

Randomized Assortment Optimization

When a firm selects an assortment of products to offer to customers, it uses a choice model to anticipate their probability of purchasing each product. In practice, the estimation of these models is subject to statistical errors, which may lead to significantly suboptimal assortment decisions. Recent work has addressed this issue using robust optimization, where the true parameter values are assumed unknown and the firm chooses an assortment that maximizes its worst-case expected revenues over an uncertainty set of likely parameter values, thus mitigating estimation errors. In this paper, we introduce the concept of randomization into the robust assortment optimization literature. We show that the standard approach of deterministically selecting a single assortment to offer is not always optimal in the robust assortment optimization problem. Instead, the firm can improve its worst-case expected revenues by selecting an assortment randomly according to a prudently designed probability distribution. We demonstrate this potential benefit of randomization both theoretically in an abstract problem formulation as well as empirically across three popular choice models: the multinomial logit model, the Markov chain model, and the preference ranking model. We show how an optimal randomization strategy can be determined exactly and heuristically. Besides the superior in-sample performance of randomized assortments, we demonstrate improved out-of-sample performance in a data-driven setting that combines estimation with optimization. Our results suggest that more general versions of the assortment optimization problem—incorporating business constraints, more flexible choice models and/or more general uncertainty sets—tend to be more receptive to the benefits of randomization.