Non-Parametric Data-Driven Policies for Stochastic Inventory Control Models

We consider the classical newsvendor problem and its multi-period extension, but under the assumption that the explicit demand distributions are not known and the only information available is a set of independent (possibly censored) random samples drawn from the true distributions. Under the assumption that the demand distributions are given explicitly, these models are well-studied and usually are relatively straightforward to solve. However, in most real-life scenarios, the true demand distributions are not available or they are too complex to work with. We shall describe several sampling-based policies, that is, policies that are computed based only on observed demand samples and without any access to, or assumptions on, the true demand distributions. These policies are computationally attractive and have strong theoretical properties:

(i) For the model with uncensored demand samples, we obtain bounds on the number of samples required to guarantee that with high probability, the expected cost of the sampling-based policies is arbitrarily close (i.e., with arbitrarily small relative error) compared to the expected cost of an optimal policy that is computed based on the true demand distributions. The bounds that we develop are general, easy to compute and surprisingly do not depend at all on the specific demand distributions.

(ii) For the model with censored demand samples, we describe an adaptive policy that is based on a variant of the well-known Kaplan-Meier statistic. We show that the policy converges almost surely to the true newsvendor quantile, and that computationally it outperforms previously known policies.

The talk is based on several papers in joint work with Tim Huh, Jim Orlin, Georgia Perakis, Robin Roundy, Paat Rusmevichientong, David Shmoys and Joline Ann Villaranda Uichanco.

BIO

Retsef Levi is Associate Professor of Management at the Sloan School of Management, MIT. He is a member of the Operations Management Group at Sloan and affiliated with the Operations Research Center and the Computational for Design and Optimization Program. Before coming to MIT, he spent a year in the Department of Mathematical Sciences at the IBM T.J. Watson Research Center as the holder of the Goldstine Postdoctoral Fellowship. He received a Bachelor's degree in Mathematics from Tel-Aviv University (Israel) in 2001, and a PhD in Operations Research from Cornell University in 2005. Levi spent more than 11 years in the Israeli Defense Forces as an Office in the Intelligence Wing. After leaving the Military, Levi joined and emerging new Israeli hi-tech company as a Business Development Consultant.

Levi's current research is focused on the design and the performance analysis of efficient algorithms for fundamental stochastic and deterministic optimization models, arising in the context of supply chains and inventory, revenue management, logistics and healthcare management. These fundamental, multistage stochastic models are typically very hard to solve optimally, both theoretically and in practice. Hence, it is important to develop efficient heuristics that provide provably near-optimal policies for these hard models. Levi has special interest in Cost-Balancing techniques, data-driven (sampling-based) algorithms, and modern Linear-Programming-based approximation techniques applied to models in the above domains. In addition, he is interested in stochastic and combinatorial optimization and mathematical programming in their broad definition, and especially in their intersection with problems that arise in the context of real-life applications.

Levi has received the NSF Faculty Early Career Development award and the 2008 INFORMS Optimization Prize for Young Researchers.