

Dissertation Defense

Canan Gunes

Thursday, December 9, 2010

9:00 am

324 GSIA (West Wing)

Essays on Operations Management

My thesis focuses on the design and analysis of discrete-event stochastic simulations involving correlated inputs, input modeling for stochastic simulations, and application of OM/OR techniques to the operations of food banks. This thesis contributes to the stochastic simulation theory by describing how to jointly represent stochastic and parameter uncertainties in stochastic simulations with correlated inputs, and decompose the variance of the simulation output into terms related to stochastic uncertainty and parameter uncertainty. Such a decomposition would be beneficial for developing data collection schemas to reduce the parameter uncertainty in stochastic simulations. Furthermore, this thesis contributes to the vehicle routing theory by being the first work to rigorously study the 1-Commodity Pickup and Delivery Vehicle Routing Problem (1-PDVRP) that arises in the context of food rescue programs of food banks. A synopsis of the three chapters of the thesis follows.

Chapter 1: Accounting for Parameter Uncertainty in Large-Scale Stochastic Simulations with Correlated Inputs

This chapter considers large-scale stochastic simulations with correlated inputs having Normal-To-Anything (NORTA) distributions with arbitrary continuous marginal distributions. Examples of correlated inputs include processing times of workpieces across several workcenters in manufacturing facilities and product demands and exchange rates in global supply chains. Our goal is to obtain mean performance measures and confidence intervals for simulations with such

correlated inputs by accounting for the uncertainty around the NORTA distribution parameters estimated from finite historical input data. This type of uncertainty is known as the parameter uncertainty in the discrete-event stochastic simulation literature. We demonstrate how to capture parameter uncertainty with a Bayesian model that uses Sklar's marginal-copula representation and Cooke's copula-vine specification for sampling the parameters of the NORTA distribution. The development of such a Bayesian model well suited for handling many correlated inputs is the primary contribution of this chapter. We incorporate the Bayesian model into the simulation replication algorithm and the Bayesian simulation replication algorithm for the joint representation of stochastic uncertainty (i.e., the uncertainty that is due to the random numbers used in the simulation) and parameter uncertainty in the mean performance estimate and the confidence interval. Using the Bayesian simulation replication algorithm, we further decompose the variance of the simulation output into terms related to stochastic uncertainty and parameter uncertainty. Such a variance decomposition might be beneficial for developing data-collection schemas to reduce parameter uncertainty in stochastic simulations. A comprehensive numerical analysis shows that our model improves both the consistency of the mean line-item fill-rate estimates and the coverage of the confidence intervals in multi-product inventory simulations with correlated demands.

Chapter 2: Comparison of Least-Squares and Bayesian Inferences for Johnson's Bounded and Lognormal Distributions

Johnson translation system is a flexible system of distributions with the ability to match any finite first four moments of a random variable. This chapter considers the problem of fitting lognormal and bounded distributions of the Johnson translation system to finite sets of stationary, independent and identically distributed input data. The focus on the Johnson translation system is due to the flexibility it provides in comparison to the standard input models that are built in commercial input-modeling software. Specifically, Johnson's lognormal family is a positively skewed distribution with one-sided bounded support, while Johnson's bounded family contains two-sided bounded distributions capturing a wide variety of unimodal and bimodal distributional shapes. The bounded supports of these distributions make it difficult to obtain robust parameter estimates via the use of the maximum likelihood estimation method. However, the least-squares estimation method used for fitting Johnson translation system does not suffer from the existence of bounded supports; and it outperforms well-known fitting methods of matching moments and percentiles. Another fitting method that is frequently used for distributions with bounded supports as an alternative to the maximum likelihood estimation method is the Bayesian method. The main goal of this chapter is to investigate the relative performance of the least-squares estimation method and the Bayesian method in fitting the parameters of the distributions from Johnson's lognormal and bounded families, and provide guidelines to the simulation practitioner on when to use each fitting method.

Chapter 3: Food Banks Can Enhance Their Operations with OR/OM Tools: A Pilot Study with Greater Pittsburgh Community Food Bank

Food assistance programs have been challenged to serve an increasing number of low-income families in the recent economic downturn. Soaring demand is combined with diminishing supply (donations) attributed to both recession and donors' improved inventory management. As a remedy for food demand-supply mismatch, many food banks, whose primary goal is to reach as many needy people as possible, are trying to purchase more food by reducing their operational costs and by improving fundraising. In this study, we work together with our local food bank, Greater Pittsburgh Community Food Bank (GPCFB), in order to achieve the following two goals: First, by using the limited available data we illustrate the extent to which GPCFB is being affected by the recent economic downturn. We identify how they can collect better data for future use. We believe this will help GPCFB in their fundraising efforts. Second, we focus on the 1-PDVRP that arises in the context of the food rescue program of GPCFB. We present a thorough study on the state-of-the-art solution methods for the 1-PDVRP, utilizing technologies Mixed Integer Programming, Constraint Programming, and Constraint-based Local Search, and evaluate potential cost savings with respect to the current practice of GPCFB. The results indicate substantial cost savings, being at least 10% on the largest instance, which can be used to purchase more food and to reach more people in need. In addition to the practical value of this work for GPCFB, this study contributes to the theory by being the first academic work to provide a rigorous treatment of the 1-PDVRP. Overall, this study not only seeks to help GPCFB, but it is intended as a good starting place for other food banks around the U.S as they struggle with similar issues.