Operations Research

In This Issue

To cite this article:
(2009) In This Issue. Operations Research 57(3):iv-vii. https://doi.org/10.1287/opre.1090.0723

OPERATIONS RESEARCH
Vol. 57, No. 3, May–June 2009, pp. iv–vii
ISSN 0030-364X | EISSN 1526-5463 | 09 | 5703 | 00iv

informs®

DOI 10.1287/opre.1090.0723

In This Issue

Moving Retail Supply Chain Management from Art to Science

Retailers face the challenge of knowing their customers and providing them with the right products at the right price and in the right quantities. Years ago, when the "mom and pop" format was dominant, "knowing customers" had a very literal meaning as retailers personally knew their customers and used this direct knowledge combined with human judgment to understand what products their customers would like. As the mom and pops grew into today's mega chains, decision making continued to be more art than science, with the result that many retailers suffer from both frequent stock outs and a high level of excess inventory that needs to be cleared with deep end-of-season markdowns. But in the last decade much new data has become available to retailers to help them better understand their customers, including Point-of-Sales (POS) data, loyalty cards, customer satisfaction survey results, click stream records, and grocery smart carts. As a result, retailers have begun to apply analytic models to this huge volume of data to enable better decisions. This industry-wide movement resembles the transformation of Wall Street that occurred in the 1970s when physicists and other "rocket scientists" applied their analytic skills to investment decisions, and is slowly tipping the balance of retail supply chain management from mostly art to a judicious balance of art and science. The Consortium for Operational Excellence in Retailing (COER) (codirected by Marshall Fisher of The Wharton School and Ananth Raman of Harvard Business School) is a group of academics working with about 50 leading retailers to assess their progress towards rocket science retailing and to accelerate that progress through selected research projects. After some brief comments on the current state of industry practice in retail supply chain management, in "Rocket Science Retailing: The 2006 Philip McCord Morse Lecture" M. Fisher describes examples of COER research in four areas: assortment planning, pricing, inventory optimization, and store execution. This article is the archival record of the INFORMS Philip M. Morse Lecture delivered in 2006.

Dynamic Hedging of Risky Assets

There is strong evidence that risky assets in real markets follow jump diffusion processes. In this case, the standard Black-Scholes delta hedge will produce poor results when hedging a contingent claim. In "Dynamic Hedging Under Jump Diffusion with Transaction Costs" by J. S. Kennedy, P. A. Forsyth, and K. R. Vetzal, a dynamic hedging strategy is devised that uses the underlying asset and liquidly traded options, where transaction costs are present due to bid-ask spreads. Simulations of this procedure are carried out, which show that the proposed dynamic hedging strategy provides good protection against the diffusion and jump risk while not incurring large transaction costs.

Stable Portfolio Selection

It is well known that, due to estimation error, mean-variance portfolios perform poorly. One of the reasons for this is that the mean-variance portfolios computed using the traditional sample estimators tend to have extreme portfolio weights that fluctuate substantially over time. To address this issue, in "Portfolio Selection with Robust Estimation," V. DeMiguel and F. J. Nogales combine methodologies from the fields of optimization and robust statistics to propose a class of portfolio policies that have better stability properties than the traditional portfolios. The proposed portfolios are constructed using certain robust estimators; that is, estimators that are insensitive to deviations of the empirical return distribution from normality. The authors show both analytically and empirically that the resulting portfolio weights are more stable than those of the traditional portfolios.

Defeating Insurgencies Is Hard

From Afghanistan through Colombia to Iraq, insurgencies are the most common type of military conflict nowadays. Small, loosely organized and ill-equipped groups of insurgents confront a larger, better organized, trained and equipped government force that, in spite of its quantitative and qualitative advantages, experiences difficulties in controlling the insurgents. For example, during the first half of 2007 there was a steady increase of U.S. troops in Baghdad (Brookings Iraq Index, October 2007) with mixed results. In certain regions of Afghanistan the Taliban are still strong, and in Colombia government actions seem to have little effect on the insurgents. So, why is it that the larger and stronger force cannot easily defeat the smaller and weaker one? In "Why Defeating Insurgencies Is Hard: The Effect of Intelligence in Counterinsurgency Operations —A Best-Case Scenario," M. Kress and R. Szechtman suggest a possible answer to this question, and offer some insights about insurgency situations. They model the insurgency as a dynamic system, with special emphasis on the effects of intelligence and population response to government actions against the insurgents. Under very general assumptions, they show that the best the government can do is to contain the insurgency. They identify conditions under which the insurgency ultimately wins, and situations where a stalemate emerges.

Changing Plans Without Disrupting Operations

Often a transportation company's day-to-day operations are dictated by a schedule or an operating plan. For example, each airline has a flight schedule, and passengers fly using this schedule. A passenger railroad has a railway timetable, and passengers use trains according to this timetable. A freight railroad or a trucking company also has an operating plan and carries freight per this plan. The operating plan for such a company dictates the movement of its resources (crews, planes, buses, locomotives, railcars, trucks). When the operating plan is not optimal, the transportation company might wish to change it in a way that is not too disruptive to its day-to-day operations. That is, the company would like to make an incremental change to its current plan and move in the right direction. Once this change has been made and benefit reaped, it will be ready for another incremental change. In "Incremental Network Optimization: Theory and Algorithms," O. Şeref, R. K. Ahuja, and J. B. Orlin study such incremental optimization problems where a given solution is incrementally improved within the given degree of freedom.

Assessing the Quality of Maximin Latin Hypercube Designs

Latin hypercube designs (LHDs) play an important role when approximating computer simulation models. To obtain good space-filling properties, the maximin criterion is frequently used. Unfortunately, constructing maximin LHDs can be quite time consuming when the number of dimensions and design points increases. In these cases, heuristical maximin LHDs are used. In "Bounds for Maximin Latin Hypercube Designs," E. R. van Dam, G. Rennen, and B. Husslage construct bounds for the separation distance of certain classes of maximin LHDs. These bounds are useful for assessing the quality of heuristical maximin LHDs.

Effective Design of Less-than-Truckload Transportation Network

Less-than-Truckload (LTL) carriers specialize in transporting shipments that are usually between a few hundred to a few thousand pounds and do not individually fill entire trailers. Such carriers are able to provide a low-cost transportation option to their customers by consolidating shipments using a hub-and-spoke system that is similar in motivation to that used in the airline industry. A major challenge for an LTL carrier is to design an effective network that is capable of routing all anticipated freight from origin to destination while minimizing the fixed-charge and variable costs for the network's links and meeting customer service standards. In "Large-Scale, Less-than-Truckload Service Network Design," A. I. Jarrah, E. Johnson, and L. C. Neubert formulate this large-scale problem as a massive multicommodity fixed-charge network design with side constraints model. They propose a novel decomposition approach that enables the generation of high-quality solutions for medium and large networks. Extensive runs document the savings that can be reaped and the computational effort required.

The Benefits of Postponing Order-Quantity Decisions

In many logistics systems, inventory replenishment by retailers consists of two distinct stages. In the first stage, retailers may make a delivery request in advance without making a firm commitment on order quantity. In the second stage, the retailers are allowed to make an adjustment to the quantity requested just before the shipping schedule is finalized, or even allowed to adjust the quantity at the time the delivery vehicle has arrived at their receiving docks. In contrast to more traditional approaches by which delivery request and quantity decisions must be made at the same time, this postponement strategy allows retailers to postpone their quantity decision until better demand information is available, and, as a result, the mismatch between supply and demand is reduced. In "Optimal Policies for Inventory Systems with Separate Delivery-Request and Order-Quantity Decisions," Q. Li, X. Wu, and K. L. Cheung study the inventory problem faced by a retailer who replenishes its stock by using this approach. They characterize the optimal policy for delivery requests and for ordering quantities and the benefit from applying this postponement strategy.

Computing Bid Prices in Network Revenue Management

Major airlines operate complex flight networks and usually carry their passengers through a set of connecting flights. As a result, passenger demands for different flight legs tightly interact. When airlines need to decide whether a certain fare class should remain open or be closed, they need to consider the impact of their decisions over the whole flight network. In other words, to obtain globally optimal results, the capacity management decisions need to be made with a view of the entire flight network. In "Using Lagrangian Relaxation to Compute Capacity-Dependent Bid Prices in Network Revenue Management," H. Topaloglu develops an algorithm that makes the capacity management decisions by concentrating on one flight leg at a time but still obtains results that are close to being globally optimal. The primary advantage of the algorithm is its scalability. When one tries to maintain a view of the flight network, the problem easily becomes intractable as the size of the flight network grows. In contrast, problem size is more manageable when one concentrates on one flight leg at a time.

Incorporating Estimation Risk and Trading Constraints in Portfolio Optimization

Two of the well-documented limitations of the Markowitz asset selection model are that it does not consider the uncertainty in the estimation of assets returns and a number of real-life trading restrictions. In "An Exact Solution Approach for Portfolio Optimization Problems Under Stochastic and Integer Constraints," P. Bonami and M. A. Lejeune propose a new probabilistic portfolio selection model and develop an exact solution method that simultaneously account for the estimation risk and for the requirements imposed by trading restrictions. The authors support their theoretical results with computational tests that show the scalability, robustness, and speed of the approach. They discuss how their model and approach have been used in the construction of four fund-of-funds that are now traded on the major stock exchanges.

Upgrading Customers and Rationing Products

Manufacturers and retailers often offer a range of products that can be ordered from lowest to highest quality, and higher-quality products can be used to satisfy demand for lower-quality products. For example, a computer manufacturer may ship a large disk drive to a customer who has requested a smaller size that is not available. In "Dynamic Capacity Management with Substitution," R. A. Shumsky and F. Zhang describe how capacity for such products should be managed, given that capacity allocation decisions must be made before all customers have arrived and customers may only be upgraded by one product level. They show that a simple threshold rationing policy for higher-quality products is optimal, and they describe heuristics that can be used to efficiently and accurately calculate optimal rationing levels. They apply the heuristics to numerical examples, and these examples reveal the conditions under which the use of optimal policies for capacity replenishment, upgrading, and rationing have a significant impact on profits.

Staffing Call Centers

When firms outsource their call centers, they enter into service level agreements that specify acceptable performance on the part of the call center. Often these indicate penalties that will be assessed for poor performance. The time frame over which these service levels are measured is an important parameter. While researchers have previously considered the probability that an individual customer will be delayed in or abandon a queue, the average performance of a call center for all customers arriving during a period such as a rush hour may be a more salient measure of performance. In "Staffing to Maximize Profit for Call Centers with Alternate Service-Level Agreements," O. Baron and J. Milner study the impact of the way service levels are measured on the staffing chosen by a call center. They develop an approximation for the probability that customers arriving during a period are served according to stated service levels, and they demonstrate that measuring call centers based on this probability results in staffing levels more robust to uncertainty in the mean demand than measures based on either individual expectations or long-term expectations.

Generalized Geometric Programs with Mixed Free-Sign Variables

Many nonlinear optimization problems, such as capital investment, digital circuit design, batch plant modeling and others, can be formulated as generalized geometric programs (GGP). These GGP problems often contain continuous and discrete functions where the variables may represent the sizes of investment, thicknesses of steel plates, and temperatures. All these variables may have positive, zero, or negative signs. Many current nonlinear programming algorithms cannot be applied directly to solve these GGP problems, especially when the signs of variables are free and the number of discrete variables is large. In "Global Optimization for Generalized Geometric Programs with Mixed Free-Sign Variables," H.-L. Li and H.-C. Lu propose a novel method to solve GGP problems with various sign variables to obtain a globally optimal solution. The proposed method is computationally more effective than current methods in treating programs with a large number of discrete variables.

Data-Driven Staffing in Telephone Call Centers

The problem of setting optimal staffing levels is central to the management of modern telephone call centers. In "On a Data-Driven Method for Staffing Large Call Centers," A. Bassamboo and A. Zeevi describe a systematic approach to this problem. A salient feature of the proposed method is that it relies on past data rather than detailed structural and probabilistic assumptions, and can be easily applied in complex large-scale settings with multiple customer types and multiple agent skill sets. The paper first describes how the staffing problem can essentially be "decoupled" from the control problem that involves real-time routing of calls to agents. It then proposes a nonparametric staffing prescription that uses past data on call arrival epochs. The method is shown to achieve near-optimal performance in an operating regime of practical interest.

Sampling Points from Arbitrary Distributions

Simulations of complex systems usually include Monte Carlo procedures for emulating the various chance events whose occurrences influence the evolution of the underlying system. In "Discrete Hit-and-Run for Sampling Points from Arbitrary Distributions Over Subsets of Integer Hyperrectangles," by S. Baumert, A Ghate, S. Kiatsupaibul, Y. Shen, R. L. Smith, and Z. B. Zabinsky, the authors provide an algorithm that efficiently generates these multidimensional chance events for extremely general discrete distributions. The engine of the algorithm entails two independent random walks that terminate when each steps outside of the feasible domain. Applications are discussed for the optimization of discrete systems modeled through black-box performance functions and feasibility oracles.

Technology Adoption Under Uncertainty

Consumers or firms contemplating purchasing a new product or adopting a new technology are often plagued by uncertainty: Will the benefits outweigh the costs? Should we buy now or wait and gather more information? With today's rapidly evolving technological landscape and abundant sources of information, it is particularly important to understand how information and uncertainty affect adoption and purchase decisions. In "Uncertainty, Information Acquisition, and Technology Adoption," C. Ulu and J. E. Smith develop and analyze a dynamic programming formulation of this problem in a general setting where the consumer's beliefs about a technology's benefits are described by a probability distribution. They establish a number of theoretical results that provide insights into these adoption decisions. In particular, they show that improving the technology's benefit need not make the consumer better off once the information gathering costs and adoption process are taken into account. They characterize "more optimistic" beliefs and better information sources that make the consumer better off and investigate how changes in beliefs and information sources, among other things, affect optimal adoption policies.

Lead-Time Quotation and Scheduling in Service Operations

Many service systems are characterized by costly delays and limited capacity. In "Near-Optimal Dynamic Lead-Time Quotation and Scheduling Under Convex-Concave Customer Delay Costs," B. Ata and T. L. Olsen examine the effect the shape of the delay cost function has on lead-time quotation and order sequencing. They focus on the trade-off between committing future capacity now or reserving it for potential later higher revenue (lower cost) customers.

Accounting for Customer Choice Behavior in Revenue Management

Traditionally, revenue management has been built upon the independent demand paradigm. For instance, models used to assume that demand for different booking classes in the economy cabin of a commercial flight was mutually independent. However, customers behave in a more complex way, and the bookings that the airline observes are a function of the different products made available at different points in time during the booking horizon. During the past few years there has been a growing interest to enrich these models by accounting for customer choice behavior. This extension involves both modeling and computational challenges. In "A Column Generation Algorithm for Choice-Based Network Revenue Management," J. J. M. Bront, I. Méndez Díaz, and G. Vulcano devise a procedure to solve the firm revenue maximization problem under a general market model of demand segmentation. Their computational results show that the methodology is quite effective and leads to high quality implementable solutions.

Inventory Competition in a Multiperiod Environment

Retailers often worry about competition resulting from customers willing to search competitors when products stock out. Typically, this competition is difficult to analyze due to limited information that a particular retailer may have about his competitor's inventory policies. In "A Multiperiod Model of Inventory Competition," M. Nagarajan and S. Rajagopalan show that in a duopoly, under certain reasonable conditions on the cost parameters and correlation of demands, in equilibrium, the two firms ignore their competitor's inventory levels and potential substitution demand in making their own inventory decisions, i.e., their inventory decisions are decoupled. This result holds in long lifecycle products (i.e. multiperiod models) as well as where in equilibrium the retailers will simply solve a single-product dynamic newsvendor problem that ignores substitution demand.

The Censored Newsvendor Problem

A censored newsvendor does not observe lost sales, and thus has only incomplete information about his demand. This results in a partially observed system for which a use of unnormalized probabilities simplifies the analysis considerably. In "A Note on 'The Censored Newsvendor and the Optimal Acquisition of Information,"' A. Bensoussan, M. Çakanyıldırım, and S. P. Sethi provide a concise proof of earlier claims that the optimal order quantity of the censored newsvendor is always more than or equal to the myopic order quantity.