ORIE 4154: Pricing and Market Design **Fall 2026**

Lecture 17 — Assortment Optimization: Revenue Ordering and Beyond

Lecturer: Sid Banerjee *Date: October 22, 2026 (draft)*

17.1 More options can mean less revenue

Lecture 16 asked how a customer chooses from an offer set S. Now the seller chooses that set. The first surprise is that offering another product may *reduce* expected revenue, even if it attracts some additional purchases.

Consider three products with unit revenues r_i and MNL attraction weights v_i, plus a no-purchase weight $v_0 = 4$:

	1	2	3
r_i	100	60	20
v_i	2	6	8

The MNL model from Lecture 16 gives

$$R(\{1\}) = \frac{100 \cdot 2}{4+2} = \frac{200}{6} \approx 33.33,$$
$$R(\{1,2\}) = \frac{100 \cdot 2 + 60 \cdot 6}{4+2+6} = \frac{560}{12} \approx 46.67,$$
$$R(\{1,2,3\}) = \frac{100 \cdot 2 + 60 \cdot 6 + 20 \cdot 8}{4+2+6+8} = \frac{720}{20} = 36.$$

Adding product 2 helps, but then adding the attractive, low-revenue product 3 hurts. Some customers who would have bought a more expensive product switch to 3. This is *cannibalization*; whether a larger set is better depends on the entire substitution pattern.

17.2 The seller's optimization problem

Let $N = \{1, \ldots, n\}$, and let $q_i(S)$ be the probability that a customer chooses product i from S, with 0 representing no purchase. A sale of i earns $r_i \geq 0$. The expected revenue from one customer is

$$R(S) = \sum_{i \in S} r_i q_i(S), \qquad R^* = \max_{S \subseteq N} R(S). \tag{17.1}$$

We assume a known choice model and a fixed offered set for this one-customer decision. With arbitrary $q_i(S)$, there are 2^n possible sets to evaluate. Structure in the choice model can drastically reduce that search.

Under MNL, with $v_0 > 0$ and $v_i > 0$,

$$R(S) = \frac{\sum_{i \in S} r_i v_i}{v_0 + \sum_{i \in S} v_i}. \tag{17.2}$$

The numerator is attraction-weighted product revenue. The denominator includes the attraction of *all* offered choices and the outside option. Adding product j changes both quantities, not just the revenue associated with j.

17.3 The marginal value of adding a product

Write $A = \sum_{i \in S} r_i v_i$ and $D = v_0 + \sum_{i \in S} v_i$. Then $R(S) = A/D$ and

$$\begin{aligned} R(S \cup \{j\}) - R(S) &= \frac{A + r_j v_j}{D + v_j} - \frac{A}{D} \\ &= \frac{v_j (D r_j - A)}{D(D + v_j)} \\ &= \boxed{\frac{v_j}{D + v_j}\big(r_j - R(S)\big)}, \qquad j \notin S. \end{aligned} \tag{17.3}$$

Thus an available product with positive attraction helps exactly when its own unit revenue exceeds the *current expected revenue per customer* $R(S)$. At equality it changes neither expected revenue nor the optimum value. It can still change who buys what.

In the opening example, $r_2 = 60 > R(\{1\}) \approx 33.33$, so product 2 helps. But $r_3 = 20 < R(\{1, 2\}) \approx 46.67$, so product 3 hurts. The effect is about diversion as well as incremental sales.

17.4 An exact theorem under MNL

Sort products so that $r_1 \geq r_2 \geq \cdots \geq r_n \geq 0$, and define the *revenue-ordered prefixes*

$$S_0 = \varnothing, \qquad S_k = \{1, \ldots, k\} \quad (k = 1, \ldots, n).$$

Theorem 17.1 (Revenue ordering under MNL). *With $v_0 > 0$, positive product attraction weights, and nonnegative product revenues, at least one revenue-ordered prefix maximizes (17.2). Ties may create other optimal assortments.*

Proof: Let $R^* = \max_S R(S)$. For *every* S, the definition of R^* and positivity of the denominator give

$$\begin{aligned} R(S) \le R^* \quad &\iff \quad \sum_{i\in S} r_i v_i \le R^*\Big(v_0 + \sum_{i\in S} v_i\Big) \\ &\iff \quad \sum_{i\in S} v_i(r_i - R^*) \le v_0 R^*. \end{aligned} \tag{17.4}$$

For an optimal set S^*, equality holds. If $j \notin S^*$ has $r_j > R^*$, adding j makes the left side of (17.4) strictly larger than $v_0 R^*$, contradicting the inequality for all S. Thus every product with $r_j > R^*$ is in S^*.

Likewise, if $j \in S^*$ has $r_j < R^*$, removing its *negative* term makes the left side strictly larger than $v_0 R^*$, again a contradiction. Thus no product with $r_j < R^*$ is in S^*. Products with $r_j = R^*$ contribute zero and may be added or removed without changing the equality or revenue. Include all of them, if desired. The resulting optimal set is exactly a prefix of the products sorted by revenue. ■

The threshold here is the *optimal expected revenue* R^*, not an externally chosen unit price. The proof does not need to know R^* in advance; it establishes a structural fact that lets us search over prefixes.

17.5 From exponentially many sets to a short list

After sorting, maintain two prefix sums:

$$V_k = \sum_{i=1}^{k} v_i, \qquad W_k = \sum_{i=1}^{k} r_i v_i, \qquad V_0 = W_0 = 0.$$

Then $R(S_k) = W_k/(v_0 + V_k)$. A complete algorithm is:

1. Sort products by decreasing r_i.
2. Starting from $V_0 = W_0 = 0$, update $V_k = V_{k-1} + v_k$ and $W_k = W_{k-1} + r_k v_k$.
3. Return the prefix with the largest $W_k/(v_0 + V_k)$, including $k = 0$.

Sorting costs $O(n \log n)$ and all prefix evaluations cost $O(n)$, so the overall time is $O(n \log n)$. In the opening example the four candidates have revenues 0, 33.33, 46.67, and 36; the optimum is $\{1, 2\}$.

17.6 Does revenue ordering survive beyond MNL?

Not exactly. Consider three products with $r_1 = 100$, $r_2 = 90$, and $r_3 = 10$. Half the customers have a fixed preference ranking

$$2 \succ 1 \succ 0 \succ 3,$$

and half have

$$3 \succ 0 \succ 2 \succ 1.$$

Each customer chooses the highest-ranked offered alternative, so this is a random-utility choice model and satisfies regularity. Yet

$$R(\{1,3\}) = \tfrac{1}{2}(100) + \tfrac{1}{2}(10) = 55,$$

while the revenue-ordered prefixes yield

$$R(\{1\}) = 50, \qquad R(\{1,2\}) = 45, \qquad R(\{1,2,3\}) = 50.$$

Checking the remaining sets gives revenues at most 50, so the non-prefix set $\{1,3\}$ is optimal. Product 2 mostly cannibalizes product 1; product 3 reaches a different group. Exact revenue ordering depends on MNL's specific substitution structure.

17.7 A guarantee under regularity alone

Berbeglia and Joret ask a weaker question: even when revenue ordering is not *exact*, how much revenue can its best prefix guarantee? We now assume only the basic choice probabilities (17.1) and *regularity* from Lecture 16; the model need not be MNL or even random utility.

Let $a_1 > \cdots > a_k > 0$ be the k distinct positive product-revenue levels. For each level a, define its revenue-ordered assortment

$$A_a = \{i \in N : r_i \geq a\}, \qquad R_{\mathrm{RO}} = \max_{a \in \{a_1,\ldots,a_k\}} R(A_a).$$

Products with zero revenue need not be included in these candidate sets. Under regularity, adding them cannot improve the revenue of positive-revenue products.

Lemma 17.2 (A useful threshold inequality). *Let S^* be any assortment, and let X be the unit revenue of the product chosen under S^*, with $X = 0$ on no purchase. For each positive revenue level a,*

$$R(A_a) \geq a\,\mathbb{P}_{S^*}(X \geq a). \tag{17.5}$$

Proof: Set $B = S^* \cap A_a$. First compare the *nested* sets $B \subseteq S^*$. For each $i \in B$, regularity gives $q_i(B) \geq q_i(S^*)$, so

$$1 - q_0(B) = \sum_{i \in B} q_i(B) \geq \sum_{i \in B} q_i(S^*) = \mathbb{P}_{S^*}(X \geq a).$$

Next compare the other nested pair $B \subseteq A_a$. Regularity for the outside option gives $q_0(A_a) \le q_0(B)$, hence

$$1 - q_0(A_a) \ge 1 - q_0(B) \ge \mathbb{P}_{S^*}(X \ge a).$$

Every product in A_a earns at least a, so $R(A_a) \ge a(1 - q_0(A_a))$. Combining the inequalities proves the claim. Comparing S^* directly with A_a would be invalid: they need not be nested. ■

Theorem 17.3 (Berbeglia–Joret $1/k$ guarantee). *For any regular choice model with nonnegative product revenues and $k \ge 1$ distinct positive revenue levels,*

$$R_{\mathrm{RO}} \ge \frac{R^*}{k}. \tag{17.6}$$

Proof: Take an optimal S^* and partition its expected revenue by revenue level:

$$R^* = \sum_{\ell=1}^{k} a_\ell \sum_{i \in S^* : r_i = a_\ell} q_i(S^*).$$

For each ℓ, the corresponding summand is at most

$$a_\ell \, \mathbb{P}_{S^*}(X \ge a_\ell) \le R(A_{a_\ell}) \le R_{\mathrm{RO}},$$

where the middle inequality is (17.5). There are k summands, so $R^* \le k R_{\mathrm{RO}}$. ■

The distinction between k and n matters: k counts *distinct positive revenues*, not products. If many products share only two fare levels, the guarantee is $1/2$, however many products there are. The theorem does not say revenue ordering is always optimal; the ranking example above shows otherwise.

17.8 When the assortment itself is constrained

Suppose shelf space or a display interface permits at most b products. Even under MNL, the best *feasible* set need not be a revenue prefix: a prefix might be too large, and omitting a high-revenue but weakly attractive item may make room for a better combination.

The fractional MNL objective still has useful structure. For a proposed revenue level z, (17.2) gives

$$R(S) \ge z \quad \Longleftrightarrow \quad \sum_{i \in S} v_i (r_i - z) \ge v_0 z. \tag{17.7}$$

Under $|S| \le b$, the left side is maximized by selecting up to b products with the largest *positive* adjusted contributions $v_i(r_i - z)$. Thus we can test a candidate z efficiently and search over z. This is a preview, not a complete exact-arithmetic algorithm; the detailed constrained-assortment problem belongs in a later treatment.

17.9 What the two lectures have established

1. MNL's fractional revenue formula makes the effect of adding a product transparent: compare its unit revenue with the current expected revenue.
2. Under MNL, an optimal assortment is revenue ordered, reducing an exponential search to $n + 1$ candidates.
3. More expressive choice models can invalidate exact revenue ordering. Under regularity alone, however, the best revenue-ordered assortment earns at least R^*/k when there are k positive revenue levels.

Later connections include learning MNL attraction weights while offering assortments (MNL bandits), choosing assortments while preserving inventory (choice-based revenue management), and seeking robust guarantees when a precise parametric choice model is not credible. These are new questions; neither the exact MNL theorem nor the regularity approximation automatically solves them.

Recommended reading

Exact MNL result.
Banerjee, ORIE 6154 Assortment Optimization notes, for the transformed-objective proof and further examples. Talluri and van Ryzin, *The Theory and Practice of Revenue Management*, Section 7.2.2, supplies choice-model background.

General-choice guarantee.
Berbeglia and Joret, "Assortment Optimisation Under a General Discrete Choice Model: A Tight Analysis of Revenue-Ordered Assortments", Section 3. The $1/k$ guarantee here is adapted from that paper; its stronger revenue-range bound and broader results are optional.

Optional constrained extension.
Banerjee, ORIE 6154 *Constrained Assortment Optimization* notes (also in the course archive). The cardinality discussion here is only a preview.