



Operations Research

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

In This Issue

To cite this article:

(2021) In This Issue. Operations Research 69(2):ii-iv. <https://doi.org/10.1287/opre.2021.2137>

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 © 2021, 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.2021.2137>

Copyright: © 2021 INFORMS

Interpretable, Computationally Tractable Approximate Parameter Estimation for Corporate Defaults

Measuring corporate default risk is a classic problem in finance with a rich history spread over the last 50 years. Typically, default probabilities of firms are assigned a parametric structure, and the parameters are estimated from data using computationally demanding algorithms. In “Credit Risk: Simple Closed-Form Maximum Likelihood Estimation,” Deo and Juneja observe that because these probabilities over short time periods are small, a closed-form, interpretable approximation to the maximum likelihood estimator (MLE) exists. In particular, the approximation identifies the key data set that determines the underlying parameters. The authors analyze the approximation and the MLE in a realistic asymptotic regime where the default probabilities are small, but the number of firms and time horizon of available data are large. Their key conclusion, corroborated by experiments on a representative corporate data set as well as through simulations, is that both the proposed estimator and the MLE have similar accuracy.

Vehicle Routing with Multiple Depots on Directed Graphs

In “Asymmetric Multidepot Vehicle Routing Problems: Valid Inequalities and a Branch-and-Cut Algorithm,” uit het Broek, Schrottenboer, Jargalsaikhan, Roodbergen, and Coelho present a generic branch-and-cut framework to solve routing problems with multiple depots on directed graphs. They present new valid inequalities that eliminate subtours, enforce tours to be linked to the same depot, and enforce bounds on the number of customers in a vehicle tour. This is embedded in a branch-and-cut scheme that also contains generalized and adapted versions of valid inequalities that are well known for related routing problems. The authors show that the new inequalities tighten root node relaxations considerably. In combination with a simple but effective upper-bound procedure, only requiring a MIP solver and a smart reduction of the problem size, the authors show that the overall framework solves instances of considerably larger size to optimality than have been reported in the literature.

Selfish Routing is Efficient Under High Traffic Demand

The price of anarchy (PoA) is a standard measure for the inefficiency of selfish routing in the static Wardrop traffic model. Empirical studies and a recent analysis reveal a surprising property that the PoA tends to one when the total demand T gets large. These results are extended by a new

framework for the limit analysis of the PoA in arbitrary nonatomic congestion games that apply to arbitrary growth patterns of T and all regularly varying cost functions. For routing games with Bureau of Public Road (BPR) cost functions, the convergence follows a power law determined by the degree of the BPR functions, and a related conjecture need not hold. In “Selfishness Need Not Be Bad,” Wu, Möhring, Chen and, Xu, show these findings are confirmed by an empirical analysis of traffic in Beijing.

Accounting for Travel Time Correlations When Dispatching a Fleet of Vehicles

Travel times on a road network are commonly affected by some degree of uncertainty. The duration of vehicle routes may then vary according to the observed travel time realizations. For managerial reasons, however, a fleet operator is generally interested in designing routes whose duration is not too dispersed with respect to their expected value. In this context, correlations among travel times play an important role, and failing to account for them might result in a large underestimation of the variance of the route duration.

In “Branch-Price-and-Cut Algorithms for the Vehicle Routing Problem with Stochastic and Correlated Travel Times,” Rostami, Desaulniers, Errico, and Lodi adopt a mean-variance model leading to a quadratic binary program. The authors propose two alternative set partitioning reformulations, which are then solved by branch-price-and-cut algorithms. The obtained solutions significantly reduce the time variability when compared with solutions of the classical vehicle routing problem.

An Equitable Gale-Shapley Algorithm

In a foundational paper, Gale and Shapley (1962) introduced the deferred acceptance algorithm that achieves a stable outcome in a two-sided matching market by letting one side of the market make proposals to the other side. What happens when both sides of the market can propose? In “Deferred Acceptance with Compensation Chains,” Dworczak answers this question by constructing an equitable version of the Gale-Shapley algorithm in which the sequence of proposers can be arbitrary. The main result of the paper shows that the extended algorithm, equipped with so-called compensation chains, is not only guaranteed to converge in polynomial time to a stable outcome, but—in contrast to the original Gale-Shapley algorithm—achieves all stable matchings (as the sequence of proposers vary). The proof of convergence uses a novel potential function. The algorithm may find applications in settings where both stability and fairness are desirable features of the matching process.

Vehicle Routing for Maximal Delivery Service

On-time delivery is of utmost importance in today's urban logistics. However, travel times are uncertain and classical deterministic routing solutions often fail to ensure timely delivery. In "Robust Data-Driven Vehicle Routing with Time Windows," by Y. Zhang, Z. Zhang, Lim, and Sim, a robust solution that exploits travel times data to determine the best routes for maximal timely delivery is proposed. A new decision criterion is introduced, the service fulfillment risk index (SRI), which accounts for both the late arrival probability and its magnitude. Together with Wasserstein distance-based ambiguity in travel times, SRI can be evaluated efficiently in closed form. In addition, an exact branch-and-cut approach and a meta-heuristic algorithm are developed to minimize SRI with a given travel cost. Simulation studies demonstrate that handling uncertainty improves service punctuality, and that incorporating ambiguity prevents overfitting. Most importantly, SRI outperforms the canonical decision criteria of lateness probability and expected lateness duration.

Dynamic Reserve Prices Optimization: Learning from Bids

How can an auctioneer optimize revenue by learning the reserve prices from the bids in the previous auctions? How should the long-term incentives and strategic behavior of the bidders be taken into account? Motivated in part by applications in online advertising, in "Incentive-Compatible Learning of Reserve Prices for Repeated Auctions," Kanoria and Nazerzadeh investigate these questions. They show that if a seller attempts to dynamically update a common reserve price using the bidding history, buyers will shade their bids, which can hurt the revenue. However, when there is more than one buyer, using personalized reserve prices, the auctioneer can achieve a near-optimal revenue. In their proposed mechanism, the personal reserve price for each buyer is determined using the historical bids of other buyers.

Defender-Attacker-Defender Games

Multilevel programming can provide the right mathematical formulations for modeling sequential decision-making problems. In such cases, it is implicit that each level anticipates the optimal reaction of the subsequent ones. Defender-attacker-defender trilevel programs are a particular case of interest that encompasses a fortification strategy, followed by an attack, and a consequent recovery defensive strategy. In "Multilevel Approaches for the Critical Node Problem," Baggio, Carvalho, Lodi, and Tramontani study a combinatorial sequential game between a defender and an attacker that takes place in a network. The authors propose an exact algorithmic framework. This work highlights the significant improvements that the defender can achieve by taking the three-stage game into account instead of considering fortification and recovery as isolated. Simultaneously, the paper contributes to advancing the methodologies for solving trilevel programs.

Demand Shaping Through Bundling and Product Configuration: A Dynamic Multiproduct Inventory-Pricing Model

In today's digital age with advanced information technology and analytic tools, such as the internet and data mining, many companies use dynamic bundling to influence demand to match up with the operational status, especially in industries with short product life cycles. In "Demand Shaping Through Bundling and Product Configuration: A Dynamic Multiproduct Inventory-Pricing Model," Song and Xue present a novel dynamic model to analyze how exactly the inventory dynamics impact the bundling strategy and, in turn, how the bundling strategy affects the firm's inventory decisions. They also characterize how the optimal bundling decision depends on item complementarity, cost structure, inventory levels, demand uncertainty, and supply responsiveness.

Network Effects and Competitive Strategies

It has been realized for a long time that network effects play an important role in how market participants compete with each other. Arguably, companies like Facebook and Google are able to gain immense market power by leveraging the network effects of their consumers, despite potential competitors. In "Duopoly Competition with Network Effects in Discrete Choice Models," N. Chen and Y-J. Chen investigate how the dynamics play out in duopoly competition. They find that when the network effects per unit of consumption are weak, the competitors can co-exist and gain even market shares. As network effects become stronger, it is unstable, and even impossible, for the firms to coexist, and one firm emerges victorious, taking the majority of the market. The study provides a theoretical analysis for commonly observed market phenomena. It may also have implications for antitrust legislation: Special policies need to be created to maintain a competitive market structure for products and services with strong network effects.

Dynamic Pricing with Limited Resource and Unknown Demand Parameters

When the underlying demand model's parameters are unknown, how should a seller adjust prices of multiple products which are subject to finite resource constraints? In "Technical Note—Joint Learning and Optimization of Multi-Product Pricing with Finite Resource Capacity and Unknown Demand Parameters," Chen, Jasin, and Duenyas revisit the classic learning and earning problem in the context of network revenue management with a continuum of feasible prices and develop two computationally efficient pricing heuristics. For the general parametric demand models, their first heuristic attains an asymptotic regret bound which exactly matches the known theoretical lower bound under any feasible pricing control. If the underlying demand model family also satisfies a separable structure, their second heuristic leverages this

structural property to attain a much better regret. The key design features underpinning the two heuristics could be powerful ideas for developing effective pricing policies in other applications.

Measuring People's Beliefs is Really Simple.

People's beliefs about uncertain events play a pivotal role in real-world decisions. Examples include entrepreneurs who have to assess the chance that their business activities will be successful, patients who have to decide on a risky treatment, and policy makers who have to assess the risks of climate change. Existing methods to measure beliefs are complex and make restrictive assumptions that limit their use in assisting and understanding everyday decision-making. In "Measuring Beliefs Under Ambiguity," Abdelloui, Bleichrodt, Kemel, and l'Haridon propose a simple method to measure beliefs that is valid under general assumptions and that requires only three measurements. The ease of the method makes it straightforward to use in practice. An experiment of expected temperatures shows that the beliefs elicited by the method are well calibrated and similar to those elicited by more complex and time-consuming methods. This study provides an easy tool to help both individuals and policy makers make better decisions.

Games with Multiplicative Payoffs

Trade is win-win: The buyer gets a product that the seller can provide more easily. Selling more and buying more benefits both. This mutual benefit can sometimes be represented multiplicatively. But the buyer also pays a price to the seller. This money transfer is win-lose, a zero-sum game. The sum of the two payoff matrices is a matrix of rank one. In "Fast Algorithms for Rank-1 Bimatrix Games," Adsul, Garg, Mehta, Sohoni, and von Stengel show new algorithms that analyze such rank-1 games. The computation time for finding one Nash equilibrium is polynomial. Finding all Nash equilibria takes time comparable to finding a single equilibrium of a general bimatrix game. This puts games in computational reach that are economically much more interesting than zero-sum games (which are of rank zero), whereas solving general games is computationally hard. The paper's detailed structural analysis of games based on their rank enhances our understanding of this computational complexity.

Hardness of Making Rational Group Decisions

In "Bayesian Decision Making in Groups is Hard," Hązła, Jadbabaie, Mossel, and Rahimian study the computational complexity of popular models of rational group decision making. In such models, rational agents exchange their beliefs or actions over a social network to collectively perform inference or make a decision. The inference could

be on the suitability of a political candidate, quality of a product, or an important group decision in a firm or organization. The authors show that there are worst-case examples of networks and distribution of private observations that make such inferences computationally hard. They also describe a natural search algorithm to compute agents' actions through elimination and show that for certain network structures satisfying a transitivity property, the computations simplify. Although behavioral experiments had already shown that making rational decisions is hard for even small groups of individuals, this is the first result to show that such problems are inherently hard in a computational sense.

Patient-Type Bayes-Adaptive Treatment Plans

Treatment decisions that explicitly consider patient heterogeneity can lower the cost of care and improve outcomes by providing the right care for the right patient at the right time. "Patient-Type Bayes-Adaptive Treatment Plans," by Skandari and Shechter analyze the problem of designing ongoing treatment plans for a population with heterogeneity in disease progression and response to medical interventions. The authors create a model that learns the patient type by monitoring patient health over time and updates a patient's treatment plan according to the information gathered. They formulate the problem as a multivariate state space partially observable Markov decision process (POMDP). They provide structural properties of the optimal policy and develop several approximate policies and heuristics to solve the problem. As a case study, they develop a data-driven decision-analytic model to study the optimal timing of vascular access surgery for patients with progressive chronic kidney disease. They provide further policy insights that sharpen existing guidelines.

Matching While Learning

Platforms face a *cold start problem* whenever new users arrive: namely, the platform must learn attributes of new users (explore) in order to match them better in the future (exploit). How should a platform handle cold starts when there are limited quantities of the items being recommended? For instance, how should a labor market platform match workers to jobs over the lifetime of the worker, given a limited supply of jobs? In this setting, there is one multiarmed bandit problem for each worker, coupled together by the constrained supply of jobs of different types. A solution is developed to this problem. In "Matching While Learning," by Johari, Kamble, and Kanoria, it is found that the platform should estimate a shadow price for each job type, and for each worker, adjust payoffs by these prices (i) to balance learning with payoffs early on and (ii) to myopically match them thereafter.