



Operations Research

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

In This Issue

To cite this article:

(2022) In This Issue. Operations Research 70(5):iii-vi. <https://doi.org/10.1287/opre.2022.2393>

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 © 2022, 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.2022.2393>

Copyright: © 2022 INFORMS

Predictive Analytics for Retail

High accuracy in demand prediction allows retailers to effectively manage their inventory and mitigate stock-outs and excess supply. A typical retail setting involves predicting the demand for hundreds of items simultaneously, some with abundant historical data and others with scarce data. In “Data Aggregation and Demand Prediction,” Cohen, Zhang, and Jiao propose a novel practical method, called *data aggregation with clustering* (DAC), which balances the tradeoff between data aggregation and model flexibility. DAC empowers retailers to predict demand while optimally identifying the features that should be estimated at the item, cluster, and aggregate levels. Theoretically, DAC yields a consistent estimate, along with improved prediction errors relative to the benchmark that estimates a different model for each item. Practically, DAC yields a higher demand prediction accuracy relative to many common benchmarks using a real data set from a large online retailer.

How Reactive Capacity Affects Product Quality and Firm Profitability

The COVID-19 pandemic exacerbates both supply and demand uncertainties, which have strategically influenced firms’ product quality and targeting decisions. For example, during the pandemic, automakers faced supply shortages for automotive chips because of the upstream suppliers’ limited parts inventory and production capacities, which have prompted the automakers to increase the quality (e.g., producing higher trims with more optional upgrade features) and price of their products to target fewer (high-valuation) consumers. In “Effects of Reactive Capacity on Product Quality and Firm Profitability in an Uncertain Market,” Jiang and Tian study the impacts of the supplier’s reactive capacity and demand uncertainty on product quality and firm profitability under pull contracts in the supply chain. They find that an increase in the supplier’s reactive capacity can lead to higher or lower equilibrium product quality, benefiting the retailer but potentially reducing the supplier’s profit. Interestingly, both the retailer and the supplier can be worse off with a higher probability for the high market state (with more high-valuation consumers). Further, a higher probability for the high market state can lead to lower product quality.

Avoiding Risk Concentration

One of the mantras of risk measurement is the avoidance of risk concentration. However, most formal approaches to the topic actually require more than this. In “Star-Shaped Risk Measures,” Castagnoli, Cattelan, Maccheroni, Tebaldi, and Wang study this property “in purity” for monetary risk measures. They show that it unites value at risk and convex risk measures, it is amenable to aggregation of opinions, and it leads to treatable optimization, thanks to a meaningful functional representation. They also show its ubiquitous presence in several fields of decision making under uncertainty.

Vehicle Routing Strategies with Overlapped Routes

Motivated by logistical problems faced by a large supply chain software company, in “A New Approach for Vehicle Routing with Stochastic Demand: Combining Route Assignment with Process Flexibility,” Ledvina, Qin, Simchi-Levi, and Wei study a vehicle routing problem where some routes are allowed to overlap. The paper proposes a class of simple and effective strategies to design overlapped routes with customer sharing, which combines ideas from the process flexibility in manufacturing and traditional vehicle routing literatures. Through theoretical analysis and numerical simulations, the paper illustrates the advantage of an overlapped routing strategy with a small amount of customer overlaps. In particular, it shows that such a strategy can provide consistent route assignments to drivers, while achieving a similar expected travel distance as the theoretical benchmark in the fully reoptimized setting. The strategy is in contrast to the traditional fixed routing strategy, which provides consistent route assignments to drivers, but incurs a much higher expected travel distance.

Online Job Scheduling with General Cost Functions

In “An Optimal Control Framework for Online Job Scheduling with General Cost Functions,” Etesami devises online speed-augmented competitive algorithms for minimizing the generalized completion time on a single and multiple unrelated machines for a very general class of cost functions. To that end, a novel optimal control formulation for the offline version of the problem is developed. Such a formulation allows using tools from optimal control theory, such as the minimum principle and Hamilton-Jacobi-Bellman equation, to set the dual variables as close as possible to the optimal dual variables and leverage them to design primal-dual online algorithms as an iterative application of the offline problem. The analysis can achieve state-of-the-art competitive ratios for several special cases and provide new competitive ratios which are the first in their settings. In particular, the analysis offers a principled method of estimating dual variables in a general setting of online job scheduling.

Endogenous Inverse Demand Functions

Buying or selling assets in a financial market impacts the prices upward or downward. Quantifying these price impacts is fundamental to many problems within finance (e.g., optimal liquidation and systemic risk). In “Endogenous Inverse Demand Functions,” Bichuch and Feinstein construct an equilibrium risk sharing problem that results in an endogenous inverse demand function. This is taken in contrast to the often assumed exogenous linear or exponential forms for the inverse demand functions. The authors determine sufficient joint properties for the financial market and assets to replicate these common exogenous forms. The properties of the general endogenous inverse demand functions, found via the equilibrium risk sharing

problem, are investigated; the authors deduce that price cross-impacts, which are often assumed to be zero, naturally arise in this equilibrium setting.

Optimizing Station-sizes in Bike-sharing Systems

Station-based bike-sharing systems have changed the urban landscape of major American cities including New York City, Boston, Washington DC, and Chicago. In these systems, stations are capacitated: a station's capacity is given by the number of docks it consists of. Each dock can either be filled with a bike or empty with the number of bikes at the station fluctuating throughout the day. Interestingly, in these systems, stockouts can occur both when stations are empty (making rentals impossible) and when stations are full (making returns impossible). Avoiding stockouts of either type is the most pressing operational concern faced by these systems. In "Minimizing Multi-modular Functions and Allocating Capacity in Bike-Sharing Systems," Freund, Henderson, and Shmoys work on capacity management in bike-sharing systems studies the strategic question of how to size stations in these systems in order to minimize stockouts of either kind. Their findings suggest that reallocating a small portion of the capacity within a system is a cost-efficient way to significantly improve service levels.

How Testing Agents Manipulate a Firm's Concept Selection Process

How should a firm structure its concept testing processes when testing efforts must be delegated to self-interested testing agents? In "Delegated Concept Testing in New Product Development," Schlapp and Schumacher analyze different configurations of testing processes and compare their relative benefits. The authors show that the optimal number of agents to be involved in a testing process critically depends on the costs of testing, the informational quality of test outcomes, and the agents' tendency to strategically prolong the testing process. The authors also provide guidelines on how to manage the testing process: which product concepts a firm should test (or not), and how it should structure the incentives of agents. Rather surprisingly, the authors also find that the delegation of testing activities can increase a firm's chances of successful concept selection. The article thus provides a comprehensive analysis of a key organizational issue that firms encounter during their concept testing processes, and it challenges some common decision rules.

Individualized Patient Monitoring Under Alarm Fatigue

Hospitals are rife with alarms, many of which are false. This leads to alarm fatigue, in which clinicians become desensitized and may inadvertently ignore real threats. In "Individualized Dynamic Patient Monitoring Under Alarm Fatigue," Piri, Huh, Shechter, and Hudson study the problem of personalizing alarm thresholds for vital signs at a hospital while considering the "boy who cried wolf" effect of false alarms. The authors create a model that learns patients' personal alarm thresholds during their hospital stay and updates their alarm settings dynamically. They formulate the problem as a partially observable Markov decision process. They provide structural properties of the optimal policy and perform a numerical case study based on clinical data from an intensive care unit. The authors show that dynamic methods of alarm settings that explicitly consider the feedback loop of false positives can significantly reduce patient harm when compared with current methods of alarm settings.

Revenue-Maximizing Auctions: A Bidder's Standpoint

A vast part of the Internet economy is powered by advertising, much of which is sold at auction. A key question for sellers is how to optimize the auction mechanism they use. Bidders, conversely, try to optimize their bidding strategy. Incentive compatible auctions are a sweet spot: theory predicts that it is in the bidders' interest to bid their values, making it relatively easy for them to bid optimally. However, as they learn bidders' value distributions, sellers can progressively optimize their mechanism and extract more revenue from bidders. In "Revenue-Maximizing Auctions: A Bidder's Standpoint," Nedelec, Calauzènes, Perchet, and El Karouic show that, in sharp contrast with most results in the academic literature, bidders should not be bidding their value in incentive compatible auctions when there is no commitment from the seller about using a fixed auction. They provide a mix of theoretical and numerical results and practical methods that can easily be deployed in practice.

Assortment Planning for Two-Sided Sequential Matching Markets

Matching platforms, like those for labor and dating, operate by presenting users with a set of recommended partners and allowing them to choose which of these to pursue a match with. When optimizing over the recommended partners shown to their users to maximize the number of matches, the platform faces a trade-off. On the one hand, increasing the number of recommended partners that a user sees increases the chances that she finds an acceptable one to contact (call this contacted user Alice). On the other hand, if Alice is shown too much, she may receive many requests; this may increase the chances she finds an acceptable partner (increasing matches), but may result in other conflicting requests being rejected, resulting in lost matches. In "Technical Note—Assortment Planning for Two-Sided Sequential Matching Markets," Ashlagi, Krishnaswamy, Makhijani, Saban, and Shiragur propose a novel stylized model that captures this tradeoff, and propose a simple algorithm that achieves a constant-factor approximation to the optimal number of matches.

Allocation with Weak Priorities and General Constraints

Allocating scarce resources with (weak) priority and complex constraints has many applications ranging from course allocation, and healthcare rationing to refugee resettlement. Its generality, however, mostly leads to impossibility results. In "Allocation with Weak Priorities and General Constraints," Lin, H. Nguyen, T. Nguyen, and Altinkemer offer a positive result with a mechanism based on a new concept called competitive stable equilibrium. This is a natural extension of competitive equilibrium with endowed budgets that accommodates weak priorities. Specifically, an agent only needs to pay for a resource if it belongs to the last tier among the agents currently consuming the resource. This concept applies to both frameworks of one-sided and two-sided markets in a unifying way and allows for a more flexible tie-breaking rule by giving agents different budgets. They empirically apply their mechanism to reassign season tickets to families in the presence of social distancing. Their simulation results show that their method outperforms existing ones in both efficiency and fairness measures.

Capacity-Constrained Assortment Optimization under Nested Logit Preferences

In "Technical Note—Approximation Schemes for Capacity-Constrained Assortment Optimization Under the Nested Logit Model," Segev proposes a carefully crafted dynamic programming approach for capacitated assortment optimization under

the nested logit model in its utmost generality, potentially including partially captured nests and possibly synergistic products. Specifically, the author shows that the optimal revenue can be efficiently approached within any degree of accuracy through synthesizing ideas related to continuous-state dynamic programming, state space discretization, and sensitivity analysis of modified revenue functions. These developments allow the author to devise the first fully polynomial-time approximation scheme in this context, thus resolving fundamental open questions posed in earlier literature.

A Novel and Promising Approximation for Network Revenue Management

In “Technical Note—Product-Based Approximate Linear Programs for Network Revenue Management,” R. Zhang, Samiedaluie, and D. Zhang propose a novel separable piecewise linear (SPL) approximation for the network revenue management problem. The coefficients of the proposed SPL approximation can be interpreted as each product’s revenue contribution to the value of each resource in a given period, which provides more granular information compared with the existing resource-based SPL approximation in the literature. The new approximation provides more flexibility for policy construction. Furthermore, the new approximation opens the opportunity to derive a set of valid inequalities to further improve the computational performance and achieve additional gains in the expected revenue. Computational experiments with instances of various network structures and parameters demonstrate its efficacy: The new approximation leads to bid-price policies generating higher expected revenues and demonstrates better performance in terms of both computational efficiency and numerical stability.

How to Rank Distributions by Knowing Only Their Means and Variances

In “Technical Note—Ranking Distributions When Only Means and Variances Are Known,” Müller, Scarsini, Tsetlin, and Winkler address the question of ranking distributions when only the first two moments—that is, means and variances—are known. This is important in decision making under uncertainty, with potential applications in economics, finance, statistics, and other areas. Previous results require some assumptions about the shape of the distributions, while this paper’s approach is to impose bounds on how much marginal utility can change, thus constraining risk preferences. Such a ranking is consistent with almost stochastic dominance and provides a new connection between the Sharpe and Omega ratios from finance.

Patrolling in Uniform or Undercover

Police cars patrol streets, security officers patrol museums, and soldiers patrol the perimeter of a military installation. If an attacker can hide and learn about the patrol pattern by observing patrollers going by, then how does the attacker time his attack to avoid getting caught? How effective is patrolling in uniform compared with patrolling undercover? In “Technical Note—Optimal Patrol of a Perimeter,” Lin formulates a two-person zero-sum Stackelberg game, in which a defender decides when to dispatch patrollers and how fast they move along the perimeter to maximize the probability of detecting an attack. Somewhat surprisingly, patrolling in uniform can be just as effective as patrolling undercover, as long as the patrol pattern is randomized properly.

Finding Minimum Volume Circumscribing Ellipsoids Using Generalized Copositive Programming

The problem of finding the Lowner–John ellipsoid (i.e., an ellipsoid of minimum volume covering a given set) has applications

in many areas of operations research, engineering, and statistics. However, this problem is NP hard for many of the sets of practical importance. In “Finding Minimum Volume Circumscribing Ellipsoids Using Generalized Copositive Programming,” Mittal and Hanasusanto present a generalized copositive program whose optimal solution produces the minimum volume ellipsoid. This generalized copositive program is used to obtain near-optimal semidefinite approximations, which are theoretically proven to provide ellipsoids of lower volume than the state-of-the-art methods. Their numerical tests demonstrate that the resulting approximations provides the best balance between accuracy and solution time.

A New Family of Valid Inequalities for Dantzig-Wolfe Reformulation of Mixed Integer Linear Programs

In “Consistency Cuts for Dantzig-Wolfe Reformulation,” Vinther Clausen, Lusby, and Ropke present a new family of valid inequalities to be applied to Dantzig-Wolfe reformulations with binary linking variables. They show that, for Dantzig-Wolfe reformulations of mixed integer linear programs that satisfy certain properties, it is enough to solve the linear programming relaxation of the Dantzig-Wolfe reformulation with all consistency cuts to obtain integer solutions. An example of this is the temporal knapsack problem; the effectiveness of the cuts is tested on a set of 200 instances of this problem, and the results are state-of-the-art solution times. For problems that do not satisfy these conditions, the cuts can still be used in a branch-and-cut-and-price framework. In order to show this, the cuts are applied to a set of generic mixed linear integer programs from the online library MIPLIB. These tests show the applicability of the cuts in general.

Regret Minimization Using Adjustable Robust Optimization Techniques

Although the stochastic optimization paradigm exploits probability theory to optimize the tradeoff between risk and returns, robust optimization has gained significant popularity by reducing computation requirements through the optimization of the worst-case scenario in a set. An appealing alternative to stochastic and robust optimization consists in optimizing decisions using the notion of regret. Although regret minimization models are generally perceived as leading to less conservative decisions than those produced by robust optimization, their numerical optimization is a real challenge in general. In “Adjustable Robust Optimization Reformulations of Two-Stage Worst-Case Regret Minimization Problems,” Poursoltani and Delage show how to reduce a two-stage worst-case absolute/relative regret minimization problem to a two-stage robust optimization one. This opens the way for taking advantage of recent advanced approximate and exact solution schemes for these hard problems. Their experiments corroborate the high-quality performance of affine decision rules as a popular polynomial-time approximation scheme, from which, under mild conditions, one can even expect exact regret-averse decisions.

Nonconvex Stochastic Optimization

Nonconvex stochastic optimization problems arise in many machine learning problems, including deep learning. The stochastic gradient Hamiltonian Monte Carlo (SGHMC) is a variant of stochastic gradients with a momentum method in which a controlled and properly scaled Gaussian noise is added to the stochastic gradients to steer the iterates toward a global minimum. SGHMC has shown empirical success in practice for solving nonconvex stochastic optimization problems. In “Global Convergence of Stochastic Gradient Hamiltonian Monte Carlo for Nonconvex Stochastic

Optimization: Nonasymptotic Performance Bounds and Momentum-Based Acceleration,” Gao, Gürbüzbalaban, and Zhu provide, for the first time, the finite-time performance bounds for the global convergence of SGHMC in the context of both population and empirical risk minimization problems and show that acceleration with momentum is possible in the context of global nonconvex stochastic optimization.

Online Linear Programming: Dual Convergence, New Algorithms, and Regret Bounds

In “Online Linear Programming: Dual Convergence, New Algorithms, and Regret Bounds”, Li and Ye study the online linear programming (LP) problem under a random input model. They first identify an equivalent form of the dual problem for the underlying LP that relates the LPs with a sample average approximation problem to a stochastic program. Then, the authors establish the convergence of the dual optimal solutions for the sequence of the scaled LPs solved throughout the online procedure. Based on these findings, they propose a new online LP algorithm, action-history-dependent learning algorithm, which improves the previous algorithm performance by taking into account the past input data and the past decisions/actions. An $O(\log n \log \log n)$ regret bound is derived (under a locally strong convexity and smoothness condition) for the proposed algorithm against the $O(\sqrt{n})$ bound for typical dual-price learning algorithms, where n is the number of decision variables.

An Exponentially Faster Algorithm for Submodular Maximization Under a Matroid Constraint

In “An Optimal Approximation for Submodular Maximization Under a Matroid Constraint in the Adaptive Complexity Model,” Balkanski, Rubinfeld, and Singer study the problem of submodular maximization under a matroid constraint. It is known since the 1970s that the greedy algorithm obtains a constant-factor approximation guarantee for this problem. Twelve years ago, a breakthrough result by Vondrák obtained the optimal $1 - 1/e$ approximation. Previous algorithms for this fundamental problem all have linear parallel runtime, which was considered impossible to accelerate until recently. The main contribution of this paper is a novel algorithm that provides an exponential speedup in the parallel runtime of submodular maximization under a matroid constraint, without loss in the approximation guarantee.

A Model for New Venture Creation

New ventures go through multiple stages: In the early stage, there is a business concept and preliminary evidence supporting the concept. In later stages, there are revenues and sales. In each stage, there are usually milestones for the venture to meet in order for investors to provide additional funding. Otherwise, the venture is abandoned. The entrepreneur can engage in a set of costly activities that aim to create value and reach the appropriate milestones. In “New Venture Creation: A Drift-Variance Diffusion Control Model,” Wang and Zenios develop a framework for new venture creation. The authors provide theoretical guidance on the optimal policy, which is relatively simple to describe. Their analysis reveals a trade-off between how costly an activity is and how much upside

potential the activity generates, and their result shows how a new venture creator can manage that trade-off.

Extremizing and Antiextremizing in Bayesian Ensembles of Binary-Event Forecasts

Many organizations combine forecasts of probabilities of binary events to support critical business decisions, such as the approval of credit or the recommendation of a drug. To aggregate individual probabilities, the authors offer a new method based on Bayesian principles that can help identify why and when combined probabilities need to be extremized. Extremizing is typically viewed as shifting the average probability farther from one half; they emphasize that it is more suitable to define extremizing as shifting it farther from the base rate. In “Extremizing and Antiextremizing in Bayesian Ensembles of Binary-Event Forecasts,” Lichtendahl Jr., Grushka-Cockayne, Jose, and Winkler introduce the notion of antiextremizing, cases in which it might be beneficial to make average probabilities less extreme. Analytically, they find that their Bayesian ensembles often extremize the average forecast but sometimes antiextremize instead. On several publicly available data sets, they demonstrate that our Bayesian ensemble performs well and antiextremizes anywhere from 18% to 73% of the cases. Antiextremizing is required more often when there is bracketing with respect to the base rate among the probabilities being aggregated than with no bracketing.

Allocating Resources Across Systems Coupled by Shared Information

Many sequential decision problems involve repeatedly allocating a limited resource across subsystems that are jointly affected by randomly evolving exogenous factors. For example, in adaptive clinical trials, a decision maker needs to allocate patients to treatments in an effort to learn about the efficacy of treatments, but the number of available patients may vary randomly over time. In capital budgeting problems, firms may allocate resources to conduct R&D on new products, but funding budgets may evolve randomly. In many inventory management problems, firms need to allocate limited production capacity to satisfy uncertain demands at multiple locations, and these demands may be correlated due to vagaries in shared market conditions. In “Dynamic Programs with Shared Resources and Signals: Dynamic Fluid Policies and Asymptotic Optimality,” Brown and Zhang develop a model involving “shared resources and signals” that captures these and potentially many other applications. The framework is naturally described as a stochastic dynamic program, but this problem is quite difficult to solve. They develop an approximation method based on a “dynamic fluid relaxation”: In this approximation, the subsystem state evolution is approximated by a deterministic fluid model, but the exogenous states (the signals) retain their stochastic evolution. The authors develop an algorithm for solving the dynamic fluid relaxation. They analyze the corresponding feasible policies and performance bounds from the dynamic fluid relaxation and show that these are asymptotically optimal as the number of subsystems grows large. They show that competing state-of-the-art approaches used in the literature on weakly coupled dynamic programs in general fail to provide asymptotic optimality. Finally, the authors illustrate the approach on the aforementioned dynamic capital budgeting and multilocation inventory management problems.