geq p)$$

In order to well analyze the problem, we need some assumptions on demand functions:

- $d(p) \geq 0, \forall p$
- $d(p)$ is strictly decreasing in p and has an inverse function $p(d)$
- $d(p)$ is differentiable with respect to p
- $\lim_{p\to\infty} d(p) = 0$

Several examples of demand functions are as follows:

- Linear: $d(p) = (a - bp)^+$, which is concave in feasible region
- Exponential: $d(p) = exp(a - bp)$, which is not concave
- Logit: $d(p) = N\frac{exp(a-bp)}{1+exp(a-bp)}$, which is quasi-concave

Remark. If we consider the revenue function over d, they are all concave.

To solve the problem, we look at the first-order necessary condition such that optimal p^* satisfy

$$p^* d'(p^*) + d(p^*) = cd'(p^*)$$

which means marginal revenue equals the marginal cost.

In order to better interpret this formula, we introduce the notion of Price Elasticity of Demand, which is defined as

$$\epsilon(p) = \frac{d'(p)p}{d(p)} = \frac{\Delta d}{\Delta p}\frac{p}{d} = \frac{\Delta d/d}{\Delta p/p}$$

The PED measures how sensitive the demand is to the change of price in percentage, and even for the same product, different groups of customers may have very different PED. Given this concept, we can re-write the optimality condition as

$$\epsilon(p^*) = \frac{d'(p^*)\,p^*}{d(p^*)} = -\frac{p^*}{p^* - c}$$

When the cost is negligible, the optimal price is found at PED equals -1. Normally, we assume PED to be decreasing so as to ensure a unique optimal solution.

Remark. In practice, the effect of cost is less important compared with demand function such as airline and hotel.

2.2 Price Discrimination

Since customers have different WTP's, one way to improve the revenue is to segment the consumers into different types with different prices, which is called price discrimination.

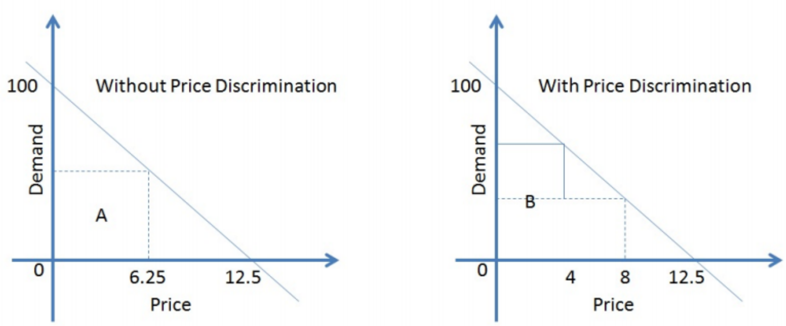

Figure 1: The power of price discrimination

If illustrating by demand function, the introduction of price discrimination helps extract higher proportion of potential maximal revenue, instead of one single rectangle. Ideally, if customers are perfectly discriminated, the maximal revenue can be achieved, which is called first degree price discrimination, but hard to implement. In practice, people usually segment customers into discrete groups by some criteria such as regions, purchase time, purchase channel, purchase effort or even age.

Mathematically, the pricing problem with inventory constraints can be easily formulated:

$$\begin{aligned} \text{maximize} \quad & \sum_{i=1}^{n} (p_i - c)\, d_i\,(p_i) \\ \text{s.t.} \quad & \sum_{i=1}^{n} d_i\,(p_i) \leq N \end{aligned}$$

If we write d as decision variables, it becomes a concave objective function with linear constraint.

In reality, unfortunately, people have to deal with cannibalization that high-value customers may reach out for low price. Specifically, we may either impose the fence to guarantee segmentation or incorporate possibility in modeling. Moreover, it becomes popular to let consumers "self-select" into appropriate groups, such as using coupons, product versioning/damaging, and especially opaque selling.

In order to analyze the effect of opaque selling, let us consider a simple model with following settings:

- A monopoly sells two types of products A and B
- Each consumer's valuation for product A and B are v_A and v_B , respectively
- v_A is uniformly distributed in $[0, 1]$ and $v_B = 1 - v_A$
- The utility of getting a product is the valuation minus the price
- Each consumer will choose a product with highest utility given it is positive; otherwise choose not to buy
- The monopoly wants to maximize the total expected revenue

Traditionally, the monopoly can optimize its expected revenue as

$$p_B \min \left\{ \frac{1 - p_B + p_A}{2}, 1 - p_B \right\} + p_A \left(1 - \max \left\{ \frac{1 - p_B + p_A}{2}, p_A \right\} \right)$$

with optimal price: $p_A = p_B = 1/2$ and optimal expected revenue $1/2$.

Now in the opaque-only selling, assume it sells an opaque product with price p_O with probability α the product is A; otherwise the product is B, and the condition for consumer purchasing the opaque product is

$$\alpha v_A + (1 - \alpha)(1 - v_A) - p_O \geq 0$$

In this way, the optimal price is $p_O = 1/2$ and optimal expected revenue is $1/2$, which are the same as before.

However, magic happens when consider opaque product together with original two products. In particular, the price of products A and B can be raised to make more revenue. Intuitively, the consumers are voluntarily segmented into three groups: {preferring A, preferring B, OK with both}.

Recall Fig.1, it seems that company makes more revenues by squeezing out consumers' welfare(surplus), which actually may not always be the case. If people greedily maximize the revenue then consumers' welfare will be squeeze out, but if people extract more revenue while considering consumers' welfare, it can be a win-win strategy, which are shown in Fig.2 and Fig.3

Intuitively, the key idea is to increase the sales, which is beneficial for company to get more revenue as well as for consumers to have more chances to buy.

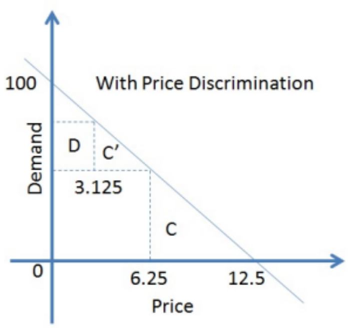

Figure 2:

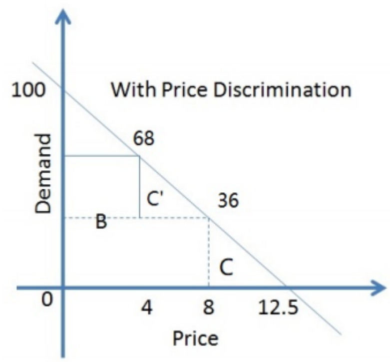

Figure 3:

2.3 Dynamic Pricing

Beyond single-period pricing, it's straightforward to extend to pricing under some finite time horizon, which is referred as dynamic pricing. Under this circumstance, the customer behavior along the time is essential. They may be myopic that just respond to current price or be strategic that refer to other factors including past prices or future information, or even other consumers' choices. Moreover, the market can be quite complicated such as competition level. To start with, we first consider the simplest case:

- There is a finite selling horizon T
- There is a fixed inventory x, which cannot be replenished during the selling season
- There is no value for the inventory after time T
- The seller needs to post a price p at each moment in the selling season
- The customers purchase according to a Poisson process with instantaneous arrival rate $\lambda(p)$
- The objective is to maximize the expected total revenue in T

Then we can formulate the problem as dynamic programming:

$$V_t(x) = \max_p \left\{\lambda(p)\Delta t\left(p + V_{t-\Delta t}(x-1)\right) + (1-\lambda(p)\Delta t)V_{t-\Delta t}(x)\right\}$$

which is equivalent to

$$\frac{\partial V_t(x)}{\partial t} = \max_p \left\{\lambda(p)\left(p + V_t(x-1) - V_t(x)\right)\right\}$$

Normally, this PDE is hard to solve but when the demand function is exponential with

$$\lambda(p) = ae^{-p}$$

We can solve it in closed-form as

$$V_t(x) = \log\left(\sum_{i=0}^{x}\left(\frac{at}{e}\right)^i \frac{1}{i!}\right)$$
$$p^*(t,x) = V_t(x) - V_t(x-1) + 1$$

For general cases, we may refer to discretization such that Δt is fixed and one should maximize

$$\lambda(p)\left(p + V_{t-1}(x-1) - V_{t-1}(x)\right)$$

Solving DP is hard so we may consider a deterministic case that the demand will be a fluid of rate $\lambda(p)$, so as for the inverse function $p(\lambda)$. Then the problem can be written as

$$\begin{array}{ll} \text{maximize} & \int_0^T p\left(\lambda_t\right) \lambda_t dt \\ \text{subject to} & \int_0^T \lambda_t dt \leq C \end{array}$$

The optimal solution to this deterministic problem is a constant pricing policy, along with optimal value $V_t^D(x)$, which is also an upper bound of DP value function. A natural question is how good is fixed price policy in original stochastic setting. It has been proved that this policy is asymptotically optimal as

$$\frac{V_t^{FP}(x)}{V_t^D(x)} \geq 1 - \frac{1}{2\sqrt{\min\left\{x, \lambda^* t\right\}}}$$

where $V_t^{FP}(x) = p^* \mathbb{E}\left[\min\left\{x, N_{\lambda(p^*)t}\right\}\right]$. Intuitively, fixed price policy captures the first-order effect of the problem while dynamic pricing adjusts to the demand fluctuation which captures the second-order effect.

3 Single-Resource Capacity Control

Beyond price and demand, revenue management usually has to deal with another important concern: limited resource, or specially capacity control. The most classical problem in this context arises from seat control in airline operations. Normally, there are several fare classes for the same flight to target on different types of customers. For example, leisure customers have low WTP and business customers have high WTP. In this chapter, we will study the control policy for revenue maximization with limited seats and given prices.

3.1 Static Model

The basic static model has following setting

- There are D_1 business passengers who afford p_1 dollars
- There are D_2 leisure passengers who afford $p_2 < p_1$ dollars
- Assume that leisure passengers arrive prior to business passengers

We need to decide a capacity limit K with p_2 for a certain flight with C seats. Clearly, there is a trade-off in choosing K and the problem is not trivial when D_1 and D_2 are random in practice. We can formulate the problem as

$$r(K) = p_2 \cdot \mathbb{E} \min\left(D_2, K\right) + p_1 \cdot \mathbb{E} \min\left(D_1, C - \min\left(D_2, K\right)\right)$$

If demands of different passengers are independent, we may denote $y = C - K$ as the number of seats that are reserved for the high-value passengers, and derive the optimality condition

$$P\left(D_1 \geq y\right) = \frac{p_2}{p_1}$$

Intuitively, it reaches some equilibrium that when you accept a low fare passenger, the immediate reward is p_2, but the opportunity cost is $p_1 P\left(D_1 \geq y\right)$. This model is known as Littlewood's model, which can be extended to problem with n classes easily. Now, we have to decide a sequence of protection levels $y_1, y_2, ..., y_{n-1}$ for $n-1$ classes. Clearly, we can use recursion to analyze the solution. Since the solutions of last two periods are the same as above, we can directly investigate stage 3 when we have to reserve seats for class 1 and 2. The value function can be represented as

$$V_2(x) = p_2\mathbb{E}\min\left(D_2, x - y_1^*\right) + p_1\mathbb{E}\min\left(D_1, \max\left(x - D_2, y_1^*\right)\right)$$

Follow the idea before, the derivative of value function(opportunity cost) should equal to p_3 in optimal condition. Thus, we get

$$P\left(D_1 \geq y_1^*, D_1 + D_2 \geq y_2^*\right) = \frac{p_3}{p_1}$$

and in general

$$\frac{p_{j+1}}{p_1} = P\left(D_1 \geq y_1^*, D_1 + D_2 \geq y_2^*, \ldots, D_1 + \cdots + D_j \geq y_j^*\right)$$

Although this nice optimal policy can be computed efficiently, most RM systems in practice use some heuristics, such as EMSR-a and EMSR-b. Suppose we are in the stage 3, the idea of EMSR-a is to reserve space for the remaining two classes separately using Littlewood's rule. In particular, it reserves $\overline{y}_2$ seats for class 2 with $P\left(D_2 \geq \overline{y}_2\right) = p_3/p_2$ and reserves $\overline{y}_1$ seats for class 1 with $P\left(D_1 \geq \overline{y}_1\right) = p_3/p_1$. Then, the idea goes for all previous stages to get 'optimal' protection level. Unfortunately, it's only powerful for a short while but less accurate in the long term because it ignores the pooling effect of different classes when dealing with uncertainty. Moreover, EMSR-b aims to reserve seats for remaining classes as a group with a weighted-average price

$$\overline{p}_2 = \frac{p_1 E\left[D_1\right] + p_2 E\left[D_2\right]}{E\left[D_1\right] + E\left[D_2\right]}$$

and choose y_2^* such that $P\left(S_2 \geq y_2^*\right) = p_3/\overline{p}_2$. Numerically. On top of that, instead of assuming some distribution structure, we can use historical samples more directly and adaptively. Define the event

$$B_j\left(\mathbf{y}^*, D\right) = I\left(D_1 \geq y_1^*, D_1 + D_2 \geq y_2^*, \ldots, D_1 + \cdots + D_j \geq y_j^*\right)$$

which expected to occur at frequency p_{j+1}/p_1. Therefore, during the process of sample collection, if we observe one sample that the event is true, we want to adjust y_j^* upwards, otherwise we adjust downwards. Finally, we will get an adaptive algorithm

$$\mathbf{y}_j^{k+1} = \mathbf{y}_j^k - \gamma_k\left(\frac{p_{j+1}}{p_1} - B_j(\mathbf{y}, D)\right) \quad \forall j$$

If learning rate is properly chosen, it will converge to optimal solution.

3.2 Dynamic Model

In the previous model, we assumed that the customers arrive in the order of their WTP. However, in practice, some high fare passengers may come earlier and some low fare passengers may come later. In this section, we assume each customer comes and requests a certain class of fares, and you need to decide whether to accept them or not (without knowing the future arrival customers).

The dynamic model kicks in to solve the problem with following settings

- There are T periods in total indexed forward (the first period is 1 and the last period is T)
- There are C inventory at the beginning
- Customers belong to n class, with $p_1 > p_2 > \cdots > p_n$ and requests to buy a ticket at price p_i

- Each period is small enough so that there is at most one arrival in each period

we need to decide which fare class to accept at period t with x inventory remaining. Let $\boldsymbol{u} \in \{0,1\}^n$ to be the decision variable at period t, then we have the DP formula

$$V_t(x) = \max_{\boldsymbol{u}} \left\{ \sum_{i=1}^{n} \lambda_i \left(p_i u_i + V_{t+1}\left(x - u_i\right)\right) + \left(1 - \sum_{i=1}^{n} \lambda_i\right) V_{t+1}(x) \right\}$$

with following optimal bid-price control policy

$$u_i = \begin{cases} 1 & p_i \geq \Delta V_{t+1}(x) \\ 0 & p_i < \Delta V_{t+1}(x) \end{cases}$$

where $\Delta V_{t+1}(x) = V_{t+1}(x) - V_{t+1}(x-1)$ represents opportunity cost. It's intuitive to get two properties:

$$\Delta V_t(x+1) \leq \Delta V_t(x) \quad \Delta V_{t+1}(x) \leq \Delta V_t(x)$$

Let's consider some extensions with more practical concern. First, assume for each low fare passenger, if low fare ticker are not available, there is some probability q that he will purchase high fare ticket, which referred as buy-up effect. We can easily adjust Littlewood's rule as

$$p_2 = (1-q) p_1 P\left(D_1 > x\right) + q p_1$$

however, this closed-form expression only works for two classes.

In fact, the issue above is just a special case incorporating customer choice. Generally, for each customer, he may have more complex choice behavior given many fare classes available instead of picking the cheapest one. Mathematically, we do not distinguish different arrival rates as λ_i but an aggregated arrival rate λ instead. At each period t, a subset of classes S_t is available such that the purchase probability of class j is $P_j(S_t)$, which satisfies the property of choice model. Specifically, if it is consistent with MNL model

$$P_i(S) = \frac{\exp\left(u_i\right)}{1 + \sum_{j \in S} \exp\left(u_j\right)}$$

then the only efficient sets are revenue ordered set such that if you offer class i, you should offer all classes that are higher than i and as inventory decreases, one should gradually close the low fare classes.

The point we need to attention is that the gap between standard Littlewood's model and choice-based model can be huge, which is known as Spiral-Down effect. In particular, the demand of class 1 tickets becomes endogenous that depends on the control policy. In this case, even though the demand distribution normally will converge and is consistent with resulting observation, the revenue may be severely hurt because the empirical demand is interpreted in a wrong way, so as the control policy. Therefore, it's really critical to consider the effect of customer choice behavior for capacity control.

Last but not least, overbooking model is becoming popular in practice. A simple model to analyze has following settings:

- The capacity is C and there is only one fare class p
- Each denied boarding has a cost h
- Each passenger has probability q to show up

- Therefore, the actual arrival will be a binomial random variable $B(n, q)$, where n is the total number of tickets sold

We need to decide the actual number of available seats to sell to maximize the expected revenue and possibly satisfy some service-level constraints. The expected revenue of having y seats is

$$V(y) = py - h\mathbb{E}\left[(B(y, q) - C)^+\right]$$

which is concave in y, and the optimal booking limit y^* is the largest value of y satisfying

$$V(x) - V(x-1) = hqP(B(x-1, q) \geq C) \leq p$$

It automatically infers Type I service level guarantee such that

$$P(B(x-1, q) \geq C) \leq p/hq$$

In reality, one may more care about what fraction of passengers will be bumped, which is called Type II service level guarantee. It can be computed as

$$s_2(x) = \frac{\mathbb{E}[(B(x) - C)]^+}{\mathbb{E}[Z(x)]}$$

and one can easily make it less than a certain level.

4 Network Revenue Management

In many real-world problems, capacity control has to deal with multiple resources together such as network revenue management, which is very popular especially in airline industry. We consider a standard network capacity control problem with m resources and n products, such that we can view the resources as the flight legs, and products as itineraries with some given prices. We analyze the model with following setting:

- each product uses at most one unit of each resource, but may use multiple types
- prices p_j's are predetermined
- we could have multiple products with same resource pattern but different p_j
- the selling horizon is T and at each period t, at most one request is made by product j with probability q_{jt}

We need to decide whether to accept the request of product j at each period under resource constraint so that the total expected revenue during the T periods is maximized. Again, we can use DP to formulate this problem such that $V_t(\boldsymbol{x})$ denotes the maximum expected revenue to go at period t and with remaining inventory $\boldsymbol{x}$. We have the following Bellman equation

$$V_t(x) = \max_{\boldsymbol{u} \in U(\boldsymbol{x})} \left[\sum_{j=1}^{n} q_{it}\left(p_j u_j + V_{t+1}\left(\boldsymbol{x} - A_j u_j\right)\right) + \left(1 - \sum_{j=1}^{n} q_{jt}\right) V_{t+1}(\boldsymbol{x})\right]$$

and its feasible region

$$U(x) = \{\boldsymbol{u} \in \{0,1\}^n : \mathbf{a}_j u_j \leq \mathbf{x}, \forall j\}$$

Similar to single resource case, we can also result in the optimal decision in the manner of bid-price

$$u_j^*(t,\boldsymbol{x}) = \begin{cases} 1 & \text{if } p_j \geq V_{t+1}(\boldsymbol{x}) - V_{t+1}(\boldsymbol{x}-\boldsymbol{a}_j) \text{ and } A_j u_j \leq \boldsymbol{x} \\ 0 & \text{otherwise} \end{cases}$$

However, the number of states with multiple resources will be exponential and thus the DP is hard to solve directly, but can be approximated in some way. Recall the dynamic pricing model with DP formulation, we use deterministic model to find an upper bound. In fact, this technique also fits into network capacity control problem. We can solve a deterministic linear programming with demand represented by its expectation

$$\begin{array}{rl} V_t^{LP}(x) = \max & \sum_{j=1}^n p_j y_j \\ \text{subject to} & \sum_{j=1}^n \boldsymbol{a}_j y_j \leq \boldsymbol{x} \\ & 0 \leq y_j \leq d_{jt} \end{array}$$

where $d_{jt} = \sum_{\tau=t}^T q_{j\tau}$

Clearly, the solution of DLP provides an upper bound of the optimal revenue, which can be utilized to construct some control policies for the original problem.

- Partition allocation rule: set as limit the optimal solution y^* of DLP; suffer from demand fluctuation
- Bid price rule: set as marginal value of resource the dual price of each constraint, and by complementary slackness, we have decision rule $p_j \geq \boldsymbol{a}_j^T \boldsymbol{p}^*$; may update the dual price frequently but no guarantee to increase the expected revenue
- Certainty equivalent control: approximate the value function in the optimal decision rule with DLP $p_j \geq V_{t+1}^{LP}(\boldsymbol{x}) - V_{t+1}^{LP}(\boldsymbol{x}-\boldsymbol{a}_j)$

It has been shown that partition allocation rule and bid price rule are asymptotically optimal.

Another stream of technique tries to improve DLP directly, such as randomized linear programming. Instead of using expected demand to go, RLP keeps the demand as a random variable. We can take the expectation of this RLP objective value $\mathbb{E}[H_t(\boldsymbol{x}, D)]$ by solving LP for each realization. Intuitively, this approximation should be better than DLP if we are able to sample some demand efficiently. Although partition allocation rule is not suitable here, bid price rule is still applicable because

$$\nabla_{\boldsymbol{x}} V_t^{RLP}(\boldsymbol{x}) = \nabla_{\boldsymbol{x}} \mathbb{E}[H_t(\boldsymbol{x}, D)]$$

Interchange the derivative and expectation, we can calculate the bid price as the average of the dual prices.

Last but not least, the common technique in solving DP is ADP. First we write the equivalent LP as

$$\begin{array}{ll} \min_{\{v_{t,\mathbf{x}}\}} & v_{0,C} \\ \text{s.t.} & v_{t,x} \geq v_{t+1,x} + \sum_{j=1}^n q_{jt} u_j(t,x) \left(p_j + v_{t+1,x-\mathbf{a}_j} - v_{t+1,x}\right) \quad \forall t, \boldsymbol{x}, \boldsymbol{u}(t,\boldsymbol{x}) \end{array}$$

where $v_{t,\boldsymbol{x}}$ is the decision variable we care about. If we restrict those value functions in affline form

$$\tilde{v}_{t,\boldsymbol{x}} = \sum_{\tau=t}^T \theta_t + \lambda^T x$$

then clearly the corresponding optimal value performs an upper bound of the original DP. With this restriction, the LP can be simplified to

$$
\begin{aligned}
\min_{\theta, \boldsymbol{\lambda}, \boldsymbol{z}} & \sum_{t=1}^{T} \theta_t + \boldsymbol{\lambda}^T C \\
\text{s.t. } & \theta_t - \boldsymbol{q}_t^T \boldsymbol{z} \geq 0, \quad \forall t = 1, \ldots, T \\
& z + \boldsymbol{\lambda}^T A \geq p \\
& \mathbf{z} \geq 0, \boldsymbol{\lambda} \geq 0
\end{aligned}
$$

with its dual

$$
\begin{array}{ll}
\max_{\mathbf{y}} & \boldsymbol{p}^T \mathbf{y} \\
\text{s.t.} & A^T \mathbf{y} \leq C \\
& 0 \leq \mathbf{y} \leq \sum_{t=1}^{T} \boldsymbol{q}_t
\end{array}
$$

which is nothing but the DLP above. One way to strengthen its approximation is to consider time-dependent bid price λ_t, but the LP is not trivial to solve.

5 Discrete Choice Model

In many practical applications, customers are given multiple alternatives and only need to purchase one of them, such as flight tickets and hotel offers. Normally, we suppose the choice set is finite and a customer may not purchase any product. This type of discrete choice model is quite important in demand modeling and understanding the customer behavior so that many operational decisions like price and assortment can be made in an optimal manner.

5.1 Parametric Model

The parametric model assumes the choice probability depends on the deterministic utilities of each alternative.

A classic framework is Random Utility Model that assumes each customer has a random utility $U_i = \mu_i + \epsilon_i$ on each alternative $i \in \mathcal{N} = \{1, ..., n\}$, where μ_i is deterministic cross all customers and ϵ_i is a random variable that captures heterogeneity among customers and cannot be observed. The key assumption in RUM is that each customer picks the alternative with the highest utility

$$
q_i(\boldsymbol{\mu}) = \mathbb{P}\left(i = \operatorname{argmax}_{k \in \mathcal{N}} \left(\mu_k + \epsilon_k\right)\right)
$$

Multinomial Logit Model is the most popular choice model, which requires ϵ to be i.i.d Gumbel distributions $F(x) = e^{-e^{-x/\eta}}$, then the resulting choice probability has closed-form expression

$$
q_i(\boldsymbol{\mu}) = \mathbb{P}\left(U_i > U_j, \forall j \neq i\right) = \frac{\exp\left(\mu_i/\eta\right)}{\sum_{k \in \mathcal{N}} \exp\left(\mu_k/\eta\right)}
$$

where the scale parameter η can be interpreted as the rationality level: if it tends to infinity, the choice is completely random; if it tends to zero, the choice is dominated by deterministic utility.

Moreover, MNL preserves the Independence of Irrelevant Alternative property, which says the ratio of choice probabilities between two alternatives does not depend on other alternatives

$$
\frac{q_i(\boldsymbol{\mu})}{q_j(\boldsymbol{\mu})} = \frac{\exp\left(\mu_i/\eta\right)}{\exp\left(\mu_j/\eta\right)} = \exp\left(\left(\mu_i - \mu_j\right)/\eta\right)
$$

This property is mathematically concise but may not be realistic in some cases. In fact, this property only makes sense when the feature of different alternatives are homogeneous, which means customers compare them under some common criteria, i.e. convenience, price, time. Moreover, the only RUM that satisfy IIA property is the MNL.

In order to deal with non-homogeneous problem, we can build a two-stage(multi-stage) choice system, which is referred as Nested Logit Model. Instead of choosing alternative directly, we partition all alternatives into K nests as the intermediate with the choice probability

$$q_i(\boldsymbol{\mu}) = \frac{\left(\sum_{j \in B_k} \exp\left(\mu_j/\lambda_k\right)\right)^{\lambda_k}}{\sum_{l=1}^{K}\left(\sum_{j \in B_l} \exp\left(\mu_j/\lambda_l\right)\right)^{\lambda_l}} \cdot \frac{\exp\left(\mu_i/\lambda_k\right)}{\sum_{j \in B_k} \exp\left(\mu_j/\lambda_k\right)}$$

so customers first choose a nest, and then choose a product within the nest. Clearly, the second stage is nothing but a MNL model and IIA holds, while the first stage normally does not preserve IIA unless $\lambda_k = 1$ for all k.

Since NL is an extension of MNL, we'd like to understand the condition, under which it belongs to RUM. Specifically, instead of i.i.d ϵ across alternatives, NL model has correlated ϵ with following cdf

$$F(\epsilon) = \exp\left(-\sum_{k=1}^{K}\left(\sum_{i \in B_k} e^{-\epsilon_i/\lambda_k}\right)^{\lambda_k}\right) = \prod_{k=1}^{K} \exp\left(-\left(\sum_{i \in B_k} e^{-\epsilon_i/\lambda_k}\right)^{\lambda_k}\right)$$

if i and j belong to different nests, they are un-correlated while if belong to the same nest, they are correlated. $\lambda_k \in (0, 1]$ measures the independence of i and j in nest k.

By investigating the distribution of ϵ, we observe that NL model is actually a special case of the generalized extreme value models, where ϵ follows a generalized extreme value distribution

$$F(\epsilon) = \exp\left(-G\left(e^{-\epsilon_1}, \ldots, e^{-\epsilon_n}\right)\right)$$

where G is a generating function. In particular, we can recall the form of MNL and NL

- MNL: $G\left(y_1, \ldots, y_n\right) = \sum_{k=1}^{K}\left(\sum_{i \in B_k} y_i^{1/\lambda_k}\right)^{\lambda_k}$
- NL: $G = \sum_{i=1}^{n} y_i$

Further, we figure out what kind of G function is consistent with RUM. Motivated by NL, we can come up with following conditions:

- Non-negativity: $G(\boldsymbol{y}) \geq 0$
- Positive homogeneity: $G(\lambda \mathbf{y}) = \lambda G(\mathbf{y}), \forall \lambda \geq 0$
- Infinity: $y_i \rightarrow +\infty \Rightarrow G(\mathbf{y}) \rightarrow +\infty$ for all i
- Alternating signs of derivatives: the k -th order cross partial derivatives has sign$(-1)^{k+1}$

Beyond MNL and NL, some other RUM models are proposed such as ϵ is normal or exponentially distributed, but they are more complicated either in expression or estimation.

An important extension of RUM is to incorporate product features. For example, if we use MNL,

$$q_i = \frac{\exp\left(\boldsymbol{\beta}^T \boldsymbol{x}_i\right)}{\sum_{j=1}^n \exp\left(\boldsymbol{\beta}^T \boldsymbol{x}_j\right)}, \forall i \in \mathcal{N}$$

where $\boldsymbol{x}$ is the feature of alternative and $\boldsymbol{\beta}$ is the coefficient, which can be set to customer-dependent by clustering customers into groups. Even though we are not able to partition customer types directly, we may model β as random variables with some given distribution $g(\cdot)$ so that the choice probability is the average of these customers

$$q_i = \int \frac{\exp\left(\boldsymbol{\beta}^T \boldsymbol{x}_i\right)}{\sum_{j=1}^n \exp\left(\boldsymbol{\beta}^T \boldsymbol{x}_j\right)} g(\boldsymbol{\beta}) d\boldsymbol{\beta}$$

which called Mixed Logit model, by which any RUM can be approximated arbitrarily closely.

Although RUM is concise both theoretically and practically, it stills relies on some strong assumptions about the notion of utility, which are violated in many practical instances. For example, regularity means that if $i \in S \subseteq T$, then $q_{i,S} \geq q_{i,T}$, which is satisfied by all RUMs, but this property is disproved by halo effect, decoy effect and compromising effect. Mathematically, if $\lambda_k > 1$ is NL, we can observe some situation that is not consistent with RUM.

From another point of view, we consider Representative Agent Model where a central planner makes choice $\boldsymbol{q}$ on behalf of the market. He has to consider the expected utility $\boldsymbol{\mu}'\boldsymbol{q}$ and also wants to achieve some diversification to serve the market. Mathematically, he chooses

$$\boldsymbol{q}(\boldsymbol{\mu}) = \operatorname{argmax}_{\boldsymbol{x} \in \Delta_{n-1}} \boldsymbol{\mu}^T \boldsymbol{x} - V(\boldsymbol{x})$$

where $\Delta_{n-1} = \{x | \sum_{i=1}^n x_i = 1, x \geq 0\}$. Suppose $V(\boldsymbol{x})$ is the negative entropy $\eta \sum_{i=1}^n x_i \log x_i$, then the choice decision recovers the MNL model. In fact, for any RUM, one can always find a corresponding convex $V(\cdot)$, so RAM is one way to generalize RUM.

Another generalization of RUM is semi-parametric choice model, which assumes ϵ is from a set of distributions Θ. Thus the choice probability is

$$q_i^S(\boldsymbol{\mu}) = \mathbb{P}_{\theta^*(\boldsymbol{\mu})} \left(i = \underset{k=1,\ldots,n}{\operatorname{argmax}} \left(\mu_k + \epsilon_k\right) \right)$$

where $\theta^*(\boldsymbol{\mu}) = \underset{\theta \in \Theta}{\arg\sup} \mathbb{E}_{\boldsymbol{\epsilon} \sim \theta} \left[\max_{i=1,\ldots,n} \mu_i + \epsilon_i\right]$. Different types of Θ have been proposed including MDM, MMM and CMM. Essentially, it has been shown that RAM and SPCM are equivalent.

5.2 Non-Parametric Model

The rank list model assumes each customer has a preference list σ over all alternatives with totally $n!$ types, and one chooses the best alternative available in S. It has been shown that $q_{i,S}$ comes from a rank list model if and only if there exists a distribution of U such that

$$q_{i,S} = \mathbb{P}\left(U_i \geq U_j, \forall j \in S\right)$$

The Markov Chain model assumes that customer comes with a most preferred item i. If the item is available, then chooses that option; otherwise, the customer will switch to another choice j with probability π_{ij}. In fact, this model can recover MNL, and is a special case of rank list model.

5.3 Estimation of Choice Model

Essentially, there are two goals of estimation: identify the relative importance among features and find the correct parameters in choice model.

First, consider the feature-based choice model where the utility is a linear function of features: $\mu_i = \boldsymbol{\beta}^T \boldsymbol{x}_i$, and thus we could calculate the prediction of choice probability $q_{i,S_k}(\boldsymbol{\beta})$ for a given assortment S_k. Empirically, the frequency of i being picked in S_k is denoted as y_{i,S_k}. The optimal $\boldsymbol{\beta}$ can be found by maximizing the log likelihood function, or equivalently minimizing the cross-entropy

$$\max LL(\beta) = \sum_{k=1}^{K} m_k \sum_{i \in S_k \cup \{0\}} y_{i,S_k} \log q_{i,S_k}(\beta)$$

which is concave in many choice models including nested logit, exponomial, and MDM. Compared with least square error, this objective prevents the predicted choice probability being zero.

Second, to incorporate model-specific parameters, the log-likelihood function will extent to $LL(\beta, \gamma)$ and optimize alternatively.

6 Multi-Product Pricing and Assortment Optimization

Discrete choice models are closely related to many operational decisions, such as multi-product pricing that a firm sets prices for its n products to maximize revenue, and assortment optimization that a firm decides the subset of products to offer given the prices.

6.1 Multi-Product Pricing

When customers see the price $\boldsymbol{p}$, they would purchase the product based on a choice model $\boldsymbol{q}(\boldsymbol{u})$ where the expected utility $\boldsymbol{u}$ is a linear function of $\boldsymbol{p}$ with $\mu_i = a_i - b_i p_i, \forall i = 1, \ldots, n$. This setting is easy for analysis but in practice, the function could be more complicated that depends on the historical prices, etc. The profit maximization problem can be written as

$$\text{maximize} \quad \pi(\boldsymbol{p}) = \sum_{i=1}^{n} (p_i - c_i)\, q_i(\boldsymbol{p})$$

We may first consider the demand is MNL, due to the form of $q_i(\boldsymbol{p})$, the objective function is not even quasi-concave in $\boldsymbol{p}$. Fortunately, it has been shown that for $\boldsymbol{p}^*$ as the optimal price, it must hold that $p_i^* - c_i - 1/b_i$ is constant for all $i \in \mathcal{N}$. We could set this constant as θ, which is precisely the optimal profit, then we reduce the problem to single variable

$$\theta = \sum_{i=1}^{n} \left(\frac{1}{b_i} + \theta \right) \frac{\exp(a_i - b_i c_i - b_i \theta - 1)}{1 + \sum_{j=1}^{n} \exp(a_j - b_j c_j - b_j \theta - 1)}$$

by showing that the equation has a unique solution, it's easy to get the optimal price.

Intuitively, we could view $p_i - c_i$ as the markup(marginal profit) on top of the cost. Compared with other products, the higher price sensitivity, the lower markup. If all b_i are equal, then we have equal markup property. This result is also proved to be true for all GEV models.

Remark. recall the problem we deal with in basic single product pricing, the objective is also not concave in price and by reformulating as a function of demand, it becomes concave. Similar idea can be applied to multi-product pricing as well, where the demand is interpreted as market share or choice probability.

In practice, we also face some constraints on prices and market shares, such as the certain price range, inventory limit, etc. Based on the analysis above, we can see the difficulty adding price constraints but some convenience of market share constraints.

6.2 Assortment Optimization

In some cases, such as airline and hotel, the prices are pre-determined by outside factors, and the firm needs to decide which products to sell or offer to customers, which is so called assortment planning.

Let $q_{i,S}$ denote the probability for customers to choose product i in assortment S, then the assortment planning problem can be written as

$$\text{maximize} \quad R_S = \sum_{i \in S} (p_i - c_i)\, q_{i,S\cup\{0\}}$$

where $q_{i,S\cup\{0\}}$ comes from a choice model.

Intuitively, we may want to offer all products to provide customers with more choices. However, it may not be the case because the prices are given and they may not be appropriate under choice model. For the case of MNL, on the one hand, more products should be offered so that the probability of outside option or non-purchase can be squeezed; On the other hand, new product will 'steal' the market share of existing products, and thus if the price of new product is not set properly, the new product will cannibalize the revenue.

The explanation above motivates to simplify the assortment optimization by focusing on the price or revenue. We can order the product such that

$$p_1 - c_1 \geq p_2 - c_2 \geq \cdots \geq p_n - c_n$$

We call $S_k = \{1, \ldots, k\}, k = 0, \ldots, n$ revenue-ordered assortments, with $S_0 = \emptyset$. It has been shown in [2] that there must be one of the revenue-ordered assortment that is optimal under MNL, which reduces the computation to complexity $O(n)$. Furthermore, if the firm can display at most M products, a polynomial-time algorithm can be derived; if there is a capacity for each product, then the problem is equivalent to a linear program.

7 Dynamic Revenue Management with Network Effect and Path Dependency

7.1 Network Effect

First type of network effect is local network effect, whereby the purchase decision of a consumer is affected by his neighbors actions. Suppose there is a social network and the connections between friends will change the valuation of the goods, then the sellers may charge different prices to different buyers in the network to take advantage of the network effect and thus maximize the revenue. A popular model considers n customers embedded in a network by an influence matrix G, with $g_{ij} \geq 0, g_{ii} = 0$. The utility of customer i is

$$u_i\left(x_i, x_{-i}, p_i\right) = a_i x_i - b_i x_i^2 - p_i x_i + x_i \cdot \sum_{j=1}^{n} g_{ij} x_j$$

Given the prices, all the consumers are involved in a game to maximize their own utility. For ease of analysis, we will look for equilibrium solution, in which

$$x_i = \operatorname{argmax}_{y_i} u_i\left(y_i, x_{-i}, p_i\right)$$

we can get the equilibrium solution by solving a set of linear equations

$$x_i = \frac{a_i - p_i}{2b_i} + \frac{1}{2b_i} \sum_{j \neq i} g_{ij} x_j, \quad \forall i$$

Then the optimal pricing problem can be written as

$$\begin{array}{ll} \operatorname{maximize}_{\boldsymbol{p}, \boldsymbol{x}} & (\boldsymbol{p} - \boldsymbol{c})^T \boldsymbol{x} \\ \text{subject to} & (\Lambda - G)\boldsymbol{x} = \boldsymbol{a} - \boldsymbol{p} \end{array}$$

with the optimal solution

$$\boldsymbol{p}^* = \boldsymbol{a} - (\Lambda - G)\left(\Lambda - \frac{G + G^T}{2}\right)^{-1} \frac{\boldsymbol{a} - \boldsymbol{c}}{2}$$

by assuming $\boldsymbol{a} = a\boldsymbol{e}, \boldsymbol{b} = b\boldsymbol{e}$, we can rewrite as

$$\boldsymbol{p} = \frac{a + c}{2} 1 + \frac{a - c}{8b} G \mathcal{K}\left(\frac{G + G^T}{2}, \frac{1}{2b}\right) - \frac{a - c}{8b} G^T \mathcal{K}\left(\frac{G + G^T}{2}, \frac{1}{2b}\right)$$

the first term is uniform across all customers to capture basic utility structure; the second and third terms reflect the network effect.

Moreover, setting many different prices for different customers is usually impractical. A common situation is one full price together with the other discounted price, whereby the problem becomes

$$\begin{array}{ll} \text{maximize} & (\boldsymbol{p} - \boldsymbol{c})^T (\Lambda - G)^{-1} (\boldsymbol{a} - \boldsymbol{p}) \\ \text{subject to} & p_i \in \{p_L, p_H\} \end{array}$$

which is equivalent to a MAX-CUT problem.

Second type of network effect is global network effect, whereby the purchase decision is affected by the aggregate purchase behavior of others. In this case, there is no limit on number of customers, and thus at each time step k, a new participant will enter the market, observe the market share, make the choice and update the market share.

$$U_i = \overbrace{a_i - b_i p_i}^{\text{Intrinsic: } f_i} + \overbrace{\gamma h\left(q_i(k)\right)}^{\text{Network}} + \overbrace{\epsilon_i}^{\text{Noise}}$$

The stochastic process of the market share has been shown to asymptotically follow a deterministic path and converge to an equilibrium.

$$\bar{q}_i = \frac{\exp\left(a_i - b_i p_i + \gamma q_i\right)}{\sum_j \exp\left(a_j - b_j p_j + \gamma q_j\right)}$$

beyond this choice model, several related questions are studied such as multi-product pricing and assortment.

7.2 Path Dependency

Polya's Urn model is used to characterize stochastic choices in a dynamic process. Consider an example of trial-offer market, whereby at the first stage, a product is sampled from a MNL model with position bias, and at the second stage, a purchase/download decision is sampled according to a Bernoulli trial with a probability based on product quality.

Specifically, the customer decides whether to listen a song with probability $P_i(\sigma, d_t) = \frac{v_{\sigma_i}(\alpha a_i + d_{i,t})}{\sum_{j=1}^{n} v_{\sigma_j}(\alpha a_j + d_{j,t})}$, where v captures position bias, a denotes product appeal, and d counts the number of purchase. If we use $X_{i,t}$ to describe the number of balls of color i at time t, then the sampled probability for color i can be represented as $G_i(X_t) = \frac{v_{\sigma_i} X_{i,t}}{\sum_{j=1}^{n} v_{\sigma_j} X_{j,t}}$. With some probability, a new ball of color i will be added, which corresponds to the purchasing action. Under some condition, it can be shown that the market converges to a monopoly for the product of highest quality.

8 Online Optimization for Revenue Management

For the previous models, they rely on the fact that we know demand function or arrival distribution in advance. If these information are not known or only partially known, online optimization kicks in to get efficient policy. We will investigate two standard settings: competitive ratio and regret analysis. Interestingly, to this point, despite the wide variety of algorithms developed and analyzed in the two literatures, there are no algorithms that can guarantee good performance with respect to both the dynamic optimal(competitive ratio) and the static optimal solutions(regret).

8.1 Competitive Ratio without Data

Consider the basic problem in capacity control that we need to come up with a policy to serve two fare classes with limited capacity and no demand information. In this case, we are interested to analyze the worst-case performance with any potential arrival sequence

$$\Upsilon = \inf_{\mathcal{I} \in \Omega_\Upsilon} \frac{\nu'_\Upsilon(\mathcal{I})}{\nu^*_\Upsilon(\mathcal{I})}$$

where for any $\mathcal{I} \in \Omega_\Upsilon$, let $\nu'_\Upsilon(\mathcal{I})$ be the objective value achieved by the online algorithm for input I and let $\nu^*_\Upsilon(\mathcal{I})$ be the objective value achieved by an optimal offline algorithm. The key assumption is pointed out for analysis: demand for different fare classes are independent so that there is no cannibalization, no decision-dependent demand response, and customer type is fixed.

A typical result is shown by [1] that for the continuous two-fare problem, no deterministic online booking policy has a competitive ratio larger than $b(r) = 1/(2-r)$, where $r = f_2/f_1$. The result is proved by considering two extreme arrival sequences, along with their performance. Moreover, it shows for the two-fare booking problem, the policy with protection level $\theta_1 = (1-b(r))n$ has competitive ratio $b(r)$ exactly. The proof basically discusses two cases when either a high-fare class is rejected or a low-fare class is rejected, along with their worst-case analysis.

8.2 Regret with Online Data

Now we consider a different setting that given some domain $\mathcal{K}$, at each step t, we have to choose a decision $x_t \in \mathcal{K}$, and then the nature will give us a convex function f_t and its corresponding loss $f_t(x_t)$. Compared with previous model, we always get information updated for optimizing decision. In this case, we are interested to analyze the worst-case performance with any sequence of f_t

$$\text{regret }_T := \sum_{t=1}^{T} \left[f_t\left(x_t\right) - f_t\left(x^*\right)\right]$$

where x^* is the optimal stationary solution.

The main theorem for this technique starts from [3], in which a sublinear regret is proved under some proper conditions. Suppose $\max_{x \in \mathcal{K}} \|\nabla f_t(x)\| \leq G, \max_{x,y \in \mathcal{K}} \|x-y\| \leq D$. Online Gradient Descent with step sizes $\eta_t = D/(G\sqrt{t})$, for $t = 1, \ldots, T$, guarantees:

$$\text{regret }_T = \sum_{t=1}^{T} f_t\left(x_t\right) - \min_{x^* \in \mathcal{K}} \sum_{t=1}^{T} f_t\left(x^*\right) \leq \frac{3}{2} GD\sqrt{T}$$

where the OGD is achieved as

$$\begin{aligned} \mathbf{y}_{t+1} &= \mathbf{x}_t - \eta_t \nabla f_t\left(\mathbf{x}_t\right) \\ \mathbf{x}_{t+1} &= \Pi_{\mathcal{K}}\left(\mathbf{y}_{t+1}\right) \end{aligned}$$

The proof relies on the convexity of f_t and non-expansion property of projection to get a bound with respect to η_t, which can be set appropriately to achieve sublinear regret.

Essentially, the solution x_{t+1} in OGD can be obtained by solving a regularized linear approximation of $f_t(x)$:

$$x_{t+1} = \operatorname{argmin}_{x \in \mathcal{K}} \left\{\nabla f_t\left(x_t\right) \bullet x + \frac{1}{2\eta_t} \left\|x - x_t\right\|_2^2\right\}$$

More generally, we may consider other regularization term or Bregman Divergence associated with some function ϕ such that we can replace $\|w - v\|_2^2$ by

$$B_\psi(w, v) = \psi(w) - \psi(v) - \nabla\psi(v) \bullet (w - v)$$

which measures the gap between the true function value and its linear approximation. For example, let $\psi(w) = \sum_{j=1}^{d} w_j \log\left(w_j\right)$, it just measures KL divergence. Moreover, we can simplify the formulation to

$$x_{t+1} = \operatorname{argmin}_{x \in \mathcal{K}} B_\psi\left(x, y_{t+1}\right)$$

where $\nabla f_t\left(x_t\right) + \frac{1}{\eta_t}\left(\nabla\psi\left(y_{t+1}\right) - \nabla\psi\left(x_t\right)\right) = 0$

9 Data-driven Pricing

We aim to fit a RAM with aggregated sales data and then optimize the price. Normally, we have the following representation

$$\max_{\mathbf{x} \in \Delta_N} \sum_{j=0}^{N} \left(v_{ij} - \alpha p_j\right) x_{ij} - C(\mathrm{x}) \quad \Delta_N = \left\{\mathrm{x} \in \mathbb{R}_{+}^{N+1} \middle| \sum_{j=0}^{N} x_{ij} = 1\right\}$$

Suppose we are able to get the marginal distributions of utility residual for each product, and by separating $C(x)$, we end up with MDM that maximizes the expected utility under all joint distributions with given marginal distributions of $\tilde{\epsilon}_{ij}$:

$$\max_{\mathbf{x} \in \Delta_N} \sum_{j=0}^{N} \left((v_{ij} - \alpha p_j) x_{ij} + \int_{1-x_{ij}}^{1} F_j^{-1}(t) dt \right)$$

The optimality condition for MDM

$$p_j^* = \frac{v_{ij} + F_j^{-1}\left(1 - x_{ij}^*\right) - F_0^{-1}\left(1 - x_{i0}^*\right)}{\alpha}, \forall j = 1, 2, \ldots, N$$

which characterizes the relation between p and x. Then we can formulate the pricing problem with respect to market share

$$\begin{array}{ll} \max_x & -\sum_{j=1}^{N} w_j x_j + \frac{1}{\alpha} \sum_{j=1}^{N} x_j F_j^{-1}\left(1 - x_j\right) - \frac{1}{\alpha}\left(1 - x_0\right) F_0^{-1}\left(1 - x_0\right) \\ \text{s.t.} & \sum_{j=0}^{N} x_j = 1 \\ & x \geq 0 \end{array}$$

For tractability, we assume that the following two conditions hold:

C1. $xF_j^{-1}(1-x)$ for each product $j = 1, \ldots, N$ is a concave function.

C2. $(1-x)F_0^{-1}(1-x)$ for the outside option is a convex function. Then, the pricing problem is a convex optimization problem in the market share x variables and the optimal prices are computable in polynomial time.

Actually these two conditions are well satisfied for common distributions including normal, logistic, exponential, extreme value, Laplace distributions. To estimate $xF^{-1}(1-x) = y(x)$, we can estimate those discrete values $y(x_t)$ corresponding to the data, and then construct a piece-wise linear function to optimize in pricing problem. Apparently, some side constraints including monotonicity and convexity should be added in estimation process.

10 Bundle Pricing

If the seller has n products to sell, he could use multi-product pricing technique for substitutable products, but sometimes customers want more than one product, which can be regarded as different kinds of bundles. Compared with naive component pricing, the common method to capture this effect is pure bundling that customers either buy the bundle of all products or nothing.

Suppose customers have i.i.d valuation distribution over n products with all possible valuations larger than cost c, then their purchasing decision comes from a joint distribution, with some known mean $E[s_n]$ and variance $Var[s_n]$. From Chebyshev's inequality, we know

$$P\left(|s_n - E[s_n]| \geq k\sqrt{Var[s_n]}\right) \leq \frac{1}{k^2}$$

if we set the price $p = E[s_n] - k\sqrt{Var[s_n]}$, we have the lowerbound of the profit

$$\pi \geq (E[s_n] - k\sqrt{Var[s_n]} - nc)\left(1 - \frac{1}{k^2}\right)$$

By setting $k = n^{1/6}$, we have

$$RR = \frac{\pi^* - \pi}{n} \leq (1-c)n^{-1/3}$$

so when n is large, the regret rate tends to zero and pure bundling is asympotically optimal.

Bundle Size Pricing has been proposed recently and getting more and more popular

- Customers choose their own bundles
- The price of the bundle depends only on the size of the bundle
- The firm offer m sizes of bundles ($m \leq n$)
- Note the set of all offered sizes as $\mathcal{S}$
- Pure bundling is a special case of bundle size pricing

The benefit of bundle size pricing is two-folded. Intuitively, it's profitable because it can capture customer heterogeneous valuations and reduce the variance through bundling, which is similar as risk pooling effect. Also, it's operationally simple. However, since different customers' valuations, they may choose different bundle size for maximal surplus.

In general, most bundle prices could be determined by a mixed integer program (MIP). We take the mixed bundling as an example

- Product set $\mathcal{N} = \{1, \ldots, n\}$
- There are K customers. Each customer k has a valuation over n products $u_{k1}, \ldots, u_{kn}$
- Bundles are indexed as $S \subseteq \mathcal{N}$
- The firm set a price p_S for all $S \subseteq \mathcal{N}$
- The surplus of the customer of a bundle is $u_{ki} - p_S$
- Given the price, each customer choose a bundle that maximize her surplus
- Cost for bundle $S : c_S = \sum_{i \in S} c_i$

$$\begin{array}{ll}
\max & \sum_{k=1}^{K} \sum_{S \subseteq \mathcal{N}} (p_S - c_S)\, x_{k,S} \\
\text{s.t.} & u_k \leq \sum_{i \in S} u_{ki} - p_S + M\,(1 - x_{k,S}) \\
& u_k \geq \sum u_{ki} - p_S \\
& \sum_{S \subseteq \mathcal{N}} x_{k,S} = 1 \\
& p_S + p_T \geq p_{S \cup T} \\
& x_{k,S} \in \{0, 1\}
\end{array}$$

We can linearize the problem to standard MIP with big-M, but this formulation is not general enough.

Let's consider a more general case that customer valuations over items $\tilde{\boldsymbol{u}} = (\tilde{u}_1, \ldots, \tilde{u}_n)^T$ is jointly distributed as F, and we aim to get the optimal prices for different bundle sizes.

Denote the ith order statistic of $\tilde{\boldsymbol{u}}$ as $\tilde{u}_{(i)}$, i.e., the i th largest value in $\tilde{\boldsymbol{u}}$. Given $\tilde{\boldsymbol{u}}$, the valuation for size j bundle is

$$\tilde{w}_j = \sum_{k=1}^{j} \tilde{u}_{(k)} \quad \forall j \in \mathcal{S}$$

The induced demand for bundle size i is

$$q_i^*(\boldsymbol{p}) = \mathbb{P}_{\tilde{\boldsymbol{W}} \sim G}\left(i = \underset{j \in \{0\} \cup S}{\arg\max} \tilde{w}_j - p_j\right)$$

where G is the joint distribution of $\tilde{\boldsymbol{w}}$.

The firm maximizes its profit

$$\begin{array}{ll} \max_{\boldsymbol{p} \geq 0} & \sum_{i \in \mathcal{S}} (p_i - c_i)\, q_i^*(\boldsymbol{p}) \\ \text{s.t.} & q_i^*(\boldsymbol{p}) = \mathbb{P}_{\tilde{\boldsymbol{w}} \sim G}\left(i = \underset{j \in \{0\} \cup \mathcal{S}}{\arg\max} \tilde{w}_j(\tilde{\boldsymbol{u}}) - p_j\right) \quad \forall i \in \mathcal{S} \end{array}$$

The problem is obviously hard to solve, but we can handle it by CMM. Basically, CMM aims to optimize

$$Z^* = \sup_{\theta \in \Theta} \mathbb{E}_\theta\left(\max_{j \in \{0\} \cup S} (\tilde{w}_j - p_j)\right)$$

with mean and covariance information. The choice probability in CMM is defined as

$$q_j^*(\boldsymbol{p}) = \mathbb{P}_{\theta^*}(j = \arg\max{}_{j \in \{0\} \cup S} (\tilde{w}_j - p_j))$$

which is exactly what we want. Then since we know the equivalence to RAM, we can formulate the choice probability as

$$\boldsymbol{q}^*(\boldsymbol{p}) = \arg\max_{x} \left\{ (\boldsymbol{a} - \boldsymbol{p})^T \boldsymbol{x} - f(x) : \boldsymbol{x} \in \boldsymbol{\Delta}_n^o \right\}$$

where

$$\boldsymbol{\Delta}_n^\circ \triangleq \left\{ x | \boldsymbol{e}^T x \leq 1, x \geq 0 \right\}$$

and

$$f(x) = -\operatorname{trace}\left(\left(\boldsymbol{\Sigma}^{1/2} \boldsymbol{S}(x) \boldsymbol{\Sigma}^{1/2}\right)^{1/2}\right)$$

By arguing that the optimal solution can not be at the boundary, we can characterize the problem as unconstrained convex optimization with optimal condition. Then we can replace $\boldsymbol{p}$ in the original objective function with $\boldsymbol{q}^*(\boldsymbol{p})$ and thus the final objective becomes concave.

References

[1] Michael O Ball and Maurice Queyranne. "Toward robust revenue management: Competitive analysis of online booking". In: *Operations Research* 57.4 (2009), pp. 950–963.

[2] Kalyan Talluri and Garrett Van Ryzin. "Revenue management under a general discrete choice model of consumer behavior". In: *Management Science* 50.1 (2004), pp. 15–33.

[3] Martin Zinkevich. "Online convex programming and generalized infinitesimal gradient ascent". In: *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*. 2003, pp. 928–936.