

Big, Small, and Small+ Data-Driven OM

Ming Hu

Professor Joseph L. Rotman School of Management, University of Toronto

Abstract: We discuss three data-driven decision-making scenarios in classical operations management (inventory or pricing) settings, using big or small data or small data but with one chance of experimentation. First, with big data, we study a contextual-based newsvendor problem using deep neural networks (DNN). Empirical process theory is pivotal in ensuring that the asymptotic behavior of observed data converges to the true underlying distribution as the sample size increases. We provide theoretical guarantees in terms of excess risk bounds for the DNN solution, characterized by the network structure and sample size, validating the applicability of DNNs in relevant OM contexts. These excess risk bounds exhibit polynomial growth in the feature dimension and attain the minimax convergence rate (with respect to the sample size) in expectation. Second, with small data, traditional frequentist methods may be ineffective, and we propose using the empirical Bayes (EB) method to achieve transfer learning in estimating unknown parameters using data across many products and subsequently making decisions based on these estimates. We illustrate this approach with a multi-product pricing problem, employing a hierarchical feature-based demand model and a nonparametric maximum likelihood method to derive the prior from the data. The effectiveness of the EB method is demonstrated by characterizing the regret bound using an oracle benchmark that presumes prior knowledge of the underlying distribution. Third, we study the benefits of one-shot price experimentation for a seller in setting a price who only knows the exact purchase probability associated with a single historical price and aims to maximize the worst-case revenue ratio compared to an oracle with complete knowledge of the value distribution. We analytically characterize the optimal distributionally robust experimental and final price points, obtain their tight performance guarantee for any historical purchase probability, and then evaluate the value of experimentation, which exhibits a two-modal behavior with respect to the historical purchase probability.