



Operations Research

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

In This Issue

To cite this article:

(2023) In This Issue. Operations Research 71(5):iii-vii. <https://doi.org/10.1287/opre.2023.inthisissue.v71.n5>

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 © 2023, 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.2023.inthisissue.v71.n5>

Copyright: © 2023 INFORMS

The Importance of Considering System Congestion When Acquiring Used Items for Remanufacturing

The condition of used items obtained by remanufacturers often varies widely and may only be truly known after an inspection. In “Technical Note—Optimal Procurement in Remanufacturing Systems with Uncertain Used-Item Condition,” Nadar, Akan, Debo, and Scheller-Wolf study the procurement problem by taking into account this pre-inspection uncertainty as well as the quantities of available used items of different quality levels. The authors provide a novel optimal policy structure for a Markov decision process representation of this problem. They prove that the optimal acquisition decision depends on the system congestion level — the total remanufacturing load weighted by the quality levels. Their results show that it becomes less desirable to acquire a used item as the number of available items or the quality level of any available item increases. To prove their results, the authors introduce a new functional characterization that is a relaxation of supermodularity and discrete convexity.

The Pricing Problem of Vehicle Routing Problems Is Strongly NP-Hard

The pricing problem of vehicle routing problems is strongly NP hard. Many algorithms for finding optimal solutions to various vehicle routing problems rely on a subroutine to solve a so-called pricing problem of a set partitioning formulation. Only exponential time algorithms are known for these pricing problems. Therefore, the set partitioning formulation is oftentimes relaxed at the expense of weaker LP bounds, so that the resulting pricing problem can now be solved using a pseudopolynomial or polynomial time algorithm. In “Technical Note—The Complexity of the Pricing Problem of the Set Partitioning Formulation of Vehicle Routing Problems,” Spliet proves that the pricing problem is strongly NP hard for many different vehicle routing problems. This means that, unless $P = NP$, no pseudopolynomial or polynomial time algorithm exists for the pricing problem, justifying the common use of relaxations.

The Impact of the Prohibition of Trade-Through

In “Does the Prohibition of Trade-Through Hurt Liquidity Demanders?,” Chen, Gao, and Kou explore the effects of the order protect rule (OPR) in the United States, which prohibits any trade-through in the stock market that does not execute at the best possible price among fast trading venues. The study found that, whereas trade-throughs may offer flexible trading strategies that can benefit liquidity demanders, the benefits are insignificant for most cases, especially for small trades and stocks with fast resilience. The authors suggest that the current separate regulations for fast and slow venues may be extended to differentiate stocks with fast and slow resilience speeds. The results are supported by a case study using the data from the Australian Securities Exchange. Overall, this study supports the regulation of the OPR.

When are Coupled Networked Markets Prone to Sudden Collapse?

Many important financial instruments are traded in large networked markets, whose fragility played a role in the Great Financial Crisis. In “Exit Spirals in Coupled Networked Markets,” Aymanns, Georg, and Golub study strategic market participation to show that fragility is exacerbated when several networked markets are *coupled*, so that activity in one market facilitates participation in others. Such coupling occurs, for example, when participants use relationships in one market to borrow funds to trade in others. The authors ask when market participation is prone to sudden collapse in an *exit spiral*. They explain when such exit spirals emerge and find that market coupling is a pervasive cause of fragility, creating exit spirals even between networks that are individually robust. The fragility of coupled markets can be mitigated if one of two coupled markets becomes centralized or if links become more correlated across markets.

Operations Management in a Repair Facility

Capital goods are machines or products that are used in the production of goods or service deliveries. Maintenance service providers keep spare parts on stock so that, whenever a critical component of a capital good breaks down, the broken part is replaced with a spare part to prevent long and costly downtime. In “Joint Inventory and Scheduling Control in a Repair Facility,” Özkan and van Houtum study inventory and repair scheduling decisions of a maintenance service provider for repairable capital goods. In case of a stock-out, the service provider should decide whether to back-order the demand or execute an emergency repair, which is an urgent but expensive repair operation for a broken part. The authors derive a simple and intuitive decision rule stating if the emergency repairs are necessary to achieve a close-to-optimal system performance. Moreover, they propose a simple, intuitive, and easy-to-implement heuristic control policy that performs well in numerical experiments.

Demand Learning via Exploration Boosts

In the Bayesian newsvendor problem, it is known that the optimal decision is always greater than or equal to the myopic decision. As a result, the optimal decision can be expressed as the sum of the myopic decision plus a nonnegative “exploration boost.” In “Bayesian Inventory Control: Accelerated Demand Learning via Exploration Boosts,” Chuang and Kim characterize the form of the exploration boost in terms of basic statistical measures of uncertainty. This characterization expresses in clear terms the way in which the statistical learning and inventory control are jointly optimized; when there is a high degree of parameter uncertainty, inventory levels are boosted to induce a higher chance of observing more sales data to more quickly resolve statistical uncertainty, and as parameter uncertainty resolves, the exploration boost is reduced.

A Nonparametric Algorithm for Optimal Stopping Based on Robust Optimization

Optimal stopping is a fundamental class of stochastic dynamic optimization problems with numerous applications in finance and operations management. In “A Nonparametric Algorithm for Optimal Stopping Based on Robust Optimization,” Sturt introduces an approach based on robust optimization for computing near-optimal solutions for computationally demanding stochastic optimal stopping problems with known probability distributions. Through this new use of robust optimization, the paper develops new algorithms for solving stochastic optimal stopping problems that are practical and can outperform state-of-the-art simulation-based algorithms in the context of pricing high-dimensional financial options.

Dangers of Ignoring Market Friction When Leveraging Financial Portfolios

Portfolio leveraging is a standard industry practice to target higher fund returns, for example, in risk parity asset allocation. However, the existing models of optimal bet sizing fail to integrate analytically the impact on leveraged portfolio selection resulting from market liquidity issues. In “Optimal Leveraged Portfolio Selection Under Quasi-Elastic Market Impact”, Edirisinghe, Chen, and Jeong consider a market in which both temporary and permanent impact on trading prices are present with the former impact being sufficiently large relative to the latter. Their analytical conclusions, supported by a case study that uses even relatively more liquid U.S. exchange-traded fund assets, demonstrate that fund managers are ill advised to ignore market friction when leveraging to achieve target higher returns. Not only risk-adjusted returns significantly deteriorate, but also those losses become steeper when setting higher targets requiring increased levels of leverage. Moreover, leverage-constrained and less risk-averse investors ignoring liquidity costs ex ante face the most losses in expected utility ex post.

A Better and More Reliable Exact Algorithm for Capacitated Location-Routing

In “Nonrobust Strong Knapsack Cuts for Capacitated Location Routing and Related Problems,” Liguori, Mahjoub, Marques, Sadykov, and Uchoa present a novel BCP algorithm for the CLRP and for other problems that share a nested knapsack structure. It outperforms existing exact algorithms in the literature, making it a powerful tool for solving instances with a large number of depot locations. A key methodological contribution is the introduction of RLKCs, a family of nonrobust cuts derived from the “outer” knapsack constraints. These cuts are strong in the sense that they contain all facets of the master knapsack polytope, dominating the cover cuts by Dabia S, Ropke S, Van Woensel T (2019) Cover inequalities for a vehicle routing problem with time windows and shifts. *Transportation Sci.* 53(5):1354–1371. By exploring their monotonicity and superadditivity properties, it is possible to adapt the labeling algorithm for handling RLKCs efficiently. The overall positive impact of RLKCs on the BCP performance varies depending on the problem and instance characteristics, but they have proven particularly effective for CLRP instances with tight depot capacities, making the final BCP algorithm more reliable.

Pricing in On-Demand and One-Way Vehicle-Sharing Networks

In “Technical Note–Pricing in On-Demand and One-Way Vehicle-Sharing Networks”, Benjaafar and Shen introduce a new method for evaluating the performance of static pricing in one-way vehicle-sharing systems. Our approach, based on a well-known recursive relationship, leads to a series of increasingly

tight bounds on the performance of the static pricing policy. These bounds are valid for systems with multiple locations, nonzero travel times, and an arbitrary number of vehicles. They also apply to systems where the static pricing policy does not lead to a fully connected network. The author’s method results in a family of asymptotically optimal static pricing policies that improve upon previous results in the literature. The approach applies to the case of a single location and yields a bound that is at least as tight as the best known bound.

The Role of Capacity Constraint in Dynamic Mechanisms

When the number of tasks is large, how should a firm design reward and penalty schemes to incentivize its employees? In “Technical Note–Dynamic Mechanism Design with Capacity Constraint,” He studies the role of capacity constraint in a project assignment problem, where a principal needs to assign multiple tasks to an agent. The author fully characterizes the optimal mechanism via a sequence of deadlines. This characterization is used to show that the presence of the capacity constraint reduces the principal’s payoff and delays the completion of projects. It further illustrates that the widely adopted no-capacity constraint framework may provide inaccurate results in various dynamic problems.

Information Disclosure in Platforms: Optimizing Volumes and Prices Can Harm Consumers

Two-sided platforms reduce frictions and facilitate trade in many sectors, and, in doing so, they increasingly collect and process data about supply and demand. In “Strategic Release of Information in Platforms: Entry, Competition, and Welfare”, Bimpikis and Mantegazza show that platforms can increase their profits by strategically disclosing (coarse) information that they collected about demand to the supply side. However, this practice may also adversely impact the welfare of consumers. By designing its information disclosure policy, a platform can influence the entry and pricing decisions of its potential suppliers. In general, it is optimal for platforms to disclose their information only partially to either “nudge” entry when it is costly for suppliers to join or discourage it when suppliers do not have valuable outside options. On the other hand, consumers may end up worse off due to higher prices compared with when the platform refrains from sharing any demand information.

Adaptive Discretization in Reinforcement Learning

Performance guarantees for RL algorithms are typically for worst case instances, which are pathological by design and not observed in meaningful applications. Moreover, many domains (such as computer systems and networking applications) have large state-action spaces and require algorithms to execute with low latency. This phenomenon highlights a trifecta of goals for practical RL algorithms: low sample, storage, and computational complexity. In “Adaptive Discretization in Online Reinforcement Learning”, Sinclair, Banerjee, and Yu develop an algorithmic framework for nonparametric RL with data-driven adaptive discretization. The author’s framework has probably better sample, storage, and computational complexity than uniform discretization or kernel regression methods. Moreover, we highlight how the performance guarantees are min-max optimal with respect to a novel instance-specific complexity measure that captures structure in facility location and newsvendor models.

Novel Insights on Proactive Care: The Case of Transfers to the Intensive Care Units

Patients whose transfer to the intensive care unit (ICU) unplanned are prone to higher mortality rates. In “Robustness of Proactive Intensive Care Unit Transfer Policies,” Grand-Clément, Chan, Goyal, and Escobar study the problem of finding robust patient transfer policies to the ICU, which account for uncertainty in statistical estimates because of data limitations when optimizing to improve overall patient care. Under general assumptions, it is shown that an optimal transfer policy has a threshold structure. A robust policy also has a threshold structure, and it is more aggressive in transferring patients than the optimal nominal policy, which does not consider parameter uncertainty. The sensitivity of various hospital metrics to small changes in the parameters is highlighted using a data set of close to 300,000 hospitalizations at 21 Kaiser Permanente Northern California hospitals. This work provides useful insights into the impact of parameter uncertainty on deriving simple policies for proactive ICU transfer that have strong empirical performance and theoretical guarantees.

Optimizing Mobile Food Pantry Operations Under Demand Uncertainty

Managing complex systems often involves making trade-offs between different objectives. A common example is seeking fairness guarantees in sequential resource allocation problems. For example, mobile food pantries are tasked with allocating resources under demand uncertainty with the goal of simultaneously minimizing inefficiency (leftover resources) and envy (deviations in allocations). In “Sequential Fair Allocation: Achieving the Optimal Envy-Efficiency Trade-off Curve,” Sinclair, Jain, Banerjee, and Yu tackle a problem established from a partnership with the Food Bank of the Southern Tier in optimizing their mobile food-pantry operations. They provide an exact characterization of the achievable (envy, efficiency) pairs, showing that any algorithm achieving low envy must suffer from poor inefficiency and vice versa. The authors complement this exact characterization with a simple algorithm capable of achieving any desired point along the trade-off curve.

The Price of Incentivizing Exploration: A Characterization via Thompson Sampling and Sample Complexity

In “The Price of Incentivizing Exploration: A Characterization via Thompson Sampling and Sample Complexity,” Sellke and Slikvins consider “incentivized exploration”: a version of multiarmed bandits where the choice of arms is controlled by self-interested agents. The algorithm can only issue recommendations and needs to incentivize the agents to explore, even though they prefer to exploit. The algorithm controls the flow of information and relies on information asymmetry to create incentives. This line of work, motivated by misaligned incentives in recommendation systems, has received considerable attention in the “economics and computation” community. The authors focus on the “price of incentives”: the loss in performance, broadly construed, incurred for the sake of incentive compatibility. They prove that Thompson sampling, a standard bandit algorithm, is incentive compatible if initialized with sufficiently many data points. The performance loss because of incentives is, therefore, limited to the initial rounds when these data points are collected. The problem is largely reduced to that of sample complexity. How many rounds are needed to collect even one sample of each arm? The authors zoom in on this question, which is perhaps the most basic question one could ask about incentivized exploration. They characterize the dependence on agents’ beliefs and the number of arms (which was essentially ignored in prior work), providing matching upper

and lower bounds. Typically, the optimal sample complexity is polynomial in the number of arms and exponential in the “strength of beliefs.”

Using Hospital Admission Predictions at Triage for Improving Patient Length of Stay in Emergency Departments

In emergency departments (EDs), one of the major reasons behind long waiting times and crowding overall is the time it takes to move admitted patients from the ED to an appropriate bed in the main hospital. In “Using Hospital Admission Predictions at Triage for Improving Patient Length of Stay in Emergency Departments,” Chen, Argon, Bohrmann, Linthicum, Lopiano, Mehrotra, Travers, and Ziya develop a methodology that can be used to shorten these times by predicting the likelihood of admission for each patient at the time of triage and starting the process of identifying a suitable hospital bed and making preparations for the patient’s eventual transfer to the bed right away if the predicted probability of admission is deemed high enough. A simulation study suggests that the proposed methodology, particularly when it takes into account ED census levels, has the potential to shorten average waiting times in the ED without leading to too many false early bed requests.

A Simple and Provably Good Inventory Policy Led to a Significant Reduction in Platelet Outdates

Platelets are critical blood products. The management of platelet inventory is particularly challenging because of its perishable nature with a short shelf life. In “Inventory Sharing for Perishable Products: Application Platelet Inventory Management in Hospital Blood Banks,” Zhang, Ayer, White, Bodeker, and Roback study how the wastage of platelets and, more broadly, perishable products can be reduced through inventory sharing. The authors derive structural results on the optimal ordering and transshipment policies for a two-location perishable inventory system and prove that an easy-to-implement myopic transshipment policy is optimal for a few special cases relevant to practice and serves as a lower bound on the optimal transshipment policy for more general settings. This policy has been successfully implemented by the Emory University Hospital System, which led to a reduction of approximately 20% in platelet outdates. Moreover, the author’s also shed light on how the presence of inventory sharing affects the optimal ordering quantities and provides insights that significantly depart from existing findings for nonperishable inventory systems.

Tight Guarantees for Static Threshold Policies in the Prophet Secretary Problem

When there are only k copies of an item (e.g. COVID-19 vaccines) and more than k people who want them, how should the threshold for eligibility be determined? In “Tight Guarantees for Static Threshold Policies in the Prophet Secretary Problem”, Arnosti and Ma study the performance of static threshold policies, which fix the same eligibility threshold throughout the time horizon. Under the assumption that population needs for the item are drawn independently and that people arrive in a random order, the authors show that the performance guarantees for their static threshold policies are best-possible. Moreover, they establish different ways of setting the static threshold to achieve this guarantee, that do not require knowing precise needs (i.e. pre-vaccination mortality risk) for each person and instead only proxies (e.g. age and underlying medical conditions).

Strict Priority Policies in a Stochastic System with Abandonment

In “Technical Note—Stochastic Scheduling with Abandonment: Necessary and Sufficient Conditions for the Optimality of a Strict Priority Policy,” Chen, Gayon, and Lemaire consider a stochastic scheduling problem in which jobs abandon when their waiting time exceeds their lifetime. Such a problem arises, for example, in call centers or emergency systems. It is known that the optimal policy is a strict priority policy under some sets of conditions. The authors provide the first set of necessary and sufficient conditions for a problem with two types of jobs. They also provide conjectures to guide toward generalizations of the proposed conditions.

Dual Approach for Two-Stage Robust Nonlinear Optimization

In “Technical Note—Dual Approach for Two-Stage Robust Nonlinear Optimization,” de Ruiter, Zhen, and den Hertog study adjustable robust minimization problems where the objective or constraints depend in a convex way on the adjustable variables. They reformulate the original adjustable robust nonlinear problem with a polyhedral uncertainty set into an equivalent adjustable robust linear problem, for which all existing approaches for adjustable robust linear problems can be used. The reformulation is obtained by first dualizing over the adjustable variables and then over the uncertain parameters. The polyhedral structure of the uncertainty set then appears in the linear constraints of the dualized problem, and the nonlinear functions of the adjustable variables in the original problem appear in the uncertainty set of the dualized problem. The authors show how to recover linear decision rules to the original primal problem and how to generate bounds on its optimal objective value.

A Comprehensive Complexity Study of Covariate Balancing Problems

In an observational study there are two disjointed groups of samples, one of treatment samples and the other of control samples. Each of the samples is characterized by several observed covariates. Covariate balancing problems arise when estimating causal effects using observational data. It is desirable to replicate a randomized experiment by obtaining treatment and control groups with similar covariate distributions. Even though covariate balancing problems have been studied extensively, the complexity status of many variants has not been established. In “Algorithms and Complexities of Matching Variants in Covariate Balancing,” by Hochbaum, Levin, and Rao some of their results demonstrate that 2-covariate balancing problems are polynomial time solvable, whereas almost all problems are hard for three or more covariates. These results have practical implications, such as justifying the use of implicit enumeration techniques or heuristics for the hard cases. A new approach suggested by their results is to use the 2-covariate polynomial cases and relax the problem by aggregating covariates into two sets to be solved efficiently.

Solving Convex Discrete Optimization via Simulation via Stochastic Search Algorithms

Many decision-making problems in operations research and management science require the optimization of large-scale complex stochastic systems with high-dimensional discrete decision variables. For a number of applications, the objective function exhibits convexity in the discrete decision variables, or the problem can be transformed into a convex one. In “Gradient-Based Algorithms for Convex Discrete Optimization via Simulation,” Zhang, Zheng, and Lavaei propose provably efficient simulation–optimization algorithms for general convex optimization via simulation problems with high-dimensional discrete decision space. By utilizing

the convex structure, the developed algorithms demonstrate a polynomial dependence on the dimension and scale of the decision space. In addition, information-theoretic lower bounds are derived to show the limit of algorithm design. In the case when possibly biased gradient estimators are available, they propose simulation–optimization algorithms to achieve optimality guarantees with a reduced dependence on the dimension.

Mixed-Integer Programming Formulations for Nonconvex Piecewise Linear Functions

Piecewise linear functions are deceptively simple structures that are nonetheless capable of approximating complex nonlinear behavior. As such, they have been adopted throughout operations research and engineering to approximate nonlinear structures in optimization problems which would otherwise render the problem extremely difficult to solve. In “Nonconvex Piecewise Linear Functions: Advanced Formulations and Simple Modeling Tools,” Huchette and Vielma derive new mixed-integer programming (MIP) formulations for embedding low-dimensional nonconvex piecewise linear functions in optimization models. These formulations computationally outperform the crowded field of existing approaches in a number of regimes of interest. As these formulations are derived using recently developed machinery that produce highly performant, but uninterpretable, formulations, the authors showcase the utility of high-level modeling tools by presenting PiecewiseLinearOpt.jl, an extension to the popular JuMP optimization modeling language that implements a host of MIP formulations for piecewise linear function in a single, easy-to-use interface.

Your Choices Are Justified by Benjamin Franklin!

Is it possible to capture every nuance of the rich human choice behavior via a structured model? In “Every Choice Function Is Pro-Con Rationalizable,” Dogan and Yildiz introduce and analyze the *pro-con model* that is inspired by Franklin’s *prudential algebra*. Consider an agent who is endowed with two sets of orderings: pro- and con-orderings. For each choice set, if an alternative is the top-ranked by a pro-ordering (con-ordering), then this is a pro (con) for choosing that alternative. The alternative with more pros than cons is chosen from each choice set. Each ordering may have a weight reflecting its salience. In this case, the probability that an alternative is chosen equals the difference between the total weights of its pros and cons. Although, this is an additive model similar to the random utility model with structurally invariant primitives, authors show that every (random) choice function is (random) pro-con rational. Their technique requires a generalization of Ford-Fulkerson theorem. The connection between the random model and its deterministic counterpart demonstrates a fruitful use of classical integer programming techniques in choice theory.

Data-Driven Composite Optimization

Optimization under uncertainty and risk is ubiquitous in business, engineering, and finance. Typically, we use observed or simulated data in our decision models, which aim to control risk, and result in composite risk functionals. In “Stability and Sample-Based Approximations of Composite Stochastic Optimization Problems,” Dentcheva, Lin, and Penev address the stability of the decision problems when the composite risk functionals are subjected to measure perturbations at multiple levels of potentially different nature. They analyze data-driven formulations with empirical or smoothing estimators such as kernels or wavelets applied to some or to all functions of the compositions and establish laws of large numbers and consistency of the optimal values and solutions. This is the first study to propose and analyze smoothing in data-driven composite optimization problems. It is

shown that kernel-based and wavelet estimation provide less biased estimation of the risk compared with the empirical plug-in estimators under some assumptions.

Inducing Continuous Effort from Multiple Agents Through Resource Allocation Contracts

On online platforms, goods/services/content providers, also known as agents, introduce adverse events. The frequency of these events depends on each agent's effort level. In "Efficient Resource Allocation Contracts to Reduce Adverse Events," Liang, Sun, Tang and Zhang study continuous-time dynamic contracts that utilize resource allocation and monetary transfers to induce agents to exert effort and reduce the arrival rate of adverse events. They devise an iterative algorithm that characterizes and calculates such contracts and specify the profit-maximizing contract for the platform, also known as the principal. In contrast to the single-agent case, in which efficiency is not achievable, they show that efficient and incentive-compatible contracts, which allocate all resource and induce agents to exert constant effort, generally exist with two or more agents. Additionally, the authors provide

efficient and incentive-compatible dynamic contracts that can be expressed in closed-form and are therefore easy to understand and implement in practice.

Sampling Graphs with Partial Information

The problem of simulating graphs (networks) subject to constraints has been studied extensively across several areas. Applications of this problem include modeling inter-bank financial networks, predator-prey ecological graphs, contingency tables, and even studying larger networks such as the Internet. In "Maximum Entropy Distributions with Applications to Graph Simulation," Glasserman and Lelo de Larrea study the more general problem of sampling uniformly from product sets under linear constraints, which includes simulating bipartite, directed, and undirected graphs with given degree sequences. For this purpose, they consider two suitable probability distributions: one that maximizes the entropy of the system, and another that maximizes the minimum probability of hitting the desired target set. Although apparently different, the authors provide conditions under which both distributions coincide. In addition, they propose a simple sequential algorithm to sample medium-sized graphs with fixed degrees.