Columbia B9136 February 25, 2025

Lecture 5 — Assortment Optimization

Lecturer: Will Ma *Scribe: Jingyuan Wang and Yiquan Ji*

1 Choice Models

Why choice models are essential:

- Real-world consumers do not simply choose the most expensive or the cheapest product; their decisions follow behavioral patterns.
- Traditional linear demand models often fail to capture the effect of adding a similar product on the demand for existing products.
- Choice models help retailers and platforms determine optimal pricing and assortment strategies to maximize revenue.

Suppose a retailer hosts $j = 1, \cdots, n$ products each with revenue $r_1 \geq \cdots \geq r_n \geq 0$. The variable r_j denotes the revenue generated if product j is purchased. In the context of assortment optimization, a product with a higher r_j is more profitable; however, its inclusion in the assortment also depends on its attractiveness (w_j) and how it interacts with other products in the offered set.

Given a subset $S \subseteq [n]$[1] of these products, The set $S \subseteq [n]$ represents the collection of products offered to the customer. The decision problem involves choosing an optimal S such that the expected revenue is maximized, considering both r_j and w_j for each product $j \in S$. The customer's choice could be modeled by the *choice function* $\phi(j, S)$, which denotes the probability of the customer choosing product $j \in S$ when S is offered, $\forall S \subseteq [n], j \in S$. Note that $\sum_{j \in S} \phi(j, S) \leq 1$, and with probability $1 - \sum_{j \in S} \phi(j, S)$ (outside choice probability), the customer buys nothing.

2 Multi-Nomial Logit (MNL) Choice Model

- Each product j has a weight $w_j > 0$. This weight reflects the *attractiveness* or *appeal* of the product. In other words, the higher the weight, the more likely the product is to capture the customer's attention and be chosen when it is offered.
- The total weight of an assortment S is defined as:

$$w(S) = \sum_{j \in S} w_j.$$

This aggregate measure represents the cumulative attractiveness of the offered products and is used to normalize the choice probabilities. It ensures that all probabilities (including the outside option) sum to one. The choice function is defined as

$$\phi(j, S) = \frac{w_j}{1 + \underbrace{\sum_{j' \in S} w_j}_{=:w(S)}}, \forall S \subseteq [n], j \in S. \tag{1}$$

[1] We denote $[n] = \{1, \cdots, n\}$.

- An outside option (or the "no purchase" option) is included in the model with a normalized weight $w_0 = 1$. This fixed value provides a baseline level of attractiveness when none of the products in S are chosen. The no-choice probability (choosing "product 0" with weight $w_0 = 1$) is

$$\phi(0, S) = 1 - \sum_{j \in S} \phi(j, S) = \frac{1}{1 + w(S)}.$$

Remark 1. The outside choice weight is normalized to 1 by convention.

2.1 Assortment Optimization under under MNL

Under the MNL choice model, the assortment optimization problem is formulated as

$$\max_{S \subseteq [n]} \quad \sum_{j \in S} r_j \phi(j, S) =: \text{Rev}(S). \tag{2}$$

Writing out $\phi(j, S)$ under MNL, Problem (2) is equivalent to

$$\max_{S \subseteq [n]} \quad \sum_{j \in S} \frac{r_j w_j}{1 + w(S)} =: r^*. \tag{3}$$

To solve Problem (2), we deploy the *"denominator" trick*. Rewrite Problem (3) by subtracting r^* on both sides,

$$\max_{S \subseteq [n]} \quad \left(\sum_{j \in S} \frac{r_j w_j}{1 + w(S)} - r^* \right) = 0$$

$$\max_{S \subseteq [n]} \quad \left(\sum_{j \in S} r_j w_j - r^* (1 + \sum_{j' \in S} w_{j'}) \right) = 0$$

$$\max_{S \subseteq [n]} \quad \sum_{j \in S} (r_j - r^*) w_j = r^* \tag{4}$$

Explanation:

- The term $(r_j - r^*)$ indicates that including any product with $r_j < r^*$ would reduce the overall objective.
- This observation motivates a *threshold-based* (or revenue-ordered) selection rule: only products with $r_j \geq r^*$ should be included.
- In practice, r^* can be found using a binary search method.

To maximize the objective in (4), we need $r_j - r^*$ to be positive, and (4) is equivalent to

$$\sum_{j \in [n] : r_j - r^* \geq 0} (r_j - r^*) w_j = r^*.$$

This implies that the optimal assortment is *"revenue-ordered"*, which means the solution takes the form $\{j : r_j \geq \tau\}$ for some threshold τ. Therefore, the solution is derived by searching over the assortment $\{1\}, \{1, 2\}, \cdots, \{1, 2, \cdots, n\}$, and find S such that $\text{Rev}(S)$ is the highest.

2.2 Constrained Assortment Optimization under under MNL

There are usually two types of constraints:

Capacity Constraint: When the retailer is allowed to display at most k products, the problem becomes

$$\max_{S \subseteq [n]:\, |S| \leq k} \mathrm{Rev}(S).$$

Explanation:

- This constraint limits the size of the assortment.
- Although the unconstrained optimal assortment is revenue-ordered, under capacity constraints the optimal solution may deviate from simple revenue ranking.

Knapsack Constraint: If each product j occupies a space $d_j \in [0, 1]$ and the total available space is 1, the optimization becomes

$$\max_{S \subseteq [n]:\, \sum_{j \in S} s_j \leq 1} \mathrm{Rev}(S).$$

Explanation:

- This constraint reflects physical or display limitations.
- When r^* is unknown, one can use a binary search over r^* while verifying that the assortment satisfies the space constraint.

As before, we denote the constrained optimal value to be r^*. If the optimal value r^* is given, then the LHS of the objective in (4) can be approximately computed for the knapsack constraint.

If r^* is not given, one can deploy binary search for r^*. Start with a r^* guess:

- if LHS of the objective in (4) is greater than the RHS, the r^* guess is too low;
- if LHS of the objective in (4) is smaller than the RHS, the r^* guess is too high.

Adjust the r^* guess accordingly to search for r^*.

Remark 2. The optimal assortment under MNL may no longer be revenue-ordered. Suppose we rank the products in the LHS of (4) with descending order and choose the top k products (under capacity constraint). It could happen that r_j is slightly bigger than r^* but has large enough w_j so that product j is chosen.

2.3 "Red Bus Blue Bus" Paradox

We note that the MNL model has some limitation. Consider the following scenario:

- Car choice with $w_{\text{car}} = 1$;
- Blue Bus choice with $w_{\text{Blue Bus}} = 1$;
- Outside choice with $w_0 = 1$.

In this case, the choice probability for each transportation method is 1/3. If we add another choice of red bus, in reality, people's choice probability would not change: people who drive cars wouldn't switch to red bus just because the color of the bus is red. However, under MNL, the only reasonable thing to do is to model as

- Car choice with $w_{\text{car}} = 1$;
- Blue Bus choice with $w_{\text{Blue Bus}} = 1$;
- Red Bus choice with $w_{\text{Red Bus}} = 1$;
- Outside choice with $w_0 = 1$;

and the choice probability of each transportation method is now 1/4. This means that adding the red bus option not only cannibalize the choice probability of blue bus but also cannibalize the choice probability of other irrelevant alternatives (contradicts the IIA (Independence of Irrelevant Alternatives) assumption). This motivates other choice models.

- **Formal Definition of IIA:** For any two alternatives i and j and any set S that contains them, the ratio of their choice probabilities should remain the same when an unrelated alternative is added or removed from S.
- **Note on the "Red Bus, Blue Bus" Paradox:** The paradox illustrates that the MNL model, by strictly adhering to IIA, produces counter-intuitive effects. For example, adding a red bus (with the same weight as the blue bus) dilutes the probability of choosing the blue bus as well as other alternatives (such as car), even though in reality, consumers who drive cars would not switch to a differently colored bus.
- **Key Point:** IIA does not adequately capture the substitution patterns observed in practice.

3 Nested Logit Choice Model

- **Why Use Nests:** The nested logit model partitions products into groups (nests) so that products within the same nest are more similar and hence exhibit stronger substitution effects. For example, in a retail context, soft drinks may be nested separately from juices.
- **Within-Nest Revenue-Ordering:** Although each nest's products are ordered by revenue (or adjusted revenue), note that the overall solution is obtained by first determining optimal subsets within each nest and then combining them.
- In the nested logit model, products are grouped into disjoint sets called *nests* (denoted by $N_1, N_2, \ldots, N_m$). Products within the same nest are considered similar or substitutable. This grouping allows the model to capture the phenomenon that similar products compete more directly with each other. Partition n products into m "nests" $N_1 \dot{\bigcup} \cdots \dot{\bigcup} N_m = \{1, \cdots, n\}$[2]:

$$N_i \bigcap N_{i'} = \varnothing, \quad \text{if } i \neq i' \text{ and } \bigcup_{i=1}^{m} N_i = \{1, \cdots, n\}.$$

  Products in the same nest are "similar" and cannibalizes each other more.

[2]We use $\dot{\bigcup}$ to denote disjoint union.

- Define "dissimilarity" parameter $\gamma_i \in (0,1]$ for each nest i. For each nest i, the parameter $\gamma_i \in (0,1]$ quantifies the degree of similarity among products within that nest. A lower γ_i implies that the products are very similar (i.e., have strong substitution effects), while $\gamma_i = 1$ indicates that the nest behaves like a collection of independent products.

- Choice function: For a given nest N_i and an assortment S, define:

$$w(S_i) = \sum_{j \in S \cap N_i} w_j.$$

  This sum measures the cumulative attractiveness of the products within nest i that are included in the assortment. For any $S \subseteq [n]$, and any product $j \in S$:

$$\phi(j,S) = \underbrace{\frac{w(S_{i_j})^{\gamma_{i_j}}}{1 + \sum_{i=1}^{m} \underbrace{w(S_i)^{\gamma_i}}_{\text{weight of nest } i}}}_{\text{choice probability of nest } i_j} \cdot \underbrace{\frac{w_j}{w(S_{i_j})}}_{\text{choice probability of item } j \text{ in nest } i_j}, \tag{5}$$

  where we denote $S_i := S \bigcap N_i$ for any nest i and i_j be the unique nest that includes the product $j \in S$.

Remark 3. This Nested Logit Choice Model solved conflicts intrigued by Red Bus Blue Bus Paradox.

Remark 4. We still WLOG, assume $w_0 = 1$ be the weight of the outside choice.

Remark 5. If $\gamma_i = 1$ for all nest i, then the nests don't matter and the model is reduced to MNL.

3.1 Assortment Optimization under Nested Logit

We can solve the assortment optimization problem in this case by extending the denominator trick in Section 2.1. Under the nested logit model, the new optimization problem is

$$\max_{S \subseteq [n]} \quad \sum_{i=1}^{m} \frac{w(S_i)^{\gamma_i}}{1 + \sum_{i'=1}^{m} w(S_{i'})^{\gamma_{i'}}} \cdot \underbrace{\sum_{j \in S_i} r_j \frac{w_j}{w(S_i)}}_{=: V_i(S_i)} = r^*$$

$$\max_{S \subseteq [n]} \quad \sum_{i=1}^{m} w(S_i)^{\gamma_i} (V_i(S_i) - r^*) = r^*$$

$$\sum_{i=1}^{m} \max_{S_i \subseteq N_i} \underbrace{w(S_i)^{\gamma_i} (V_i(S_i) - r^*)}_{(*)} = r^*,$$

where $V_i(S_i)$ represents the expected revenue conditional on choosing nest i with items S_i. Note that the above derivative implies that the optimization problem is equivalent to solving $(*)$ for each $N_i, i \in [m]$.

Formally, our goal is to find small number of assortments $\mathscr{S}_i$ in N_i such that for any r^*, an optimal solution to $(*)$ lies in $\mathscr{S}_i$. In the following, we will show that the solution is "revenue-ordered", i.e., has the form $\{\{j \in N_i : r_j \geq \tau\} \mid \tau \geq 0\}$.

For simplicity, we omit the subscript i, and let $N = \{1, \cdots, |N|\}$ and let $r_1 \geq \cdots \geq r_{|N|}$. Consider the fractional relaxation,

$$\max_{S \subseteq N} w(S)^\gamma \left(\sum_{j \in S} \frac{r_j w_j}{w(S)} - r^* \right) = \max_{S \subseteq N} \frac{\sum_{j \in S}(r_j - r^*) w_j}{(\sum_{j \in S} w_j)^{1-\gamma}} \leq \max_{\mathbf{x} \in [0,1]^{|N|}} \frac{\sum_{j \in S}(r_j - r^*) w_j x_j}{(\sum_{j \in S} w_j x_j)^{1-\gamma}}.$$

In the fractional problem, the optimal solution takes the form $x_1 = 1, x_2 = 1, \cdots, x_j \in (0,1], x_{j+1} = 0, \cdots, x_{|N|} = 0$, which takes the revenue-order form. Figure 1 provides the intuition with a simple example (considering the denominator as fixed).

We now further shows that the optimal cannot be fractional, i.e., we are better off to include all of x_j if we ever include $x_j \in (0,1]$ at all. Suppose we have increased $x_1, \cdots, x_j$ to 1. The current numerator is $\sum_{j'=1}^{j}(r_{j'} - r^*) w_{j'}$ and the denominator is $(\sum_{j'=1}^{j} w_{j'})^{1-\gamma}$. Then the derivative of increasing x_{j+1} is

$$\Bigg\{ (r_{j+1} - r^*) w_{j+1} \Big(\underbrace{\sum_{j'=1}^{j} w_{j'} + w_{j+1} x_{j+1}}_{=:B} \Big)^{1-\gamma}$$
$$- \Big(\sum_{j'=1}^{j} (r_{j'} - r^*) w_{j'} + (r_{j+1} - r^*) w_{j+1} x_{j+1} \Big) (1-\gamma) \Big(\underbrace{\sum_{j'=1}^{j} w_{j'} + w_{j+1} x_{j+1}}_{=:B} \Big)^{-\gamma} w_{j+1} \Bigg\} / D^2,$$

for some variable $D = (\sum_{j'=1}^{j} w_{j'})^{1-\gamma}$ on the denominator. Since $r_1 \geq \cdots \geq r_{|N|}$, there exists some $C \geq 0$ such that

$$\Big(\sum_{j'=1}^{j} (r_{j'} - r^*) w_{j'} + (r_{j+1} - r^*) w_{j+1} x_{j+1} \Big) = (r_{j+1} - r^*) B + C,$$

and the above derivative can be further simplified to

$$\begin{aligned} & \big((r_{j+1} - r^*) w_{j+1} B^{1-\gamma} - (C + (r_{j+1} - r^*) B)(1-\gamma) B^{-\gamma} w_{j+1} \big) / D^2 \\ = & \big(- C(1-\gamma) B^{-\gamma} w_{j+1} + (r_{j+1} - r^*) w_{j+1} B^{1-\gamma} \gamma \big) / D^2. \end{aligned}$$

Note that B is increasing as x_{j+1} is increasing and C is decreasing as x_{j+1} is increasing, which implies that the numerator of the above derivative is increasing if x_{j+1} is increasing. And the denominator D is always positive. Hence, if the derivative is positive initially (which means that the numerator is positive initially), then it stays positive. Hence, if it is beneficial to set $x_j > 0$, then it is better to set $x_j = 1$, eliminating that possibility. This shows that the optimal solution is integer, i.e. corresponds to a subset of products.

This concludes the argument.

If r^* is unknown, we can solve for the following LP:

$$\begin{aligned} \min \quad & r^* \\ \text{s.t.} \quad & r^* \geq \sum_{i=1}^{m} y_i \\ & y_i \geq w(S_i)^{\gamma_i} (V_i(S_i) - r^*) \quad \forall i \in [m], S_i \in \mathscr{S}_i \\ & r^*, y_i \geq 0 \quad \forall i \in [m]. \end{aligned}$$

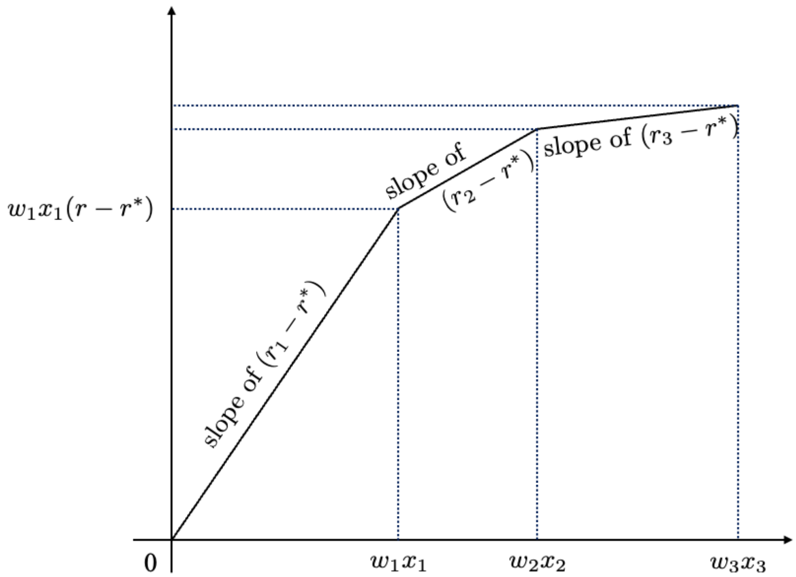

Figure 1: The fractional problem solution intuition.

4 Alternative Interpretations of Choice Model

4.1 Random Utility Model

Each customer draws random utility U_j for all $j \in [n]\bigcup\{0\}$ and then selects $\text{argmax}_{j\in S\bigcup\{0\}}U_j$. Within the random utility framework, u_j is the deterministic component of the utility associated with product j. It is typically modeled as a function of observable features or attributes of the product. The term ϵ_j captures the random variation in a customer's utility for product j. In the MNL model, these random terms are assumed to follow a Gumbel distribution, which leads to the closed-form expression for choice probabilities.

Under the MNL model, this is equivalent to setting

$$U_j = u_j + \varepsilon_j, \forall j \in [n]\bigcup\{0\}, \text{ where iid } \varepsilon_j \sim \text{Gumbel}(0,1) \text{ and } u_0 = 0.$$

This choice function is

$$\phi(j,S) = \mathbb{P}_{\varepsilon_0,\varepsilon_1,\cdots,\varepsilon_n}\big(j = \text{argmax}_{j'\in S\bigcup\{0\}}U_{j'}\big).$$

Explanation:

- u_j is the deterministic component of the utility.
- ϵ_j is a random error term that is independently and identically distributed (iid) according to the Gumbel distribution.
- This framework leads directly to the MNL formulation when computing choice probabilities.

4.2 Ranked List Model

In the ranked list model, each customer generates a random ordered list ℓ of products based on their preferences. The order reflects the customer's ranking, with the first product being the most preferred. The choice is then determined by which product appears highest in the list among those offered. Customer has random list ℓ drawn from some known distribution over ordered subsets of $[n]$. The first product in the list is most preferred; the second product is the second choice, etc. Everything not in the list is less preferred compared to the outside option. The choice function is

$$\phi(j, S) = \mathbb{P}_\ell\Big(j \text{ is the most preferred among } S \bigcup \{0\}, \text{ according to } \ell\Big).$$

We also introduce the *Plackett-Luce* Plackett (1975) model. Suppose there are three horses in a competition, each with score $w_1 = 1, w_2 = 2, w_3 = 3$. Then

$$\begin{aligned}\mathbb{P}(\text{order is 321}) =& \mathbb{P}(\text{3 comes first}) \cdot \mathbb{P}(\text{2 beats 1} \mid \text{3 comes first}) \\ =& \frac{3}{1+2+3} \cdot \mathbb{P}(\text{2 beats 1}) = \frac{3}{1+2+3} \cdot \frac{2}{1+2}.\end{aligned}$$

Similarly,

$$\mathbb{P}(\text{order is 132}) = \frac{1}{1+2+3} \cdot \frac{3}{2+3}.$$

Now suppose the list ℓ is generated via Plackett-Luce and the customer chooses the product that is most preferred according to the list, then this choice model recovers MNL.

Explanation:

- The Plackett-Luce model is a specific case where the list is generated sequentially based on the product weights.
- The resulting choice probabilities coincide with those given by the MNL model.

4.3 Choice Models Summary

A summary of choice models are given in Figure 2.

For an example of choice paralysis, consider the following. Let $n = 2$ and the choice function

$$\begin{aligned}\phi(1, \{1\}) =& 1, \\ \phi(2, \{2\}) =& 1 \\ \phi(j, \{1, 2\}) =& 0, \forall j = 1, 2.\end{aligned}$$

Note that in this case, increasing S can increase the chances of the customer buying nothing. Therefore, it is not a random utility/random list model.

For an example of decoy, consider a burger place that offers a burger for \$5 and a set of burger with fries for \$10. If we add a decoy: a burger with a single fry for \$9.5, then the burger/fries set would appear more attractive to customers.

5 Markov Chain Induced Choice Model

Define a Markov chain on n products with:

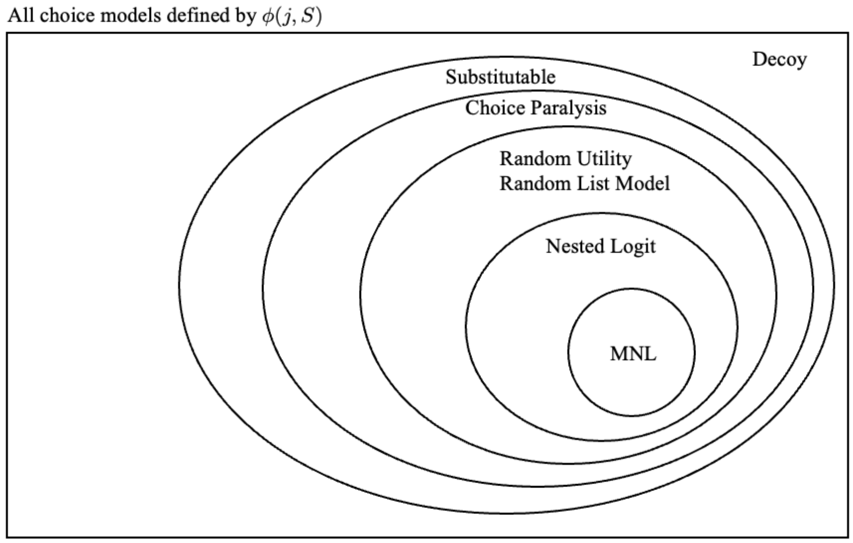

Figure 2: Summary of Choice Models

- Transition probabilities p_{ij} for all $i, j \in [n]$:

$$\sum_{j=1}^{n} p_{ij} \leq 1 \quad \forall i \in [n]$$

$$\text{with probability } 1 - \sum_{j=1}^{n} p_{ij}, \text{ transition to 0 and ends}$$

  In this model, p_{ij} denotes the probability of transitioning from product i to product j. These probabilities capture how likely a customer is to move from one product to another during their consideration process. The condition

$$\sum_{j=1}^{n} p_{ij} \leq 1$$

  ensures that there is a nonzero probability of transitioning to the outside option.

- Probability of starting on product j: λ_j:

$$\sum_{j=1}^{n} \lambda_j \leq 1$$

$$\text{with probability } 1 - \sum_{j=1}^{n} \lambda_j, \text{ start at 0}$$

  The variable λ_j represents the probability that a customer initially considers product j. This can be seen as the product's initial attractiveness or visibility in the decision process. The sum of all initial probabilities is bounded by 1, with the remainder allocated to starting with the outside option.

- List ℓ generated by

1. hop in Markov chain with probability $\{\lambda_j\}_{j=1,\cdots,n}$
2. whenever visit a node for the first time, add it to the end of the list ℓ
3. terminate when transition to 0 (assuming this happens with probability 1).

Then MNL is captured by this Markov chain, given weights $w_1, \cdots, w_n > 0$. Consider the following example, with $n = 3$. The λ_j's are $\frac{w_j}{1+\sum_{i=1}^{3} w_i}$; and the transition probability to some node j from the node i should be proportional to the weights w_j: $p_{ij} = \frac{w_j}{\Delta}$ for some normalizer $\Delta > 0$.

Explanation:

1. A customer first lands on a product j with probability λ_j.
2. The customer then follows the Markov chain. Each time a product is visited for the first time, it is added to the consideration list.
3. The process terminates when the chain transitions to the outside option.

Remark 6. This construction is more general and also gets round the "red bus blue bus" paradox by setting the transition probability between red bus and blue bus to be near to 1, that is, if we are on red bus node now, then with a high probability, the next thing on the list is the blue bus.

5.1 Assortment Optimization under the Markov Chain Induced Choice Model

We note that for random utility models, the following holds:

$$\begin{aligned}
\max_{S\subseteq[n]} \text{Rev}(S) &= \max_{S\subseteq\{2,\cdots,n\}} \text{Rev}(\{1\}\bigcup S) \quad (r_1 \geq \cdots \geq r_n) \\
&= \max_{S\subseteq\{2,\cdots,n\}} \big(r_1 \cdot \phi(1,\{1\}\bigcup S) + \sum_{j\in S} r_j \cdot \phi(j,\{1\}\bigcup S)\big) \\
&= \max_{S\subseteq\{2,\cdots,n\}} r_1\big(\phi(1,\{1\}) - \sum_{j\in S} \mathbb{P}\big((j \succeq \{0,1\})\bigcup(S\bigcap 1 \succ 0)\big) + \sum_{j\in S} r_j\phi(j,\{1\}\bigcup S)\big) \\
&= \max_{S\subseteq\{2,\cdots,n\}} r_1\big(\phi(1,\{1\}) - \sum_{j\in S} \underbrace{\mathbb{P}\big(1 \succ 0 \mid j \succeq S\bigcup\{0,1\}\big)}_{=\mathbb{P}(\text{ hit 1 before 0}|\text{starting at } j)} \phi(j,\{1\}\bigcup S) + \sum_{j\in S} r_j\phi(j,\{1\}\bigcup S)\big) \\
&= \max_{S\subseteq\{2,\cdots,n\}} \big(r_1\phi(1,\{1\}) + \sum_{j\in S} \underbrace{(r_j - r_1 \cdot \mathbb{P}(1 \succ 0 \mid j \succ \{1,0\}))}_{\text{externality adjusted revenue, doesn't depend on } S} \cdot\phi(j,\{1\}\bigcup S)\big)
\end{aligned}$$

This derivation reduces the problem to the assortment optimization on Markov chain with 1 fewer nodes and externality adjusted revenue.

- **Externality-Adjusted Revenue:** Set one node as a *sink* and then adjust the revenue of the other nodes to account for the possibility of being preempted by the sink. This involves:
  - Locking in the sink product.
  - Recalculating the revenue of other nodes by subtracting the effect of switching to the sink.

To better illustrate this intuition, consider the example with the Markov chain in Figure 3. We first let node 1 as the sink:

$$r_1 \cdot \phi(1,\{1\}) = 6 \cdot \left(\frac{3}{4} \cdot \frac{2}{3} \cdot \frac{1}{2}\right) = 1.5.$$

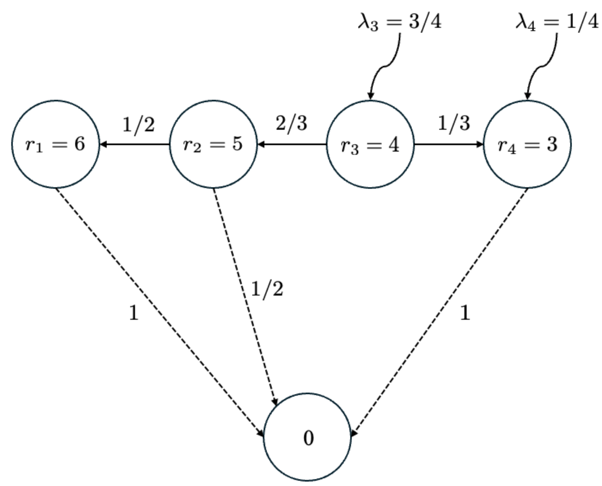

Figure 3: Markov chain induced choice model example

Then the externality adjusted revenue for other nodes are

$$\begin{aligned} r_2^{(1)} =& 5 - 6 \cdot \frac{1}{2} = 2 \\ r_3^{(1)} =& 4 - 6 \cdot \frac{2}{3} \cdot \frac{1}{2} = 2 \\ r_4^{(1)} =& 3 - 6 \cdot 0 = 3. \end{aligned}$$

Therefore the next node to add to the assortment is node 4, and

$$r_4^{(1)} \cdot \phi(4, \{1, 4\}) = 3 \cdot \left(\frac{1}{4} + \frac{1}{4} \right) = 1.5.$$

For the next round, we treat node 4 as the new sink, and the externality adjusted revenue for other nodes are

$$\begin{aligned} r_2^{(2)} =& r_2^{(1)} - 3 \cdot 0 = 2 - 0 = 2 \\ r_3^{(2)} =& r_3^{(1)} - 3 \cdot \frac{1}{3} = 2 - 1 = 1. \end{aligned}$$

Therefore we add node 2 to the assortment:

$$r_2^{(2)} \cdot \phi(2, \{4, 2\}) = 2 \cdot \frac{1}{2} = 1.$$

The optimal solution is

$$\max_S \text{Rev}(S) = \text{Rev}(\{1, 4, 2\}) = 1.5 + 1.5 + 1 = 4.$$

6 Concluding Remarks

1. Given data of agents $t = 1, \cdots, T$, each chooses j_t from the displayed set $S_t \subseteq [n]$. The log likelihood of any sample path $\{(j_t, S_t)\}_{t=1}^T$ is

$$\log \prod_{t=1}^{T} \frac{e^{u_{j_t}}}{1 + \sum_{j \in S_t} e^{u_{j_t}}} = \sum_{t=1}^{T} u_{j_t} - \underbrace{\log \left(\sum_{j \in S_t \bigcup \{u_0\}} e^{u_j} \right)}_{f(\{S_t\}_{t=1}^T)},$$

with $u_0 = 0$. To maximize the choice probability, note that f is convex and thus the above can be efficiently solved via concave maximization.

Explanation:

- Maximizing this log-likelihood via methods such as Maximum Likelihood Estimation (MLE) yields the parameter estimates for the choice model.

2. *Heterogeneous agents* model: the utility of agent t with product j is

$$U_{tj} = \underbrace{u_{tj}}_{=\beta^\top x_{tj}} + \varepsilon_{tj},$$

where x_{tj} is the observed features. The algorithms are interested in estimating β.

3. *Mixed-MNLs*: if the agents are heterogeneous but the algorithm is forced to provide the same assortment to every agent, then the objective is to

$$\max_{S \in [n]} \frac{1}{T} \sum_{t=1}^{T} \sum_{j \in S} r_j \cdot \frac{e^{\beta^\top x_{tj}}}{1 + \sum_{j' \in S} e^{\beta^\top x_{tj'}}}.$$

- **Mixed-MNL Complexity:** Note that when agents are heterogeneous (with individual parameter vectors β_t), the optimal assortment problem becomes NP-hard. This is due to the inability to use a unified formulation for all consumers.

Bibliographical notes. The proofs for MNL assortment optimization presented here are motivated by Rusmevichientong et al. (2010). The proof for Nested Logit assortment optimization is adapted from Gallego and Topaloglu (2014). The proof for Markov Chain assortment optimization based on externality adjustment comes from Désir et al. (2020).

References

A. Désir, V. Goyal, D. Segev, and C. Ye. Constrained assortment optimization under the markov chain–based choice model. *Management Science*, 66(2):698–721, 2020.

G. Gallego and H. Topaloglu. Constrained assortment optimization for the nested logit model. *Management Science*, 60(10):2583–2601, 2014.

R. L. Plackett. The analysis of permutations. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 24(2):193–202, 1975.

P. Rusmevichientong, Z.-J. M. Shen, and D. B. Shmoys. Dynamic assortment optimization with a multinomial logit choice model and capacity constraint. *Operations research*, 58(6):1666–1680, 2010.