



Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

In This Issue

To cite this article:

(2018) In This Issue. Operations Research 66(2):ii-iv. <https://doi.org/10.1287/opre.2018.1735>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2018, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

In This Issue

<https://doi.org/10.1287/opre.2018.1735>

Copyright: © 2018 INFORMS

Why Do Firms Offer Time-Locked Free Trial Periods?

Time-locked free trial periods are a ubiquitous practice among firms launching innovative products, and are particularly prevalent in the software industry, where the negligible production cost makes it a feasible business strategy. In “Signaling Product Quality through a Trial Period,” S. Wang and G. Ozkan-Seely recognize that a firm and the market may not be equally informed about product quality prior to consumption and they posit a signaling explanation for the existence of free trial periods of certain lengths offered in the marketplace. The authors find that the length of a free trial period and the price can serve as dual signals of the firm’s privately known product quality, and establish a more efficient signal than the well-recognized price alone. They show that a high-quality firm offers a longer trial period and sets a higher price, and is rewarded with a higher profit, relative to its low-quality counterpart.

Crowdsourced Prediction System: Pay All or Pay One?

In “Parametric Prediction from Parametric Agents,” Y. Luo, N. B. Shah, J. Huang, and J. Walrand, focus on a prediction scenario where a principal wants to predict the outcome of an event based on opinions elicited from agents with different capabilities and different private information. The authors design a mechanism named COPE (COst and Prediction Elicitation), which optimizes the principal’s utility that depends on the payments to the agents and the error incurred in the prediction. By employing COPE, the principal can incentivize all participating agents to exert appropriate effort levels based on their capabilities and to report their predictions truthfully. Their results reveal a striking dichotomy under Gaussian estimation noise: when the costs incurred by the agents are linear in the exerted effort, the principal should solicit service only from the best agent with the lowest reported cost. On the other hand, when the agent’s costs are quadratic, the principal should recruit multiple agents to jointly complete the task.

Modeling Traffic Congestion as a Sequential Game

Vehicle traffic in a city is determined by the decisions of many drivers who choose which road to take when going from their origin to their destination. These decisions are strategic since each driver takes into account the other drivers’ choices and is affected by them. Moreover, these decisions have consequences that spread dynamically on the road network. In “Dynamic Atomic Congestion Games with Seasonal Flows” M. Scarsini, M. Schröder, and T. Tomala model traffic congestion as a sequential game and measure the inefficiency that the selfish behavior of drivers causes to the system. To do this they use some standard inefficiency measures, like the price of anarchy or the price of stability. They show that their

dynamic model substantially differs from the more common static versions and that inefficiency is typically caused by the selfish behavior of the first drivers.

Preservation of Structural Properties in Optimization

Many operations management problems can be formulated as stochastic optimization problems in which the decision variables are truncated by some random variables. Examples include inventory control with random capacities and revenue management capacity allocation problem with random demands. An intrinsic challenge arises from the fact that the truncation by random variables may destroy convexity of the underlying optimization problem. In “Preservation of Structural Properties in Optimization with Decisions Truncated by Random Variables and Its Applications,” X. Chen, X. Gao, and Z. Pang develop a novel transformation approach that converts the original problem to an equivalent convex minimization problem. They show that their approach can be applied to prove the preservation of some desired structural properties, such as convexity, submodularity and L-natural-convexity. They then demonstrate the applications of this approach to several important operations models such as dual sourcing with random capacity, assemble-to-order systems with random capacity, and capacity allocation in network revenue management.

Robust Deployment of Public Access Defibrillators

Sudden cardiac arrest is a leading cause of death and a significant public health concern. Bystanders of cardiac arrest can improve survival outcomes by treating the cardiac arrest victim using an automated external defibrillator (AED). However, effective placement of AEDs is made challenging due to uncertainty in the locations of cardiac arrest events. In “Robust Defibrillator Deployment Under Cardiac Arrest Location Uncertainty via Row-and-Column Generation,” T.C.Y. Chan, Z.-J. Max Shen, and A. Siddiq propose a large-scale robust optimization model for deploying AEDs given uncertainty in the spatial distribution of future cardiac arrest locations. They propose an efficient and generalizable row-and-column generation algorithm for solving the AED deployment problem, which is otherwise intractable under standard Benders-based and mixed-integer programming approaches. Using real cardiac arrest location data, the authors show that the robust formulation can decrease the distance between cardiac arrest events and nearby AEDs by up to 15% over a sample average approximation approach, leading to improved survival outcomes.

Capacitated Assortment Optimization under Multinomial Logit Model

In “Capacitated Assortment Optimization under the Multinomial Logit Model with Nested Consideration Sets,”

J. Feldman and H. Topaloglu study capacitated assortment problems when customers choose under the multinomial logit model with nested consideration sets. In this choice model, there are multiple customer types and a customer of a particular type is interested in purchasing only a particular subset of products. The consideration sets of customers of different types are nested in the sense that the consideration set of one customer type is included in the consideration set of another. The goal of the assortment problem is to find a set of products to offer to maximize the expected revenue obtained from a customer, while making sure that the total space consumption of the offered products does not exceed a certain limit. The authors show that this assortment problem is NP-hard, even when there is no limit on the total space consumption of the offered products. Motivated by this complexity result, they give a fully polynomial time approximation scheme for the problem.

Risk Preference in Business-to-Business Service Contracting: An Empirical Study

The use of IT infrastructural services, from utility computing to printer management, has been rapidly increasing. Despite the growing popularity of such services, the question whether service providers are risk-averse or risk-neutral remains un-answered. In “Risk-Aversion and B2B Contracting under Asymmetric Information: Evidence from Managed Print Services,” J. Ning, V. Babich, J. Handley, and J. Keppo provided an answer by conducting empirical research in the context of managed print services (MPS). The authors constructed a structural econometric model of the contracting and usage processes of MPS under asymmetric information. Using data during 2010–2011 from Xerox Corporation, the authors found that Xerox’s sales team behaved in a risk-aversion fashion when contracting with customers. With traditional explanations for risk-aversion inapplicable for Xerox MPS, the authors found that Xerox’s commission holdback policy may be causing this risk-aversion. Using counterfactual analyses, the authors determined the consequences of rescinding the commission holdback policy, which include an increase in the service provider’s valuation and changes in contract pricing and printer model selection.

Sufficiency of Two-Prices for Dynamic Pricing in Observable Queues

Dynamic pricing is a strategy that is becoming increasingly prevalent in systems with congestion, such as in highways, ride-hailing services such as Uber, etc. As one can imagine optimal dynamic pricing policies can be quite complex to compute and also to implement. So, a natural question arises—Is there significant value to dynamic pricing over static pricing and if so, whether simple dynamic pricing schemes exist that can capture this value? In “The Value of Dynamic Pricing in Large Queueing Systems,” J. Kim and R. Randhawa find that indeed dynamic pricing can generate significant value, and somewhat surprisingly, most of this value can be realized by a simple policy that uses only two-price points—a high price when congestion is high, and a low price when congestion is low. The authors establish these results analytically using an asymptotic large system approach.

Using the Strong Law of Large Numbers to Derive Robust Inventory Management Policies

In “Robust Inventory Management: An Optimal Control Approach,” M. R. Wagner applies robust optimization techniques, coupled with uncertainty sets motivated by the strong law of large numbers for stochastic processes, to derive closed-form ordering policies that are transparent and easily implementable, in both static and dynamic settings. In the latter scenario, the issue of consistency between observed demand and robust uncertainty sets is explored, and inconsistencies are resolved using Hilbert’s Projection Theorem. Computational studies are encouraging, especially with seasonal demand.

Bias of Multilevel Monte Carlo Estimators Can Be Removed Without Loss of Efficiency

Multilevel Monte Carlo (MLMC) is a powerful technique to bring down the computational complexity in stochastic differential equation model inference (Giles 2006, *Oper. Res.*, 56(3)). Rhee and Glynn (2015 *Oper. Res.*, 63(5)) showed recently that similar order of complexity can be reached using randomized estimators, which eliminate the estimation bias entirely. In “Unbiased estimators and multilevel Monte Carlo,” M. Vihola proposes new unbiased estimators, based on a stratified randomization, which behave asymptotically like MLMC, both in terms of variance and cost, under general conditions—essentially those that guarantee canonical square root error rate for MLMC. This suggests that MLMC bias can often be eliminated entirely with a negligible additional cost. The new schemes also admit optimization criteria, which are easy to implement in practice.

Using General Scalarization Functions in Multicriteria Optimization with Stochastic Benchmarking Constraints

In multicriteria optimization, where decisions lead to vectors of outcomes, scalarization functions provide a fundamental tool for comparing outcome vectors by mapping them to scalars. Analogously, in a stochastic context scalarization functions map random outcome vectors to scalar-valued random variables, which can then be compared according to univariate preferences (based, e.g., on risk measures or stochastic dominance). While a plethora of scalarization functions are used in a deterministic context, offering a high degree of flexibility, the corresponding stochastic literature focuses overwhelmingly on linear scalarization. In “Optimization with Stochastic Preferences Based on a General Class of Scalarization Functions,” N. Noyan and G. Rudolf examine optimization problems in finite probability spaces with multivariate stochastic benchmarking constraints based on the new class of *min-biaffine* scalarization functions. This general class includes the majority of commonly used scalarizations from the deterministic literature, and the min-biaffine structure makes it possible to obtain tractable formulations by building on recent advances for solving similar problems with linear scalarization.

A New Method for Stochastic Derivative Estimation with Discontinuities

Stochastic derivative estimation plays a central role in both sensitivity analysis and gradient-based optimization. An

actively studied problem is to obtain an unbiased stochastic derivative estimator for a discontinuous sample performance. In “A New Unbiased Stochastic Derivative Estimator for Discontinuous Sample Performances with Structural Parameters,” Y. Peng, M. C. Fu, J.-Q. Hu, and B. Heidergott propose a new method to estimate the derivative for the expectation of a discontinuous sample performance. The proposed estimator is unbiased and has an analytical form. Many derivative estimation problems such as problems in statistical quality control, and financial derivatives evaluation, and problems containing probability constraints can be handled by the new method in a general framework. The theoretical results in this paper shed some light on how to generally deal with discontinuities in stochastic derivative estimation, and the developed methodology offers a systematic treatment for this long-standing problem.

Robust Newsvendor Solution is Invariant to Demand and Price Correlation

In “Profit Sharing Agreements in Decentralized Supply Chains: A Distributionally Robust Approach,” Q. Fu, C. K. Sim, and C.-P. Teo obtain a closed-form solution to a distributionally robust newsvendor problem when both demand and selling price are random with known means, variances. This extends the well-known Scarf’s Bound and generalizes various earlier results for different robust newsvendor problems. Interestingly, in the distributionally robust setting, the correlation between demand and selling price has no bearing on the order decision of the retailer. This result has numerous applications. For instance, it allows them to simplify the solution structure for a class of principal-agent model encountered in the design of profit sharing agreement, and recover the optimal robust selling price when the mean demand is a linear function of the selling price etc. Other applications are discussed in an online appendix of the paper.

Staffing to Reduce Excessive Delay in Queues with Time-Varying Arrivals

Queueing theory is a field driven by applications. But unfortunately there still remains a large gap between tractable theoretical studies and practical applications, such as call centers and health care systems, which have many realistic features (e.g., time-varying arrivals, customer abandonment, nonexponential service, etc.). In response to these changes, in “Staffing to Stabilize the Tail Probability of Delay in Service Systems with Time-Varying Demand,” Y. Liu develops analytic formulas to set the time-dependent number of servers for a queueing model having a nonstationary non-Poisson arrival process, customer abandonment, and nonexponential service distributions, in order to stabilize an important service-level indicator: the tail probability of delay (TPoD). TPoD, defined as the probability that the customer delay exceeds a customary delay target, has potential applications to customer contact centers (e.g., the 80-20 rule) and health care systems (e.g., the Canadian triage and acuity scale). To validate the effectiveness of the staffing prescription, the authors prove a limit theorem which shows that TPoD can be asymptotically stabilized at a desired constant probability target in a finite time interval as the system scale increases.

A Practicable Robust Optimization Scheme When Performance Decomposes Over Each Decision

In “A Practicable Robust Counterpart Formulation for Decomposable Functions: A Network Congestion Case Study,”

E. Delage, L. Gianoli, and B. Sansò explore a new robust optimization formulation that can circumvent computational difficulties in nonlinear optimization problems when cost decomposes as a sum of the cost of each decision. Their approach considers that the nonlinear cost function for each decision lies between a nominal and a maximum cost function, which can be calibrated based on historical data, while controlling the number of decisions for which the cost simultaneously reaches the most pessimistic value. They present how the approach can be employed to robustify packet routing on a telecommunication network with congestion. Their computational results obtained from a large number of problem instances of realistic size confirm that it is possible to identify robust solutions that outperform a deterministic approach and other natural approximation schemes in terms of both the amount of congestion and the risks of excessive congestion.

Capacity of Information Processing Systems

In “On the Capacity of Information Processing Systems,” L. Massoulié and K. Xu propose and analyze a family of information processing systems, where a finite set of experts or servers are employed to extract information about a stream of incoming jobs. Each job is associated with a hidden label drawn from some prior distribution. An inspection by an expert produces a noisy outcome that depends both on the job’s hidden label and the type of the expert, and occupies the expert for a finite time duration. A decision maker’s task is to dynamically assign inspections so that the resulting outcomes can be used to accurately recover the labels of all jobs, while keeping the system stable. Among their chief motivations are applications in crowd-sourcing, diagnostics, and experiment designs, where one wishes to efficiently learn the nature of a large number of items, using a finite pool of computational resources or human agents. They focus on the capacity of such an information processing system. Given a level of accuracy guarantee, they ask how many experts are needed in order to stabilize the system, and through what inspection architecture. Their main result provides an adaptive inspection policy that is asymptotically optimal in the following sense: the ratio between the required number of experts under their policy and the theoretical optimal converges to one, as the probability of error in label recovery tends to zero.

Private Sequential Decision-Making

In “Delay-Predictability Tradeoffs in Reaching a Secret Goal,” J. Tsitsiklis and K. Xu formulate a model of sequential decision-making, dubbed the Goal Prediction game, to study the extent to which an overseeing adversary can predict the final goal of an agent who tries to reach that goal quickly, through a sequence of intermediate actions. Their formulation is motivated by the increasing ubiquity of large-scale surveillance and data collection infrastructures, which can be used to predict an agent’s intentions and future actions, despite the agent’s desire for privacy. Their main result shows that with a carefully chosen agent strategy, the probability that the agent’s goal is correctly predicted by an adversary can be made inversely proportional to the time that the agent is willing to spend in reaching the goal, but cannot be made any smaller than that. Moreover, this characterization depends on the topology of the agent’s state space only through its diameter.