# Revenue Management with Product Retirement and Customer Selection

Adam N. Elmachtoub[1], Vineet Goyal[1], Roger Lederman[2], and Harsh Sheth[1]

[1] Department of Industrial Engineering and Operations Research and Data Science Institute, Columbia University, New York, NY 10027
{ae2516,vg2277,hts2112}@columbia.edu
[2] Amazon, 440 Terry Ave N, Seattle, WA 98109 rllederm@amazon.com**

In the classical setting of dynamic assortment optimization, the seller has the flexibility to fully personalize assortments for every arriving customer. Although powerful for online marketplaces, such dynamic policies are not suitable or even feasible for many other settings such as brick-and-mortar stores (e.g., fast-fashion, department stores), electronics marketplaces (e.g., Best Buy), and business-to-business (B2B) settings. In the former case, all customers entering the store see the same products; and for the latter two, available products offered by the sellers are common knowledge and can often be obtained from the firm's webpage. Consequently, in all cases, assortment personalization is not an option and simpler assortment policies are required for revenue maximization. Motivated by this real-world context, we investigate and prove performance guarantees for a particular class of simpler policies, which allow assortment optimization to be more widely applicable as a tool for revenue maximization. In particular, we study two novel machinery that are implementable for the aforementioned settings: *product retirement* and *customer selection*.

Product retirements are a simple way to control the set of available products, where once a product is retired, it becomes unavailable for purchase for all future customers. Sellers are required to make such retirement-type decisions in many natural settings, such as when to move a product to clearance, or when to discontinue an order model or generation of a product. Optimally retiring products helps balance the trade-off between loss of total sales, and cannibalization of sales of newer or better products by a subpar product. In B2B and similar settings, product retirements can be coupled with customer selection to enhance their impact. In such settings, sellers select potential customers every day and sales representatives then offer these customers all the available products via a sales call. Thus, the seller can jointly make decisions on when to retire older generations and which customers to approach each day to balance the key trade-off more effectively.

Formally, we study the problem where the seller has a fixed inventory of multiple substitutable products to be sold over a finite time horizon, and customers arriving sequentially choose from the available products according to a *multinomial logit* (MNL) choice model. All our policies are designed with the help of a fluid approximation linear program (LP). Our main contributions are the following:

1. For the case of homogeneous customers, we design a policy which defines a static retirement time to retire a product unless its inventory was depleted earlier. We show that our policy gives a $(1-\epsilon)$-approximation, where $\epsilon$ goes to zero as the inventory level and time horizon go to infinity.
2. For the case of heterogeneous customers, we design a policy that determines the customer sequence in addition to the product retirement timeline. The policy is asymptotically optimal when there is an optimal solution to the LP upper bound with a particular nested structure. We show this nested structure always exists and can be found efficiently when there are two products.
3. However, the nested structure needed for our previous result is not guaranteed when there are more than two products. We give an alternative policy that achieves at least $1/4-o(1)$ of the expected revenue of the optimal dynamic policy. This algorithm determines a static sequence of the customers and assortments by constructing an approximate greedy solution to the fluid approximation LP.
4. In numerical experiments, we observe that our proposed policies outperform benchmarks under most scenarios. Moreover, we see that the particular nested structure required for an asymptotically optimal policy can be found in an overwhelming majority of cases, and for these cases, our policy results in the best performance.

Full paper: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4033922

** Work of this author was conducted while employed by IBM Research.