Operations Research

In This Issue

To cite this article:
(2014) In This Issue. Operations Research 62(6):ii-iv. https://doi.org/10.1287/opre.2014.1337

OPERATIONS RESEARCH
Vol. 62, No. 6, November–December 2014, pp. ii–iv
ISSN 0030-364X (print) | ISSN 1526-5463 (online)

http://dx.doi.org/10.1287/opre.2014.1337

In This Issue

Optimizing Bank Capital Requirements

Banking regulations set minimum levels of capital for banks. These requirements are generally formulated through a ratio of capital to risk-weighted assets. In "Design of Risk Weights," Glasserman and Kang bring an optimization perspective to the problem regulators face in specifying a risk-weighting procedure. A surprising conclusion of their analysis is that risk weights should be proportional to profitability to deter banks from overinvesting in risky assets. The authors also propose an adaptive procedure in which risk weights are automatically adjusted in response to changes in bank portfolios, for example, increasing the risk weight on real estate loans as banks take on more real estate risk. This adaptive feature helps guard against the deficiencies of a fixed scheme.

Noisy Communication as a Game

When a sender transmits a message as a "codeword" over a noisy channel, the receiver's optimal "decoding" of the possibly garbled channel output is to identify the message that has most likely been sent. In "Nash Codes for Noisy Channels," P. Hernández and B. von Stengel extend this standard approach by asking if the sender's action is also optimal. That is, given the receiver's decoding strategy, is the sender most likely understood correctly by transmitting the original codeword, rather than any other channel input? A "Nash code" has this desirable stability property, where encoding and decoding define a Nash equilibrium in this sender-receiver game with fully aligned interests of the players. The paper studies this concept with examples, and shows two sufficient conditions for Nash codes: receiver-optimal codes, and, somewhat surprisingly, *arbitrary* codes for the classic binary channel (if the receiver breaks ties "monotonically").

Revealing Hidden Edge Weights and Distances by means of Path Exploration

In the paper "Economical Graph Discovery," N. Alon, Y. Emek, M. Feldman, and M. Tennenholtz introduce a novel graph discovery model that captures an input graph G with known topology and hidden edge weights. Given designated source vertex s and destination vertex t, the discovery is done by means of agents that traverse (s, t)-paths in G and report the total cost they incurred. Depending on the model variant, this cost can be simply the weighted length of the traversed path or it can be affected by congestion caused by other agents that use the same edges. Two different discovery problems are studied: (i) edge weight discovery, where the goal is to reveal all edge weights; and (ii) shortest-path discovery, where the goal is to reveal a shortest (s, t)-path. The authors establish various lower and upper bounds on the number of agents required to solve these problem expressed in terms of the number of vertices and edges.

Using Optimization to Reconcile Community with School Choice

Many public school districts in the United States allow families to submit preferences on where to send their children to school, and use these along with priorities and lotteries to assign seats. Such systems help families place their children where best matches their needs. A shortcoming is that children on the same street may be sent to many different schools, hence weakening community ties and contributing to neighbors not knowing one another. In "Improving Community Cohesion in School Choice via Correlated-Lottery Implementation," P. Shi and I. Ashlagi propose an approach to reconcile choice with community: by using optimization to correlate the lottery numbers, each individual can have the same chances at each school as before but neighbors can have higher chances of going to school together. The amount of improvement depends on how correlated families' preferences are. In Boston, where this issue was recently salient, this approach can significantly improve community cohesion without compromising choice.

Managing Elective Admissions to Reduce Variability in Hospital Workloads

Occupancy fluctuation in hospitals causes bed block that leads to overcrowding in emergency departments and intensive care units. Too often, the result is preventable mortalities, compromised quality of care, emergency patient diversions, canceled surgeries and excess cost. In "Design and Optimization Methods for Elective Hospital Admissions," J. E. Helm and M. P. Van Oyen propose optimization methods to determine the correct mix and volume of elective patients to admit over time to reduce the variance in hospital occupancy. Using network models to accurately capture stochastic patient flows through hospital wards, hospital-wide occupancies can be modeled and controlled by managing inpatient admissions. A study using admit-discharge-transfer data from hospitals in four countries reveals that the approach is generalizable. Further, in depth analysis of one of these hospitals shows that bed block can be reduced by 32% or surgical volume can be increased by 7% with the same level of bed block.

Portfolio Management with Multiple Clients

Since the seminal work of Markowitz (1952), multiple facets of the portfolio optimization problem have been examined in the literature. The vast majority of the studies have focused on a setting where a financial adviser is acting on behalf of a single client in order to optimally select, rebalance or liquidate her portfolio. In practice, however, advisers rarely manage a single portfolio, as they typically offer their services to multiple clients simultaneously. In the paper "Fairness and Efficiency in Multiportfolio Optimization," Iancu and Trichakis discuss the unique challenges that arise in the latter setting. The authors propose a novel framework that allows joint optimization of multiple portfolios, subject to SEC regulations. Their model introduces a novel method of fairly splitting market impact costs, while offering the manager the flexibility of selecting her preferred notion of fairness in balancing the performance of all portfolios she is in charge of.

Robust Portfolio Selection

The issue of ambiguous return distributions in portfolio optimization problems recently received a lot of attention in the literature. When neglected, this ambiguity coupled with an erroneously specified return distribution may lead to suboptimal out-of-sample results in portfolio selection, as is demonstrated in numerous empirical studies. In the paper "Robustifying Convex Risk Measures for Linear Portfolios: A Nonparametric Approach," David Wozabal develops a framework to robustify a range of portfolio optimization problems that incorporate convex risk measures to explicitly account for the distributional ambiguity in the problem. The resulting robust problems are essentially regularized versions of the nominal problems and remain computationally tractable. The connection between regularization and robustification allows to connect the resulting optimization problems to norm-constrained portfolio problems. An extensive out-of-sample study demonstrates that the robustified problems outperform the original/nominal problems for most of the considered data sets.

Stability and Endogenous Formation of Inventory Transshipment Networks

In a cooperative game of inventory transshipment, firms make their inventory decisions independently, and then decide collectively how to transship excess inventories to satisfy unmet demands. In "Stability and Endogenous Formation of Inventory Transshipment Networks," X. Fang and S. Cho develop a novel approach of modeling transshipment by using *networks* of firms, which offer a richer representation of relationships among firms by taking the coalitions used in prior literature as special cases. The authors show that various network structures such as complete, hub-spoke, and chain networks are stable only under certain conditions under the allocation based on firms' marginal contributions to networks. Moreover, these conditions differ across network structures, implying that a network structure plays an important role in establishing the stability of a decentralized transshipment system. Finally, they analyze the case when firms establish networks endogenously, and show that pairwise Nash stable networks underperform corresponding networks in centralized systems.

Dynamic Inventory and Pricing Control

There has been substantial literature on the joint inventory and pricing control problem for a firm that faces stochastic and price-sensitive demand. The existing results in this area require the joint concavity (in inventory and price/demand) of the expected single-period profit function and/or the stationary parameters. In "Joint Inventory and Pricing Control with General Additive Demand," H. Chen and Z. Zhang relax the concavity assumption to quasi-concavity in the sufficient conditions for the optimality of a time-dependent (s, S, p) policy. This generalization makes the structural results applicable to models broadly supported by economic theory and empirical data.

An Integrated Approach for Simulating Evacuation Dynamics in Large Public Places

Pedestrian dynamics plays an important role in public facility design and evacuation management. Modelling the dynamic escape process in large public places is an extremely complex challenge. In the paper "Simulating the Dynamic Escape Process in Large Public Places" by Ziyou Gao, Yunchao Qu, Xingang Li, Jiancheng Long, and Hai-Jun Huang, an integrated macro-micro approach has been developed to simulate the escape process. Based on the concept of the dynamic user optimal criterion, a macroscopic variational inequality model is proposed to describe the exit and route choice behaviors. At the microscopic level, a modified social force model is adopted to formulate the pedestrians' dynamic movements. In addition, a solution algorithm is proposed to solve the macro-micro integrated model and a series of experiments are carried out to validate the model. This method can be used to achieve a more realistic prediction of collective movement pattern, decision-making behavior and self-organizing phenomena.

A Unifying Framework for Distributionally Robust Optimization

Distributionally robust optimization (DRO) is an emerging field that provides scalable solution methods for optimization problems affected by uncertainty. Through the choice of appropriate risk measures and ambiguity sets, DRO problems account for the decision maker's attitude towards both risk (uncertainty in the outcome) and ambiguity (uncertainty in the probability distribution governing the outcome). While there has been significant recent progress in DRO, there is no unifying framework for modeling and solving DRO problems. In the paper "Distributionally Robust Convex Optimization," W. Wiesemann, D. Kuhn, and M. Sim develop a canonical form for DRO problems that unifies several recent contributions and at the same time extends to several new classes of DRO problems that have not yet been studied in the literature. The paper establishes conditions under which DRO problems based on their canonical form are computationally tractable, and they provide tractable conservative approximations for problems that violate these conditions.

When Is It Worth Investing in the Development of a Stochastic Model?

When one needs to decide in an uncertain context, the task of developing a compelling stochastic model is never a trivial one. It might require collecting a large amount of observations, consulting with experts, or resorting to simplifying assumptions. Alternatively, a common heuristic is to make decisions simply based on expected values. In "The Value of Stochastic Modeling in Two-Stage Stochastic Programs with Cost Uncertainty," E. Delage, S. Arroyo, and Y. Ye identify conditions under which this heuristic leads to robust decisions. The authors also provide the means of quantifying the added value of developing a legitimate stochastic model for the uncertainty in terms of how such a model might lead to a better decision. Experiments on a fleet composition problem for three airline companies indicate that this value was under 7% and might not justify to invest the resources in such modeling.

Computing Performance Bounds in Convex Stochastic Dynamic Programs

Stochastic dynamic programs are often very difficult to solve to optimality. Consequently, we often rely on heuristic policies that may not be optimal. Although we can evaluate the performance of a heuristic policy using Monte Carlo simulation, it is also useful to know how much better we could do with an optimal policy. In "Information Relaxations, Duality, and Convex Stochastic Dynamic Programs," D. Brown and J. Smith study the information relaxation approach for calculating upper bounds on the expected reward given by an optimal policy. This approach relaxes the nonanticipativity constraints that require decisions to depend only on the information available and imposes penalties that punish violations of these nonanticipativity constraints. The authors study the use of "gradient penalties" for dynamic programs with a convex structure. These penalties are useful because they lead to subproblems that can be solved using deterministic convex optimization techniques. The paper develops the theory of these methods and demonstrates their use in a network revenue management problem and in an inventory management problem with lost sales and lead times. In these challenging applications, the gaps between the upper bounds from the information relaxations with gradient penalties and the lower bounds from some simple heuristic policies are small, allowing us to conclude these simple heuristic policies are nearly optimal.

Optimization via Simulation Using Gaussian Process-based Search

Random search algorithms are often used to solve discrete optimization-via-simulation (DOvS) problems. A critical component of a random search algorithm is the sampling distribution that is used to guide the allocation of the search effort. A good sampling distribution can balance the tradeoff between the effort used in searching around the current best solution (which is called exploitation) and the effort used in searching largely unknown regions (which is called exploration). In the paper "Balancing Exploitation and Exploration in Discrete Optimization via Simulation through a Gaussian Process-based Search," L. Sun, J. Hong, and Z. Hu, propose a new scheme that derives a sampling distribution from a fast fitted Gaussian process based on previously evaluated solutions. The authors show that the sampling distribution has desired properties and it can automatically balance the exploitation and exploration tradeoff. Furthermore, they integrate this sampling distribution into a random research algorithm, called a Gaussian Process-based Search (GPS), and demonstrate the merits of the algorithm in solving DOvS problems.

Measuring Simulation Model Risk from using Estimated Distributions

Decisions based on mathematical and computer models are subject to risk because the models are never perfect reflections of reality. In stochastic computer simulation, the probability distributions that drive the simulation are a source of risk when they are estimated from (necessarily limited) observations of real-world input data. The statistics literature provides ways to quantify the error in estimated input-distribution parameters; the problem here is to assess the implied error in the simulation outputs that are functions of these inputs, and upon which the decisions will be based. In "A Bayesian Framework for Quantifying Uncertainty in Stochastic Simulation," W. Xie, B. L. Nelson, and R. R. Barton present the first fully Bayesian treatment of the input uncertainty problem in stochastic simulation, and illustrate its use in the simulation of a critical care facility. Along the way they clarify the difference between a Bayesian assessment of input uncertainty and the classical frequentist approach.

Bandwidth Sharing Networks with Rate Constraints

Bandwidth sharing networks provide a natural framework for describing the dynamic flow-level interaction among elastic data transfers in computer and communication systems. They may also be used to develop traffic pricing/charging mechanisms. In the paper "Limit Theorems for Markovian Bandwidth Sharing Networks with Rate Constraints," J. Reed and B. Zwart develop a framework to approximate bandwidth sharing networks under the assumption that the number of users as well as the capacities of the system are large, and the assumption that the traffic that each user is allowed to submit is bounded above by some rate. Under Markovian assumptions, they develop fluid and diffusion approximations for such networks which are quite tractable. In particular, for most parameter combinations, the invariant distribution of the diffusion approximation is multivariate normal, with mean and diffusion coefficients that can be computed in polynomial time as a function of the size of the network.