



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(6):iii-vii. <https://doi.org/10.1287/opre.2023.inthisissue.v71.n6>

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.n6>

Copyright: © 2023 INFORMS

A Novel Practical Stochastic Pricing Model for Multi-Interval Real-Time Markets

Practical implementations of economic dispatch with associated pricing systems are crucial for operating electricity markets. Because of the high volatility caused by the increasing integration of renewable energy, consideration of the underlying stochastic problem is becoming more important than ever. It is a challenge to incorporate the uncertain nature of real-time operations into an already complex multi-interval dynamic problem with intertemporal constraints. Because solving a standard multi-stage stochastic programming problem is too burdensome in terms of calculation time for real-time markets, it has been standard practice in electricity markets to use a deterministic approximation with varying degrees of look-ahead. In “Pricing Under Uncertainty in Multi-Interval Real-Time Markets,” Cho and Papavasiliou introduce a practical alternative method for pricing under uncertainty in multi-interval real-time markets. Using slightly different stochastic formulations, the authors propose an approach that preserves the attractive features from both the deterministic formulation (simpler calculation) and the standard stochastic formulation (better performance).

Rectangular Decomposition of Mixed Integer Programs via Decision Diagrams with Application to Unit Commitment

Despite the successful applications of decision diagrams (DDs) to solve various classes of integer programs in the literature, the question of which mixed-integer structures admit a DD representation remains open. In “On the Structure of Decision Diagram-Representable Mixed-Integer Programs with Application to Unit Commitment,” Salemi and Davarnia address this question by developing both necessary and sufficient conditions for a mixed-integer program to be DD-representable through identification of certain rectangular formations in the underlying sets. This so-called rectangularization framework is applicable to all bounded mixed-integer linear programs, providing a notable extension of the DD domain to continuous problems. As an application, the authors use the developed methods to solve stochastic unit commitment problems in energy systems. Computational experiments conducted on benchmark instances show that the DD approach uniformly and significantly outperforms the existing solution methods and modern solvers. The proposed methodology opens new pathways to solving challenging mixed-integer programs in energy systems more efficiently.

Enhancing Renewable Integration: Energy Storage and Uncertainty Management

In “Unit Commitment Problem with Energy Storage Under Correlated Renewables Uncertainty,” Cordera, Moreno, and Ordoñez introduce a novel approach to address the challenges of renewable integration. The authors acknowledge the growing variability and correlation in power availability due to renewable generation

and proposes a day-ahead unit commitment (UC) problem formulation that incorporates energy storage and considers multi-stage correlated uncertainty. Using a variant of the stochastic dual dynamic programming (SDDP) method, which can handle temporal correlations effectively, the researchers solve the complex UC problem. Results obtained from the IEEE 118-bus system demonstrate the significant advantages of considering multistage uncertainty and correlations. Applying their approach to the Chilean power system, the researchers present superior UC solutions that adapt generation to changing uncertainty at a lower cost. Additionally, they propose a more efficient deterministic UC solution that outperforms current industry practices. These advancements promise to enhance the integration of renewable energy sources, enabling a more sustainable and environmentally friendly future.

Optimizing State of Charge for Storage Resources

Storage of generated energy has existed as a contributing resource within the power sector for more than a century in the case of pumped hydro storage. Even though most regional transmission organizations/independent system operators do not currently optimize state of charge (SOC) for storage, the SOC is an essential aspect of the operating characteristics of storage. In “Optimization Formulations for Storage Devices with Disjoint Operating Modes,” Baldick, Chen, and Huang investigate the impact on computational performance of optimizing storage with explicit representation of SOC. To investigate the model, the authors consider two contexts, stand-alone and large-scale. Analysis of the stand-alone context helps to explain why the combination of features in the storage model results in a difficult problem. Numerical results for large-scale test cases verify that new tighter valid inequalities can improve the computation compared with the standard model in the literature.

Intro to the ARPA-E Grid Optimization Competition

In “Recent Developments in Security-Constrained AC Optimal Power Flow: Overview of Challenge 1 in the ARPA-E Grid Optimization Competition,” Aravena, Molzahn, Zhang, Petra, Curtis, Tu, Wächter, Wei, Wong, Gholami, K. Sun, X.A. Sun, Elbert, Holzer, and Veeramany review the state of the art in practical algorithms for scheduling power-systems operations in the short term and the results of the recent competition organized by the U.S. Advanced Research Projects Agency–Energy. They explain the mixed-integer nonlinear formulation used in the competition for nonspecialists in electrical engineering, the context and organization of the competition, and the performance of competitors. The authors find that the collective approaches and results of competitors provide support for efforts to move nonlinear optimization techniques into industrial applications, as they have proven to be a robust and efficient alternative to current linear approximation techniques.

Solving Realistic Security-Constrained Optimal Power Flow Problems

In “A Surrogate-Based Asynchronous Decomposition Technique for Realistic Security-Constrained Optimal Power Flow Problems,” Petra and Aravena propose a new algorithm for solving a classical problem in power grid operations: the security-constrained optimal power flow, considering its nonlinearities and realistic transitions between nominal and emergency post-contingency operations. Solving security-constrained optimal power flow problems accurately is a critical function, upon which depends the reliability, security, and efficiency of power systems as well as the correct functioning of other critical infrastructure dependent on electricity. The proposed algorithm was extensively tested against many state-of-the-art approaches using realistic and real instances in the ARPA-E Grid Optimization Competition Challenge 1, where it found the best-known solution for 58% of the instances, attained an average gap of less than 0.2%, and obtained the best overall scores, thereby winning all divisions of Challenge 1 with a very strong first place.

Team GO-SNIP’s Algorithm for the ARPA-E Grid Optimization (GO) Competition, Challenge 1

In “A Decomposition Algorithm with Fast Identification of Critical Contingencies for Large-Scale Security-Constrained AC-OPF” Curtis, Molzahn, Tu, Wächter, Wei, and Wong present the decomposition algorithm used by Team GO-SNIP for the ARPA-E Grid Optimization (GO) Competition Challenge 1, held from November 2018 through October 2019. The algorithm involves unique contingency ranking and evaluation strategies for determining the important contingencies to include in a master problem that approximates the original large-scale security-constrained problem. It also involves efficient strategies for handling the complementarity constraints that appear in the model and for handling the arising degeneracies. Software implementation details are described, and the results of an extensive set of numerical experiments are provided to illustrate the effectiveness of each of the used techniques. Team GO-SNIP received a second-place finish in Challenge 1 of the Go Competition

Preventing Failures and Optimizing Electric Power Systems Through Efficient Computation Techniques

When optimizing electric power system operational decisions, it is of great importance to prevent potential failures in both the system operation and the optimization algorithm. In “An ADMM-Based Distributed Optimization Method for Solving Security-Constrained Alternating Current Optimal Power Flow,” Gholami, Sun, Zhang, and Sun propose a novel two-level algorithm that (1) effectively prevents power system operational failures through consideration of impactful contingencies and (2) guarantees convergence when parallelized on a computing cluster with multiple nodes. Extensive numerical experiments suggest that the proposed algorithm is able to provide high-quality feasible solutions under the time limit of 10–45 minutes for various synthetic and industrial networks with up to 30,000 buses and 22,000 contingencies, comparable with the size of the U.S. power grid.

A Generic Way to Verify Asymptotic Optimality of Semiopen-Loop Policies For a Wide Class of MDPs with Large Lead Times

In many real-life inventory models, order lead times can result in uncertain effects of inventory decisions. However, as the lead time grows large, one would naturally postulate that the effect of the delayed order depends weakly on the current inventory level

and, thus, intuit that decoupling the delayed order with the current inventory level may provide good heuristics. Motivated by these examples, in “Asymptotic Optimality of Semiopen-Loop Policies in Markov Decision Processes with Large Lead Times,” Bai, Chen, Li, and Stolyar consider a generic Markov decision process (MDP) with one delayed control and one immediate control. For MDPs defined on general spaces with uniformly bounded cost functions and a fast mixing property, they construct a periodic semiopen-loop policy for each lead time value and show that these policies are asymptotically optimal as the lead time goes to infinity. For MDPs defined on Euclidean spaces with linear dynamics and convex structures, they impose another set of conditions under which constant delayed-control policies are asymptotically optimal.

Portfolio Risk Taking when Periodic Performance Matters

Incentives of portfolio managers in real life are typically tied to their periodic performance—a criterion that has received very little attention to date. How does the repeated nature of the rewards influence the portfolio managers’ trading strategies? In “Portfolio Selection, Periodic Evaluations, and Risk Taking,” Tse and Zheng address this question via a dynamic model of investment. The return benchmark of the periodic evaluations plays a crucial role. A reasonable benchmark in conjunction with periodic payouts can mitigate excessive risk taking during market downturns. However, an overly aggressive benchmark could result in limited portfolio growth as well as ruin risk over the long run. The theoretical findings yield policy implications over incentives management and remuneration structure design to better align the short-term periodic interest of the portfolio managers and the long-term performance goal of the stakeholders.

Efficient Learning Algorithms for Dynamic Inventory Allocation in Multiwarehouse Multistore Systems with Censored Demand

In “Inventory Control and Learning for One-Warehouse Multistore System with Censored Demand,” motivated by collaboration with a prominent fast-fashion retailer in Europe, Bekci, Gümüüş, and Miao focus their attention on the one-warehouse multistore (OWMS) inventory control problem, specifically addressing scenarios in which the demand distribution is unknown a priori. The OWMS problem revolves around a central warehouse that receives initial replenishments and subsequently distributes inventory to multiple stores within a finite time horizon. The objective lies in minimizing the total expected cost. To overcome the hurdles posed by the unknown demand distribution, the researchers propose a primal-dual algorithm that continuously learns from demand observations and dynamically adjusts inventory control decisions in real time. Thorough theoretical analysis and empirical evaluations highlight the promising performance of this approach, offering valuable insights for efficient inventory allocation within the ever-evolving retail industry.

A Data-Driven Approach to Improve Care Unit Placements in Hospitals

The choice of care unit upon hospital admission is a challenging task because of the wide variety of patient characteristics, uncertain needs of patients, and limited number of beds in intensive and intermediate care units. These decisions require carefully weighing the benefits of improved health outcomes against the opportunity cost of reserving higher level care beds for potentially more complex patients arriving in the future. In “Data-Driven Hospital Admission Control: A Learning Approach,” Zhalechian, Keyvanshokoh, Shi, and Van Oyen introduce a data-driven

algorithm to address this challenging task. By focusing on reducing the readmission risk of patients, the algorithm is designed to (i) adaptively learn the readmission risk of patients through batch learning with delayed feedback and (ii) determine the best care unit placement for a patient based on the observed information and occupancy levels to minimize total readmission risk. The algorithm is supported by a performance guarantee, and its effectiveness is showcased using real-world hospital system data.

Mixed-Integer Programming Formulations for Systemic Risk Measures

Measurement and allocation of systemic risk have become important tasks in interconnected financial networks. In “Computation of Systemic Risk Measures: A Mixed-Integer Programming Approach,” Ararat and Meimanjan consider a clearing model that considers external assets and business costs of a bank simultaneously by extending the Rogers–Veraart network model with default costs. They prove that clearing payment vectors can be calculated by solving mixed-integer programming problems in which the existence of business costs and default costs are represented by binary variables. The systemic risk measure corresponding to this model yields nonconvex sets of capital allocation vectors. The authors compute these nonconvex sets by solving a multiobjective optimization problem whose scalarizations are two-stage mixed-integer programming problems with an expectation constraint. The computational procedure is versatile enough to accommodate other applications of systemic risk where the risk factors are aggregated by solving a general mixed-integer programming problem.

Debiasing Profit Forecasts for the Newsvendor

An unbiased forecast of profit is important in most business environments. Typically, forecasts are generated from data. However, in “Technical Note—Data-Driven Profit Estimation Error in the newsvendor model,” Siegel and Wagner identify a strictly positive bias in a natural estimation of expected profit in a data-driven newsvendor model, where managers will expect more profit than will actually be realized, on average. This bias can reach significant proportions (in some cases 50%+) of the true expected profit and could therefore have undesired and damaging effects in the real world. The authors then design a data-driven adjustment that results in an unbiased estimator of expected profit, so that managers may have an accurate forecast of future profit that is free of systematic bias.

Customized Dynamic Pricing When Customers Develop a Habit or Satiation

In “Customized Dynamic Pricing When Customers Develop a Habit or Satiation,” Chen, He, and Bansal discuss optimal prices that a firm should offer to customers who develop a habit or satiation from consumption. At a high discount (or equivalently a lower price), a habit-prone customer will purchase a large amount in a current period. The firm will earn a low revenue per unit on the amount bought by the customer in this period. However, this large purchase will also create a stronger habit in the customer that the firm can exploit in the future periods. At a low discount, the firm will earn a higher per unit revenue in this period. However, the lower purchase quantity will also result in a weaker habit in the customer for the future periods. This tradeoff exists in every period. These effects are reversed for a satiation-prone customer. Using an analytical model, the authors determine the optimal profit-maximizing dynamic prices that a firm should offer to such habit- or satiation-prone customers. The results are robust for various specifications for habit formation, and yet in each specification they are obtainable tractably.

A New Connection Between Elicitation and Aggregation of Forecasts

There are many ways to elicit honest probabilistic forecasts from experts. Once those forecasts are elicited, there are many ways to aggregate them into a single forecast. Should the choice of elicitation method inform the choice of aggregation method? In “From Proper Scoring Rules to Max-Min Optimal Forecast Aggregation,” Neyman and Roughgarden establish a connection between these two problems. To every elicitation method they associate the aggregation method that improves as much as possible upon the forecast of a randomly chosen expert, in the worst case. This association maps the two most widely used elicitation methods (Brier and logarithmic scoring) to the two most well-known aggregation methods (linear and logarithmic pooling). The authors show a number of interesting properties of this connection, including a natural axiomatization of aggregation methods obtained through the connection, as well as an algorithm for efficient no-regret learning of expert weights.

Revenue Management with Unique Resources

There are a variety of revenue management systems that require making pricing or availability decisions for unique resources. For example, lodging marketplaces, boutique hotels, and bed-and-breakfasts offer unique rooms, apartments, or houses. Matching platforms for freelancers recommend differentiated workers with unique characteristics. When managing unique resources, one has to keep track of the availability of each resource at each time point in the future. Moreover, if the customers substitute between different resources, then the pricing and availability decisions for all resources become interdependent. Thus, it can be challenging to find good policies to make pricing or availability decisions. In “Revenue Management with Heterogeneous Resources: Unit Resource Capacities, Advance Bookings, and Itineraries over Time Intervals,” Rusmevichientong, Sumida, Topaloglu, and Bai consider revenue management problems when unique resources are requested for use over intervals of time under advance reservations. Using the interval structure of resource requests, they give policies with performance guarantees.

Scheduling Advertising on Cable Television

Advertisement scheduling is a daily essential operational process in the television business. Efficient distribution of viewers among advertisers allows the television network to satisfy contracts and increase ad sale revenues. Ad scheduling is a challenging multi-period, mixed-integer programming problem in which the network must create schedules to meet advertisers’ campaign goals and maximize ad revenues. Each campaign must meet a specific target group of viewers and a unique set of constraints. Moreover, the number of viewers is uncertain. In “Scheduling Advertising on Cable Television,” Souyris, Seshadri, and Subramanian solve this problem by developing and implementing a practical approach that combines mathematical programming and machine learning to create daily schedules. According to standard business metrics and the small integer programming gap, these schedules are of high quality. Using their methods, leading networks in the United States and India experience a 3% to 5% revenue increase, which translates to about \$60 million annually for one prominent user.

Optimizing Risk-Balancing Return under Discrete Choice Models

Similar to individual investors, firms balance revenue or profit risk with the expected return on investment. A stable cash flow is required to satisfy private and public shareholders or to maintain solvency. Product prices impact both the expected demand for the

products and the volatility of the demand, and consequently a firm's revenue or profit risk. In "Technical Note—Optimizing Risk-Balancing Return Under Discrete Choice Models," Li and Webster present models and results that demonstrate how risk sensitivity can inform pricing strategy. The authors show that a firm's distaste for risk not only causes the firm to discount its products in exchange for lower profit volatility, but also reduces the firm's incentive to price differentiate among the products. The findings also provide insight into how firms with varying degrees of risk aversion may price products of different quality and how this pricing strategy affects product market share.

Is a Single-Tier Subscription Optimal for Streaming Platforms?

In "Optimal Subscription Planning for Digital Goods," Alaei, Makhdoumi, and Malekian consider the problem of a streaming platform that offers different set of contents with different subscription fees with the goal of maximizing revenue. They characterize the necessary and sufficient condition for the consumer preferences under which a single-tier subscription is optimal. In particular, the authors show the preferences of the consumers for different contents must be aligned for a single-tier subscription to be optimal. This condition is often violated in practice, establishing the suboptimality of a single-tier offering. They further show that a subscription fee proportional to the number of contents is near-optimal.

Looking for Guaranteed Starting Times in Project Scheduling Under Uncertainty

In project scheduling, the durations of activities are often uncertain. Delays may cause a massive disorganization if a large number of activities must be rescheduled. In "The Anchor-Robust Project Scheduling Problem," Bendotti, Chrétienne, Fouilhoux, and Pass-Lanneau propose a novel criterion for solution stability in project scheduling under processing times uncertainty. They define anchored jobs as jobs whose starting times can be guaranteed in a baseline schedule. Finding a project schedule with bounded makespan and a max-weight set of anchors is shown to be an NP-hard robust two-stage problem. Taking advantage of the combinatorial structure of project scheduling and budgeted uncertainty, the authors obtain a compact MIP formulation for the problem. Numerical results show that the obtained MIP outperforms standard techniques from the literature. They also showcase the practical interest of anchored jobs in project scheduling.

Finite-Sample Guarantees for Wasserstein Distributionally Robust Optimization: Breaking the Curse of Dimensionality

Wasserstein distributionally robust optimization is a recent emerging modeling paradigm for decision making under data uncertainty. Because of its computational tractability and interpretability, it has achieved great empirical successes across several application domains in operations research, computer science, engineering, and business analytics. Despite its recent empirical success, existing performance guarantees for generic problems are not yet satisfactory. In "Finite-Sample Guarantees for Wasserstein Distributionally Robust Optimization: Breaking the Curse of Dimensionality," Gao develops the first finite-sample guarantee without suffering from the curse of dimensionality, which describes how the out-of-sample performance of a robust solution depends on the sample size, dimension of the uncertainty, and the complexity of the loss function class in a nearly optimal way.

On Finite Adaptability in Two-stage Distributionally Robust Optimization

In "On Finite Adaptability in Two-Stage Distributionally Robust Optimization," Han, Bandi, and Nohadani study finite adaptability with the goal to construct interpretable and easily implementable policies in the context of two-stage distributionally robust optimization problems. To achieve this, the set of uncertainty realizations needs to be partitioned. The authors show that an optimal partitioning can be accomplished via "translated orthants." They then propose a nondecreasing orthant partitioning and binary approximation to obtain the corresponding "orthant-based policies" from a mixed-integer optimization problem of a moderate size. For these policies, they provide provable performance bounds, generalizing the existing bounds in the literature. For more general settings, they also propose optimization formulations to obtain posterior lower bounds that can serve to evaluate performance. Two numerical experiments support these findings. A joint inventory-routing problem highlights the practical applicability for large-sized instances with mixed-integer recourse. A case study from a pharmacy retailer demonstrates that the orthant-based policies are less sensitive to cost parameters than optimal solutions, enabling these policies to outperform comparable methods when the realized cost ratio deviates from its nominal value.

Information Retrieval Under Network Uncertainty: Robust Internet Ranking

Ranking algorithms play a crucial role in information technologies and numerical analysis due to their efficiency in high dimensions and wide range of possible applications, including internet ranking, scientometrics, and systemic risk in finance (SinkRank and DebtRank). The traditional approach to internet ranking goes back to the seminal work of Sergey Brin and Larry Page, who developed the initial method PageRank (PR) in order to rank websites for search engine results based on linear algebra rules. But how robust is this method in times of rapid internet growth? Recent works have studied robust reformulations of the PageRank model for the case when links in the network structure may vary; that is, some links may appear or disappear, influencing the transportation matrix defined by the network structure. In "Information Retrieval Under Network Uncertainty: Robust Internet Ranking," Timonina-Farkas and Seifert make a further step forward, allowing the network to vary not only in links but also in the number of nodes. The authors focus on growing network structures and develop methods for ranking of networks uncertain both in size and in structure.

Enhanced Balancing of Bias-Variance Tradeoff in Stochastic Estimation: A Minimax Perspective

In "Enhanced Balancing of Bias-Variance Tradeoff in Stochastic Estimation: A Minimax Perspective," Lam, Xinyu Zhang, and Xuhui Zhang study a framework to construct new classes of stochastic estimators that can consistently beat existing benchmarks regardless of key model parameter values. Oftentimes biased estimators, such as finite-difference estimators in black box stochastic gradient estimation, require selection of tuning parameters to balance bias and variance and ultimately minimize overall errors. Unfortunately, this relies on model knowledge that is unknown a priori and thus leads to ad hoc choices in practice. The authors introduce a new notion called asymptotic minimax risk ratio, which is designed to compare new estimators against existing benchmarks, whose values less than one imply that the new estimators could asymptotically outperform the benchmarks regardless of the model parameter value. Based on this, the authors study an outperforming weighting scheme by explicitly analyzing

the asymptotic minimax risk ratio via a tractable reformulation of a nonconvex optimization problem.

Comparing Two Classical Approximations to Weakly Coupled Dynamic Programs

Many sequential decision problems have a weakly coupled structure in that a set of linking constraints couples an otherwise independent collection of subproblems. This structure arises in a wide variety of applications, such as network revenue management, online advertising, assortment planning, interactive marketing, optimization of power systems, and multilocation inventory management to name only a few. Such problems can be modeled as dynamic programs but are quite difficult to solve. Two widely studied approximation methods are approximate linear programs, which involve finding a best approximation of total value that is additive across the subsystems, and Lagrangian relaxations, which involve relaxing the linking constraints. It is well known that both of these approaches provide upper bounds to the optimal value, and the approximate linear programming approach is a better bound but also, more difficult to compute. In “Technical Note—On the Strength of Relaxations of Weakly Coupled Stochastic Dynamic Programs,” Brown and Zhang provide a detailed theoretical analysis of these two approximations and show that, under fairly broad conditions, these two approximations lead to upper bounds that are very close and often identical. Their

theory suggests that, between these two approximations, Lagrangian relaxations should usually be the preferred choice for researchers studying applications involving weakly coupled dynamic programs.

Existence of Average-Cost Optimal Policies in Lost-Sales Inventory Models with Incomplete Information

In many real-life situations, the inventory record may not match the actual stock perfectly. This can happen due to distortion of inventory data, such as transaction errors, misplaced inventories, and spoilage. In these cases, because the decision maker only has incomplete information about the inventory levels, many well-known inventory policies are not even admissible, and our understanding of the optimal policies, even their existence, is very limited. In “Technical Note—Average Cost Optimality in Partially Observable Lost-Sales Inventory Systems,” Bai, Chen, and Stolyar consider the classical lost-sales inventory model, in which the inventory level is only observed when it becomes zero. They formulate the cost-minimization problem as a partially observable Markov decision process. By exploiting the vanishing discount factor approach, they provide a way to verify the existence of optimal policies under the average cost criterion. The key step in their analysis is the construction of a valid policy, which, in a certain sense, copies the actions of another policy for the process starting from another initial state.