



## Operations Research

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

### In This Issue

To cite this article:

(2024) In This Issue. Operations Research 72(2):iii-vi. <https://doi.org/10.1287/opre.2024.inthisissue.v72.n2>

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](mailto: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 © 2024, 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.2024.inthisissue.v72.n2>

Copyright: © 2024 INFORMS

### New Framework Unifies Capacitated Vehicle Routing Problem Under Risk and Ambiguity

In “A Unifying Framework for the Capacitated Vehicle Routing Problem Under Risk and Ambiguity,” Ghosal, Ho, and Wiesemann propose a comprehensive and versatile framework that addresses the challenges posed by demand uncertainty in the capacitated vehicle routing problem (CVRP). This framework is able to consider and incorporate various risk measures, satisficing measures, and disutility functions, providing a unified approach to tackle different variants of the CVRP under uncertainty. By offering a unified treatment of the CVRP under risk and ambiguity, this framework enables decision makers to optimize routing decisions, accounting for the associated risks and uncertainties. One of the key advantages of this framework is its practicality for implementations. The authors demonstrate that an existing branch-and-cut algorithm can effectively solve all variants of the uncertainty-affected CVRP with minimal modifications. This scalability and adaptability make the framework applicable in practical settings.

### Stochastic Liquidity as a Proxy for Nonlinear Price Impact

Trading costs play a central role in designing and implementing quantitative trading strategies. To quantify trading costs, optimal execution and trading algorithms rely on price impact models, such as the propagator model. Empirically, price impact is concave in trade sizes, leading to nonlinear models for which optimization problems are intractable and even qualitative properties, such as price manipulation, are poorly understood. In “Stochastic Liquidity as a Proxy for Nonlinear Price Impact,” Muhle-Karbe, Wang, and Webster show that, in the diffusion limit of small and frequent orders, the nonlinear model converges to a tractable linear model. In this high-frequency limit, a stochastic liquidity parameter approximates the original impact function’s nonlinearity. This allows us to derive simple formulas for optimal trading strategies and sharp conditions on market volumes to rule out price manipulation. A detailed empirical study using high-frequency limit-order data illustrates the practical performance of the theoretical results.

### Robust Queue Inference from Waiting Times

Modeling and decision making for queueing systems have been one of fundamental topics in operations research. For these problems, uncertainty models are established by estimation of key parameters such as expected interarrival and service times. In practice, however, their distributions are unknown, and decision makers are only given historical records of waiting times, which contain relevant but indirect information on the uncertainties. Their complex temporal dependence on the queueing dynamics and the absence of distributional information on the model primitives render estimation of queueing systems remarkably challenging. In “Robust Queue Inference from Waiting Times” by Bandi, Han, and Proskynitopoulos, a new inference framework based on

robust optimization is proposed to estimate unknown service times from waiting time observations. This new framework allows data-driven, distribution-free estimation on unknown model primitives by solving tractable optimization problems.

### Dynamic Pricing and Learning with Discounting

Learning algorithms can take a substantial amount of time to converge, thereby raising the need to understand the role of discounting in learning. In “Technical Note—Dynamic Pricing and Learning with Discounting,” Feng, Dawande, Janakiraman, and Qi examine the impact of discounting on learning by examining two classic dynamic-pricing and learning problems studied in [Broder J, Rusmevichientong P (2012) Dynamic pricing under a general parametric choice model. *Oper. Res.* 60(4):965–980] and [Keskin NB, Zeevi A (2014) Dynamic pricing with an unknown demand model: Asymptotically optimal semi-myopic policies. *Oper. Res.* 62(5):1142–1167]. In both settings, the retailer initially does not know the parameters of the demand model. Given a discount factor, the retailer’s objective is to determine a pricing policy to maximize the discounted revenue over a selling horizon. The authors establish lower bounds on the regret under any policy and propose new asymptotically optimal policies that take the discount factor into consideration. They numerically examine the regret under the proposed policies and the existing policies in the aforementioned two papers.

### Best Principal Submatrix Selection for the Maximum Entropy Sampling Problem: Scalable Algorithms and Performance Guarantees

In “Best Principal Submatrix Selection for the Maximum Entropy Sampling Problem: Scalable Algorithms and Performance Guarantees,” Li and Xie study a classic maximum entropy sampling problem (MESP), which aims to select the most informative principal submatrix of a prespecified size from a covariance matrix. MESP is widely applied to many areas, including healthcare, power systems, manufacturing, and data science. By investigating its Lagrangian dual and primal characterization, the authors derive a novel convex integer program for MESP and show that its continuous relaxation yields a near-optimal solution. The results motivated them to study efficient approximation algorithms and develop their approximation bounds for MESP, which improves the best known one in the literature.

### Consumers May Benefit from Firms Tracking and Exploiting Their Data

In today’s economy, firms routinely collect, track, and leverage consumer data to make decisions. Although the effects of these practices on consumers are complex and context-dependent, one may expect that that consumers would be disadvantaged if their data are used for a purpose that does not create value for them, such as personalized pricing. In “Data Tracking Under Competition,” Bimpikis, Morgenstern, and Saban develop a game-theoretic model to explore how technologies that allow firms to

use consumer data for price discrimination affect market outcomes. Perhaps counter to intuition, they find that data-tracking practices may actually increase consumer surplus, even if consumers do not develop privacy concerns associated with disclosing their data and act myopically. However, this only occurs when multiple firms compete for consumers' purchases. The study highlights the crucial role of competition in determining the benefits that consumers may derive from the data they generate and underscores the importance of considering competition in debates about the economic consequences of data-tracking practices.

### Does Increasing the Testing Capacity Always Reduce the Spread of Infection?

A standard policy to reduce the spread of infection is increasing the use of testing that enables the isolation of infected individuals, slowing down the infection. In "Testing, Voluntary Social Distancing, and the Spread of an Infection," Acemoglu, Makhdomi, Malekian, and Ozdaglar argue the possibility of unintended behavioral consequences from increased testing: greater testing reduces voluntary social distancing or increases social activity, exacerbating the spread of the virus. They show that the effect of testing on infections is nonmonotone. This nonmonotonicity also implies that the optimal testing policy may leave some of the testing capacity of society unused and that increasing testing should be used together with mandatory social distancing to reduce the spread of infection.

### Latent Agents in Networks: Estimation and Targeting

In "Latent Agents in Networks: Estimation and Targeting," Ata, Belloni, and Candogan address the problem of estimating network effects in a setting in which data only on a subset of agents is available. In this setting, the observable agents influence each other's outcomes both directly and indirectly through their influence on the latent agents. Even in sparse networks, the combination of direct and indirect network effects yields a non-sparse influence structure that makes estimation challenging. The authors overcome this challenge and provide an estimation algorithm that performs well in high-dimensional settings. They also establish convergence rates for their proposed estimator and show that their performance guarantees are valid for a large class of networks. Finally, the authors demonstrate the application of their algorithm to a targeted advertising problem, in which it can be used to obtain asymptotically optimal advertising decisions despite the presence of latent agents.

### California Utility Firm Implements Innovative Model, Reducing Costs by 4%

In "Flattening Energy-Consumption Curves by Monthly Constrained Direct Load Control Contracts," Fattahi, Ghodsi, Dasu, and Ahmadi show a California utility firm has successfully implemented a pioneering model to balance electricity demand and supply while minimizing costs. By utilizing direct load control contracts (DLCCs), the firm can reduce energy consumption during peak hours. Researchers developed an integer stochastic dynamic optimization problem that considers monthly and annual constraints, allowing for effective execution of DLCCs. Incorporating a "reduce-to-threshold" policy to flatten energy-consumption curves during high demand, the model was verified using real data from the California Independent System Operator. When implemented, the utility firm achieved an impressive cost reduction of approximately 4%. Sensitivity analysis was conducted to enhance customer experience and improve DLCC contract features. The success of this innovative model highlights the

potential of DLCCs and advanced optimization techniques in the energy sector, offering a blueprint for other utility companies seeking to optimize grid stability and reduce costs.

### Advances in Maximum-Entropy Sampling

In "Technical Note—Masking Anstreicher's  $\text{linx}$  Bound for Improved Entropy Bounds," Chen, Fampa, and Lee show a fundamental NP-hard combinatorial-optimization in the area of statistical designs is the maximum-entropy sampling problem (MESP), which seeks to maximize Shannon's "differential entropy" over all subsets of a prespecified cardinality from a set of  $n$  Gaussian random variables. This problem has applications in many areas, such as the redesign of environmental-monitoring networks. Most algorithms for exact solution of MESP are branch-and-bound based, and one of the best upper bounds is based on Anstreicher's recent concave " $\text{linx}$  relaxation" of differential entropy. A key paradigm for improving bounds is by "masking" the covariance of the random variables with a correlation matrix. The main result establishes that in the best case, the  $\text{linx}$  bound can be improved by an amount that is at least linear in  $n$  by masking. These and other recent results on the hot topic of MESP are leading to practical algorithms for exact solution of meaningful design problems in applied areas such as environmental statistics.

### High-Order Taylor Expansion of Markov Process Stationary Distributions

Much like higher-order Taylor expansions allow one to approximate functions to a higher degree of accuracy, in "High-Order Steady-State Diffusion Approximations," Braverman, Dai, and Fang demonstrate that, by accounting for higher-order terms in the Taylor expansion of a Markov process generator, one can derive novel diffusion approximations that achieve a higher degree of accuracy compared with the classical ones used in the literature over the last 50 years.

### Efficient Learning for Clustering and Optimizing Context-Dependent Designs

Contextual simulation optimization problems have attracted great attention in the healthcare, commercial, and financial fields because of the need for personalized decision making. Besides randomness in simulation outputs, larger solution space makes learning and optimization more challenging. In "Efficient Learning for Clustering and Optimizing Context-Dependent Designs," Li, Lam, and Peng use a Gaussian mixture model (GMM) as a basic technique to deal with this difficulty. To address the computational challenge in updating GMM-based Bayesian posterior, they present a computationally efficient approximation method that can reduce the computational complexity from an exponential rate to a linear rate with respect to the problem scale. For sample allocation decision making, they propose a dynamic sampling policy to efficiently utilize both global clustering information and local performance information. The proposed sampling policy is proved to be consistent, be implementable, and achieve the asymptotically optimal sampling ratio. Numerical experiments show that the proposed sampling policy significantly improves the efficiency in contextual simulation optimization.

### Designing Sustainable Distillation Sequences

Separation of mixtures of chemicals, ubiquitous in chemical and petrochemical industries, by distillation is energy intensive. Nearly 3% of the overall energy is used for distillation in the United States. Improving the distillation process is crucial for making chemical industries more sustainable. However, designing distillation sequences is challenging because the choice set is vast and the equations governing the physical process are highly

nonconvex. Traditional design practices rely on heuristics, and often result in sub-optimal solutions. In “Advances in MINLP to Identify Energy-Efficient Distillation Configurations,” Tumbalam Gooty, Agrawal, and Tawarmalani present the first approach which reliably identifies the distillation sequence that requires the least energy for a given separation. By embedding convex hulls of substructures and adapting the reformulation-linearization technique to fractions of polynomials, they demonstrated that their approach outperforms the state-of-the-art. Their work will help the chemical industry reduce greenhouse gas emissions associated with distillation.

### High-Dimensional Simulation Metamodeling

Stochastic kriging has been widely employed for simulation metamodeling to predict the response surface of complex simulation models. However, its use is limited to cases where the design space is low-dimensional because the sample complexity (i.e., the number of design points required to produce an accurate prediction) grows exponentially in the dimensionality of the design space. The large sample size results in both a prohibitive sample cost for running the simulation model and a severe computational challenge due to the need to invert large covariance matrices. In “Sample and Computationally Efficient Stochastic Kriging in High Dimensions,” Ding and Zhang address this long-standing challenge and develop a novel methodology — based on tensor Markov kernels and sparse grid experimental designs — that dramatically alleviates the curse of dimensionality. The proposed methodology has theoretical guarantees on both sample complexity and computational complexity and shows outstanding performance in numerical problems of as high as 16,675 dimensions.

### Driver and Mitigation of Free Rider Effect in Investment in the Common Good

The free rider problem is widely observed in investments in the common good. For example, if Best Buy offers corporate social responsibility (CSR) programs, such as recycling, tech training, and supplier audit, at its own expense, other firms in the industry can benefit from these programs at no cost (i.e., they become the free riders). Such a free rider effect often manifests itself as a mixed strategy equilibrium, in which potential investors unnecessarily delay their investments. In “Investment in The Common Good: Free Rider Effect and the Stability of Mixed Strategy Equilibria,” Kim and Kwon find that recurring investment opportunities can be a major driver of delays in investment in the common good. They also show that sufficient heterogeneity in investment cost among potential investors can alleviate the free rider effect. Therefore, the policy makers may try to facilitate investments in public goods (e.g., CSR activities) through some form of selective benefits to potential investors.

### Off-Policy Evaluation Through the Lens of Distributionally Robust Optimization

Off-policy evaluation is an important topic in reinforcement learning, which estimates the expected cumulative reward of a target policy using logged trajectory data generated from a different behavior policy, without execution of the target policy. It is imperative to quantify the uncertainty of the off-policy estimate before deployment of the target policy. In “Reliable Off-Policy Evaluation for Reinforcement Learning,” Wang, Gao, and Zha leverage methodologies from (Wasserstein) distributionally robust optimization to provide robust and optimistic cumulative reward estimates. With proper selection of the size of the distributional uncertainty set, these estimates serve as confidence bounds with nonasymptotic and asymptotic guarantees under stochastic or

adversarial environments. The authors also generalize those results to batch reinforcement learning.

### On Decision Rules for Multistage Stochastic Programs with Mixed-Integer Decisions

Multistage stochastic programming is a field of stochastic optimization for addressing sequential decision-making problems defined over a stochastic process with a given probability distribution. The solution to such a problem is a decision rule (policy) that maps the history of observations to the decisions. Design of the decision rules in the presence of mixed-integer decisions is quite challenging. In “Lagrangian Dual Decision Rules for Multistage Stochastic Mixed-Integer Programming,” Daryalal, Bodur, and Luedtke introduce Lagrangian dual decision rules, where linear decision rules are applied to dual multipliers associated with Lagrangian duals of a multistage stochastic mixed-integer programming (MSMIP) model. The restricted decisions are then used in the development of new primal- and dual-bounding methods. This yields a new general-purpose approximation approach for MSMIP, free of strong assumptions made in the literature, such as stagewise independence or existence of a tractable-sized scenario-tree representation.

### Joint Inventory and Pricing for a One-Warehouse Multistore Problem: Spiraling Phenomena, Near Optimal Policies, and the Value of Dynamic Pricing

In “Joint Inventory and Pricing for a One-Warehouse Multistore Problem: Spiraling Phenomena, Near Optimal Policies, and the Value of Dynamic Pricing,” Lei, Liu, Jasin, and Vakhutinsky consider a joint inventory and pricing problem with one warehouse and multiple stores with lost sales. The retailer makes a one-time decision on the amount of inventory to be placed at the warehouse at the beginning of the selling season, followed by periodic joint replenishment and pricing decisions for each store throughout the season. The authors first analyze the performance of two popular and simple heuristic policies that directly implement the solution of a deterministic approximation of the original stochastic problem. They show that simple reoptimization of the deterministic approximation may worsen the performance by causing a “spiraling up” movement in expected lost sales quantity. The authors further propose two improved heuristic policies with provably near-optimal performance. In particular, the first policy achieves the best possible performance among all policies that rely on static pricing, and the second policy outperforms the first one because of its use of carefully designed dynamic pricing scheme.

### What Would Dantzig Do with a Quantum Computer?

It is unlikely we will ever find out the answer to this question. But we can try to understand if the simplex method can be implemented on a quantum computer, and this might have piqued Dantzig’s interest. In “Fast Quantum Subroutines for The Simplex Method,” Nannicini gives a quantum implementation of an iteration of the simplex method, in which the basis inverse is never explicitly computed: The quantum computer takes as input the current basis, and certifies optimality or outputs the next basis. Since computing the basis inverse is expensive, this can lead to an asymptotically faster algorithm in terms of the problem size: in the best case, the quantum algorithm can identify pivots in essentially linear time! This, however, comes at the cost of worse dependence on some numerical parameters: all these trade-offs are discussed in the full article.

## Data-Driven Ranking and Selection Under Input Uncertainty

In many applications, input data are collected frequently to update the simulation model of the system, whereas simulation is run to compare different designs/strategies to identify the best one with a high confidence. In “Data-Driven Ranking and Selection Under Input Uncertainty,” Wu, Wang, and Zhou consider such a simulation-based ranking and selection (R&S) problem, in which the input distribution is estimated and updated with input data arriving in batches over time. Unlike most existing works of R&S that conduct simulation under a fixed distribution, in this data-driven setting, simulation outputs are generated under different input distributions over time. A moving average estimator is introduced to aggregate simulation outputs generated under heterogeneous distributions. Then, two sequential elimination procedures are devised by establishing exact and asymptotic confidence bands for the estimator. The efficiency of the procedures can be further boosted by incorporating the “indifference zone” idea and optimizing the “drop rate” parameter of the moving average estimator.

## Black-Box Acceleration of Monotone Convex Program Solvers

In “Black-Box Acceleration of Monotone Convex Program Solvers,” London, Vardi, Eghbali, and Wierman present a framework for accelerating (speeding up) existing convex program solvers. Across engineering disciplines, a fundamental bottleneck is the availability of fast, efficient, accurate solvers. The authors present an acceleration method that speeds up linear programming solvers such as Gurobi and convex program solvers such as the Splitting Conic Solver by two orders of magnitude.

## Decision-Making Under Simultaneous Competition

Hardly any decision is made in isolation and most decision makers are dealing with fierce competition when trying to find the optimal decision for their problem. The expected outcome of such a competitive problem setting or the individually optimal course of action for each competitor is not evident. In a finite game, a finite set of decision makers simultaneously select their action from a finite set of strategies. In “Equilibrium Identification and Selection in Finite Games,” Crönert and Minner propose a solution approach enumerating all equilibria and selecting the most likely equilibrium in finite games. The approach is targeted

toward large finite games that cannot be efficiently represented in normal form. They apply their algorithm to two- and three-player knapsack and facility location and design games. Their numerical experiments show that prior approaches identifying a single equilibrium can result in unlikely outcomes.

## Optimal Investment, Heterogeneous Consumption, and Best Time for Retirement

In “Optimal Investment, Heterogeneous Consumption, and Best Time for Retirement,” Jang, Xu, and Zheng study an optimal investment and consumption problem with heterogeneous consumption of basic and luxury goods, together with the choice of time for retirement. The optimal heterogeneous consumption strategies for a class of nonhomothetic utility maximizer are shown to consume only basic goods when the wealth is small, to consume basic goods and make savings when the wealth is intermediate, and to consume almost all in luxury goods when the wealth is large. The optimal retirement policy is shown to be both universal, in the sense that all individuals should retire at the same level of marginal utility that is determined only by income, labor cost, discount factor as well as market parameters, and not universal, in the sense that all individuals can achieve the same marginal utility with different utility and wealth. They also show that individuals prefer to retire as time goes by if the marginal labor cost increases faster than that of income.

## Debiasing In-Sample Policy Performance for Small-Data, Large-Scale Optimization

In many modern large-scale decision-making problems, data can be scarce. As a result, traditional methods such as cross-validation perform poorly in evaluating the performance of decision-making policies. In “Debiasing In-Sample Policy Performance for Small-Data, Large-Scale Optimization,” Gupta, Huang, and Rusmevichientong propose a novel estimator of the out-of-sample performance for a policy in data-driven optimization. Unlike cross-validation, their approach avoids sacrificing training data for evaluation. As a result, they theoretically show the estimator is asymptotically unbiased as the problem size grows. Furthermore, the authors show that the estimator is asymptotically optimal when applied to more specialized “weakly coupled” optimization problems. Finally, using a case study on dispatching emergency medical response services, they demonstrate their proposed method provides more accurate estimates of out-of-sample performance and selects better policies.