Operations Research

In This Issue

To cite this article:
(2020) In This Issue. Operations Research 68(1):ii-iv. https://doi.org/10.1287/opre.2019.1975

http://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Vol. 68, No. 1, January–February 2020, pp. ii–iv
ISSN 0030-364X (print), ISSN 1526-5463 (online)

In This Issue

https://doi.org/10.1287/opre.2019.1975

Can We Model Blackouts?

In the article "Modeling Dependent Outages of Electric Power Plants," Malladi, Mendoza-Arriaga, and Tompaidis study simultaneous outages of many electric power plants, a problem that can have severe economic consequences. Using real data, they separate power plants into possibly overlapping groups. Each group is affected by a single cause—input fuel, geographic proximity, or location relative to the transmission network, for example. Using the mathematical technique of time subordination, the authors build dependence in outages among power plants in each group. With outage data from North American Electric Reliability Corporation, they determine these groups for Texas and the western United States and calculate the impact on system reliability and on the addition of new generation assets. Looking ahead, the methodology in the article can help account for dependence in generation from renewable sources as their share of generation capacity increases. The article will be of interest to owners of generating assets, regulators, and energy and reliability researchers.

Query Policies for Stochastic Matching with Application to Kidney Exchange

Uncertainty abounds in many general matching settings: a kidney donor might be compatible with a patient; a romantic interest might make for a good date; or an applicant might make for a profitable hire. In each setting, more can be learned about the feasibility of a potential match, but at some cost. In "Ignorance Is Almost Bliss: Near-Optimal Stochastic Matching with Few Queries," Blum, Dickerson, Haghtalab, Procaccia, Sandholm, and Sharma study this stochastic matching problem, where the goal is to find a maximum matching in a graph whose edges are unknown but can be accessed via queries. We provide polynomial-time adaptive and nonadaptive algorithms that provably achieve a large fraction of the optimal solution and use a small number of queries—for matching and its generalization to cycle packing. The authors' results apply directly to kidney exchange. We show on real data from the UNOS nationwide kidney exchange that even a very small number of nonadaptive edge queries per vertex results in large gains in expected successful matches.

Forward Markets Can Create Opportunities for Market Manipulation

The role of forward contracts in mitigating market power is an important and recurring theme in oligopolistic industries, such as electricity and gas. In "On the Inefficiency of Forward Markets in Leader–Follower Competition," Cai, Agarwal, and Wierman study the impact of forward contracting in situations where firms have capacity constraints and heterogeneous production lead times. The authors introduce a model that combines forward contracting and leader–follower competition and explicitly characterize the equilibrium outcomes with and without forward contracting. They show that the impact of forward markets is delicate—it may mitigate market power or create opportunities for market manipulation. In particular, leader–follower interactions due to heterogeneous production lead times may cause forward markets to be inefficient, even when there are a large number of followers. In fact, symmetric equilibria do not necessarily exist due to differences in market power among leaders and followers.

Optimal Recovery of a Hidden Hamiltonian Cycle with Applications in Genome Sequencing

One of the most pressing challenges in genomics is to reconstruct a long and contiguous genome sequence from short genome subsequences (contigs). Recent developments in sequencing technology obtain long-range linking measurements among the contigs. In "Hidden Hamiltonian Cycle Recovery via Linear Programming," Bagaria, Ding, Tse, Wu, and Xu propose a probabilistic model for the problem of contig assembly, where the goal is to recover a Hamiltonian cycle hidden in a complete weighted graph based on noisy edge measurements. Computing the maximum likelihood estimate in this model reduces to the traveling salesman problem (TSP). Despite the NP-hardness of TSP, they show that a simple linear programming (LP) relaxation—namely, the fractional 2-factor LP—recovers the hidden Hamiltonian cycle with high probability, whenever it is information-theoretically possible. Evaluation of the algorithm on real data shows improvements over existing approaches.

Can Public Reporting Cure Healthcare?

In the past two decades, government and private organizations have ramped up efforts to measure and publicly report clinical outcomes, a practice known as "public reporting." For example, in 2005, the Center for Medicare and Medicaid Services launched the website Hospital Compare with shared outcome data. Other sites, including CalHospitalCompare.org, the ProPublica Surgeon Scorecard, the Compare Hospitals site by the Leapfrog Group, and the hospital-rating websites of Healthgrades, Consumer Reports, Yelp, and U.S. News, have provided similar information. Despite making more and more information available, empirical evidence suggests that the reality of public reporting has fallen short of the promise. In "Can Public Reporting Cure Healthcare? The Role of Quality

Transparency in Improving Patient–Provider Alignment," Saghafian and Hopp use a model of the impact of quality transparency via public reporting to generate insights into why previous efforts have been less than fully successful and to suggest ways in which future efforts can be more effective.

When Debt Becomes Explosive: Insights from Game Theory

In "Dynamic Leveraging–Deleveraging Games," Minca and Wissel endogenize the dynamics of a large borrower's leverage based on a new type of game among the lender, the leveraging/deleveraging game. Leverage is mean reverting around a long-run level and explosive above an instability level. This is driven by the changing nature of the lenders' game from strategic substitutability to one-sided strategic complementarity: When leverage is below the instability level, the firm is not in danger of default, and the strategic substitutability in the payoff structure acts as a counterbalance on leverage, which is pushed down to the long-run level. If the leverage reaches above the instability level, then it becomes explosive: The strategic complementarity leads to spiraling effects that end in default. The authors validate their model empirically using aggregate returns of financial firms over the last two decades.

Setting Prices for Multiple Substitutable Products with Diffusion

Whether introducing and selling a new book with multiple product forms (hardcover, paperback, and electronic), smart phones with many feature variations, or multiple generations of computer microprocessors, firms encounter a challenging pricing problem. The complexity arises from interactions among the products and from the dynamics of the product diffusion process, which creates dependency of future sales on current prices. In "Optimal Pricing Under Diffusion-Choice Model," Li examines pricing in a multiproduct diffusion process, in which customer choice probabilities among the products are given by the multinomial logit model. The author presents a novel solution approach, characterizes the time path of the optimal prices, and shows how the diffusion dynamics and product-specific differences enter the optimal solution. The model applies to both simultaneous and sequential product introductions and accommodates changes to the product portfolio as a firm introduces or retires products during the planning horizon.

Ranking and Pricing Products for E-Markets

In "Approximation Algorithms for Product Framing and Pricing," Gallego, Li, Truong, and Wang propose one of the first models of "product framing" and pricing. Product framing refers to the way consumer choice is influenced by how the products are framed or displayed. They present a model in which a set of products are displayed or framed into a set of virtual web pages. They assume that consumers consider only products in the top pages with different consumers willing to see different numbers of pages. Consumers select a product, if any, from these pages following a general choice model. They show that the product-framing problem is NP-hard. They derive algorithms with guaranteed performance relative to an optimal algorithm under reasonable assumptions. The algorithms are fast and easy to implement. They also present structural results and design algorithms for pricing under framing effects for the multinomial logit model. They show that, for profit-maximization problems, at optimality, products are displayed in descending order of their value gap and in ascending order of their markups.

Analyzing the Performance of Sequential Zone Picking Systems

In "Capacity Analysis of Sequential Zone Picking Systems" by van der Gaast, de Koster, Adan, and Resing, a capacity model for sequential zone picking systems is developed. These systems are popular internal transport and order-picking systems in practice. The major disadvantage of such systems is congestion and blocking under heavy use, leading to long order throughput times. To reduce blocking and congestion, most systems use the block-and-recirculate protocol to dynamically manage workload. In this paper, the various elements of the system are modeled as a multiclass block-and-recirculate queueing network with capacity constraints on subnetworks. Because of this blocking protocol, the stationary distribution of the queueing network is highly intractable. The authors propose an approximation method based on jump-over blocking that allows realistically sized systems to be efficiently analyzed and can be used in the design phase of sequential zone picking systems. The results show that the relative error in the system throughput is typically less than 1% compared with simulation.

Solving Different Variants of the Multitrip Vehicle-Routing Problems with a Single Exact Solution Method

The need to reduce pollution and traffic in city centers requires the use of small vans, electric vehicles, and drones to distribute goods. Because of autonomy and capacity issues, these vehicles need to perform multiple trips from/to the depot during the day. The category of decision-making problems modeling such distribution problems are known as multitrip vehicle-routing problems (MTVRPs), which generalize the well-known vehicle-routing problem by allowing vehicles to perform multiple trips per day. Several MTVRPs are solved in the literature with different mathematical models and algorithms. In "An Exact Solution Framework for Multitrip Vehicle-Routing Problems with Time Windows," Paradiso, Roberti, Laganá, and Dullaert propose a single algorithm that can solve, to optimality, the MTVRP with capacity and time windows constraints and four variants of this problem featuring additional operational constraints. The proposed framework significantly outperforms the state-of-the-art algorithms from the literature.

On the Influence of Reference Point Updating on Decision Making in the Context of Portfolio Optimization

In a dynamic setting, decision makers update their reference point as a function of previous decisions and outcomes. In "Failing to Foresee the Updating of the Reference Point

Leads to Time-Inconsistent Investment," Strub and Li investigate the influence of reference point updating on decision making and, in particular, address whether a decision maker foresees the updating of the reference point in the context of discrete time portfolio optimization. By deriving and comparing optimal trading strategies under various frameworks and reference point updating rules and then simulating what typical investment behavior would look like in each setting, they come to the following conclusion: a loss-averse decision maker with a time-varying reference point exhibiting realistic investment behavior fails to foresee the future updating of the reference point, and this failure leads to time-inconsistent investment.

Online Algorithms for Hierarchical Aggregation Problems

Data and inventory aggregation problems arise in multicasting, sensor networks, communication in organization hierarchies, and in supply chain management. These problems are naturally online, in the sense that aggregation decisions need to be made without information about future requests. In "Online Algorithms for Multilevel Aggregation," Bienkowski, Böhm, Byrka, Chrobak, Dürr, Folwarczný, Jeż, Sgall, Thang, and Veselý study these problems with a general tree structure of links that can be used for deliveries. This generalizes some well-studied optimization problems: trees of depth one capture the Transmission Control Protocol acknowledgment problem, and trees of depth two capture the joint replenishment problem. For trees of depths one and two, constant-competitive online algorithms are known. The authors solve a major open problem by giving a constant-competitive algorithm for trees of arbitrary (fixed) depth. The algorithm works for arbitrary waiting cost functions, including the variant with deadlines.

Making Emulation Faster

How can one make a large and complex model fast and "small"? The simulation literature has extensively addressed this problem, and the kriging method has proven to be one of the most successful methods to deal with complex simulators. In "Faster Kriging: Facing High-Dimensional Simulators," Lu, Rudi, Borgonovo, and Rosasco propose a new kriging implementation, called "fast kriging," that copes with dimensionality issues, allowing one to deal with data sets coming from simulators with thousands of inputs.

Models for Patient-Flow in Inpatient Wards

Long delays for patients waiting to be admitted to inpatient units are currently the standard state of affairs in urban hospitals in the United States and around the world. To alleviate congestion, hospital management seeks ways to increase efficiency by improving the flow of patients between the different hospital units. To this end, effective queueing models that are specifically tailored to the hospital setting are required. In "Queueing Models for Patient-Flow Dynamics in Inpatient Wards," Dong and Perry propose a new class of queueing models for inpatient wards, taking into account the salient features associated with patient-flow dynamics. The authors demonstrate how those models can be employed to compute the optimal number of beds that are needed or to measure the impacts of changes to discharge process.

Personalized Decision Making with Big Data

Decision makers increasingly have access to rich customer-specific data, providing an opportunity to make better, *personalized* service decisions. For example, in healthcare, doctors can personalize interventions based on a patient's clinical history; and in marketing, companies can target ads based on customer purchase history. However, the increased variety of potentially relevant customer data implies that an individual's covariates may be *high dimensional*, which, in turn, poses statistical challenges for learning personalized decision-making policies. In "Online Decision Making with High-Dimensional Covariates," Bastani and Bayati introduce the LASSO Bandit, an adaptive decision-making algorithm that efficiently leverages high-dimensional user covariates by learning *sparse* models of decision rewards. The authors illustrate the practical relevance of such an approach by evaluating it against a personalized medication dosing problem, finding that the LASSO Bandit outperforms existing bandit methods and physicians in correctly dosing a majority of patients.

Bayesian Learning Works Even with Censored Information

We often learn about a problem from information that is incomplete or censored: for example, a medical treatment may cause side effects with no indication of what the right dose should have been. Bayesian belief models are useful in such settings, but cannot be constructed using traditional methods; as a result, practitioners have developed ways of constructing them approximately. These approximations have been very successful in many application domains, yet until now have lacked theoretical support. The paper "Consistency Analysis of Sequential Learning Under Approximate Bayesian Inference," by Chen and Ryzhov links approximate Bayesian learning to stochastic approximation theory. Using this link, the authors prove—for the first time—the consistency of a suite of approximate Bayesian methods culled from the literature. One highlight is an entirely new consistency proof for Bayesian logistic regression, a well-established approximation technique that essentially treats logistic regression as if it were ordinary least squares.