



Operations Research

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

In This Issue

To cite this article:

(2022) In This Issue. Operations Research 70(1):iii-viii. <https://doi.org/10.1287/opre.2021.2252>

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 © 2022, 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.2021.2252>

Copyright: © 2022 INFORMS

Bots Impact Opinions in Social Networks: Let's Measure How Much

There is a serious threat posed by bots that try to manipulate opinions in social networks. In "Assessing the Impact of Bots on Social Networks," des Mesnards, Hunter, el Hjouji, and Zaman present a new set of operational capabilities to detect these bots and measure their impact. They developed an algorithm based on the Ising model from statistical physics to find coordinating gangs of bots in social networks. The authors then created an algorithm based on opinion dynamics models to quantify the impact that bots have on opinions in a social network. They applied their algorithms to a variety of real social network data sets. The authors found that, for topics such as Brexit, the bots had little impact, whereas for topics such as the U.S. presidential debate and the Gilets Jaunes protests in France, the bots had a significant impact.

The Implications of Regime-Dependent Risk Aversion for Dynamic Portfolio Allocation

Despite the overwhelming evidence of time-varying risk aversion documented in recent literature, standard dynamic portfolio theories often adopt an assumption of constant (relative) risk aversion due to analytical tractability. In "Time Varying Risk Aversion and Dynamic Portfolio Allocation," Li, Wu, and Zhou explicitly consider the implications of time-varying risk aversion for dynamic portfolio allocation under the framework of regime-switching models. An investor with regime-dependent utility exhibits a decreasing relative risk aversion (DRRA) and has higher risk aversion when a bear market regime is more likely in the future. They develop an efficient dynamic programming algorithm that overcomes the challenges imposed by regime-dependent preference in obtaining time-consistent portfolio policies. The empirical results show that VIX is an important predictor of regime shifts and that investors with regime-dependent risk aversion achieve better investment performance than those with constant risk aversion.

Optimal Design of Combined Contingent Claims: Theory and Applications.

In "Combined Custom Hedging: Optimal Design, Noninsurable Exposure, and Operational Risk Management," Guiotto and Roncoroni develop a normative framework for the optimal design, value assessment, and operations management integration of financial derivatives. Most business and operating revenues entail a mix of financially insurable and noninsurable risk. A risk-averse firm may face them by positioning in a pair of financial derivatives with optimal bespoke payoff functions; one claim is written on the insurable term, and the other claim is written on any observable index exhibiting correlation to the noninsurable term. On a theoretical ground, the authors (1) state the problem in a general setup and prove existence and uniqueness of the optimal pair of combined claims, (2) show that the optimal payoff functions satisfy a Fredholm integral

equation, and (3) assess the incremental benefit the firm obtains by switching from the optimal single-claim custom hedge to the optimal combined custom hedge they propose. On an experimental ground, they show that (1) the optimal combined custom hedge would be empirically relevant for a highly risk-averse firm facing a market shock shown during the first period of the COVID-19 pandemic in 2020, (2) integration with the optimal procurement in a generalized newsvendor model leads to a significant improvement in both risk and return, and (3) this gain can be traded off for a substantial enhancement in operational flexibility.

Optimal Portfolio Diversification via Independent Component Analysis

Factor-risk parity is a popular investment strategy that consists of spreading the risk of a portfolio equally across a set of uncorrelated factors such as the principal components. In "Optimal Portfolio Diversification via Independent Component Analysis," Lassance, DeMiguel, and Vrin study the impact of the choice of uncorrelated factors on the performance of this approach. They find that the performance is sensitive to the choice of factors because any feasible portfolio is a factor-risk-parity portfolio for some uncorrelated factors. Moreover, relying on the principal components is arbitrary and suboptimal from a diversification perspective. Instead, they propose relying on the independent components, which are the unique set of factors that are as independent as possible. Finally, they show theoretically that the resulting portfolios have favorable tail-risk properties, and empirically that they perform well in terms of Sharpe ratio, tail risk, and turnover.

Using Data to Allocate Resources Efficiently

In city logistics systems, a fleet of vehicles is divided between service regions that function autonomously. Each region finds optimal routes for its own fleet and incurs costs accordingly. More vehicles lead to lower costs, but the trade-off is that fewer vehicles are left for other regions. Costs are difficult to quantify precisely because of demand uncertainty but can be estimated using data. In "Data-Driven Robust Resource Allocation with Monotonic Cost Functions" Chen, Marković, Ryzhov, and Schonfeld develop a principled risk-averse approach for two-stage resource allocation. The authors propose a new uncertainty model for decreasing cost functions and show how it can be leveraged to efficiently find resource allocations that demonstrably reduce the frequency of high-cost scenarios. This framework combines statistics and optimization in a novel way and is applicable to a general class of resource allocation problems, encompassing facility location, vehicle routing, and discrete-event simulation.

Robustness in Risk Measurement: The Impact of Incentives

Statistical robustness is a desirable property for a regulatory risk measure. Previous research has stressed that Value-at-Risk is more robust than Expected Shortfall if both are applied to the same financial position. In reality, however, the regulatory choice of a particular risk measure imposes certain incentives, which impact the underlying position even before a particular risk measure is applied. In “Robustness in the Optimization of Risk Measures,” Embrechts, Schied, and Wang propose a first attempt of taking such incentives into account when assessing a risk measure’s robustness properties. To this end, they develop a general methodology which they call “robustness against optimization”. The new notion is studied for various classes of risk measures and expected utility and loss. In particular, in this new notion, Value-at-Risk appears to be less robust than Expected Shortfall and many other coherent and convex risk measures, revealing a serious drawback of Value-at-Risk in the context of optimization, which was previously overlooked.

1.79-Approximation Algorithms for Two Classical Inventory Models with Lead Times

Stochastic inventory systems with lead times are often challenging to optimize, including single-sourcing lost-sales and dual-sourcing systems. Recent numerical results suggest that capped policies demonstrate superior performance over existing heuristics. However, the superior performance lacks a theoretical foundation. In “1.79-Approximation Algorithms for Continuous Review Single-Sourcing Lost-Sales and Dual-Sourcing Inventory Models,” Xin provides a theoretical foundation for this phenomenon in two classical inventory models. First, in a continuous review lost-sales model with lead times and Poisson demand, he proves that a capped base-stock policy has a worst-case performance guarantee of 1.79 by conducting an asymptotic analysis under a large penalty cost and lead time. Second, in a more complex continuous review dual-sourcing model with general lead times and Poisson demand, he proves that a similar capped dual-index policy has a worst-case performance guarantee of 1.79 under large lead time and ordering cost differences. The results provide a deeper understanding of the superior numerical performance of capped policies and present a new approach to proving worst-case performance guarantees of simple policies in hard inventory problems.

Adversarial Patrolling in a Uniform

Many of us will have seen uniformed guards patrolling in museums, airports, and other places where thefts or attacks are possible. In other similar places, we may have been unaware of undercover agents carrying out similar patrols. The latter type of patrollers has been modeled in recent literature. However, when the patroller is visible (wears a uniform), the potential thief or terrorist may be able to spot him when he is nearby and to time his theft appropriately. For example, the thief may wait a specified time after the uniformed patroller leaves his area. In “Adversarial Patrolling in a Uniform,” Alpern, Chleboun, Katsikas, and Lin study the effect on the patrolling game of having a visible (uniformed) patroller. It turns out that putting a uniform on the patroller greatly reduces his effectiveness in intercepting thefts or attacks. Of course, the visibility of a uniform may act as a deterrent to having the theft take place at all.

Simple Algorithms for Complex Multiwarehouse, Multistore Inventory Control Problems

Retailers (both brick-and-mortar and e-commerce) have always faced the problem of allocating inventories in their warehouses

(or central distribution centers) to the stores (or smaller local warehouses) in order to minimize total costs. The problem is particularly challenging when the network structure is large and complex, the selling season is long, and the replenishment is frequent. For example, giant retail chains such as Macy’s typically have many warehouses and hundreds of stores across the United States, and online retailers such as Amazon have many distribution centers and over one hundred fulfillment centers. In “Asymptotically Optimal Lagrangian Policies for Multi-Warehouse, Multi-Store Systems with Lost Sales,” Miao, Jasin and Chao develop algorithms to solve this multiwarehouse, multistore (MWMS) inventory control problem. Their algorithms are computationally efficient and asymptotically optimal as the problem becomes large and complex. This feature is very appealing to today’s fast-moving retail industry with rapidly expanding business scale.

Constrained Shortest Path in Big Networks

Motivated by the needs of modern transportation service platforms, in “Computing Constrained Shortest-Paths at Scale,” Vera, Banerjee, and Samaranyake study the problem of computing constrained shortest paths (CSP) at scale via preprocessing techniques. Our work makes two contributions in this regard: (1) They propose a scalable algorithm for CSP queries and show how its performance can be parametrized in terms of a new network primitive, the constrained highway dimension. This development extends recent work that established the highway dimension as the appropriate primitive for characterizing the performance of unconstrained shortest-path (SP) algorithms. Their main theoretical contribution is deriving conditions relating the two notions, thereby providing a characterization of networks where CSP and SP queries are of comparable hardness. (2) The authors develop practical algorithms for scalable CSP computation, augmenting our theory with additional network clustering heuristics. They evaluate these algorithms on real-world data sets to validate our theoretical findings. Their techniques are orders of magnitude faster than existing approaches while requiring only limited additional storage and preprocessing.

Profit and Loss of (Selfish) Blockchain Miners

Mining blocks on public blockchains equipped with the proof-of-work protocol is a risky business as a constant operational cost is compensated by infrequent incomes generated by block discoveries. The variance in the miners’ revenue entails a risk of insolvency. In “On the Profitability of Selfish Blockchain Mining Under Consideration of Ruin,” Albrecher and Goffard introduce a mathematical framework, inspired by insurance risk theory, to perform cost-benefit analysis and inform the decision-making process of blockchain miners. The expressions of the risk and performance indicators are amenable for numerical evaluations and allow the authors to study the opportunity for miners to implement a blockwithholding strategy called selfish mining. The main takeaway is that deviating from the prescribed protocol always endangers the miner’s solvency but may be advisable under some model settings in terms of expected profit.

Quickest Multiclass Classification

In multiclass classification, one faces greater uncertainty when the data fall near the decision boundary. To reduce the uncertainty, one can wait and collect more data, but this invariably delays the decision. How can one make an accurate classification as quickly as possible? The solution requires a multiclass generalization of Wald’s sequential hypothesis testing, but the standard formulation is intractable because of the curse of

dimensionality in dynamic programming. In “Optimal Sequential Multiclass Diagnosis,” Wang shows that, in a broad class of practical problems, the reachable state space is often restricted on, or near, a set of low-dimensional, time-dependent manifolds. After understanding the key drivers of sparsity, the author develops a new solution framework that uses a low-dimensional statistic to reconstruct the high-dimensional state. This framework circumvents the curse of dimensionality, allowing efficient computation of the optimal or near-optimal policies for quickest classification with large numbers of classes.

Fast Core Pricing for Rich Advertising Auctions

Standard ad auction formats do not immediately extend to settings where multiple size configurations and layouts are available to advertisers. In these settings, the sale of web advertising space increasingly resembles a combinatorial auction with complementarities, where truthful auctions such as the Vickrey–Clarke–Groves (VCG) auction can yield unacceptably low revenue. In “Fast Core Pricing for Rich Advertising Auctions,” Niazadeh, Hartline, Immorlica, Khani, and Lucier study and suggest core-selecting auctions, which boost revenue by setting payments so that no group of agents, including the auctioneer, can jointly improve their utilities by switching to a different outcome. Their main result is a combinatorial algorithm that finds an approximate bidder-optimal core point with an almost linear number of calls to the welfare-maximization oracle. This algorithm is faster than previously proposed heuristics in the literature and has theoretical guarantees. By accompanying the theoretical study with an experimental study based on Microsoft Bing Ad Auction data, the authors conclude that core pricing is implementable even for very time-sensitive practical use cases such as real-time online advertising and can yield more revenue than the VCG or generalized second price auction.

Core-Pricing in Combinatorial Exchanges with Budget Constraints

The computation of market equilibria is a fundamental and practically relevant problem. Although we know the computational complexity and the types of price functions necessary for combinatorial exchanges with quasilinear preferences, the respective literature does not consider financially constrained buyers. In “Core Pricing in Combinatorial Exchanges with Financially Constrained Buyers: Computational Hardness and Algorithmic Solutions,” Bichler and Waldherr show that computing market outcomes that respect budget constraints but are core stable is a problem in the second level of the polynomial hierarchy. Problems in this complexity class are rare but ignoring budget constraints can lead to significant efficiency losses and instability. They introduce mixed integer bilevel linear programs (MIBLP) to compute core-stable market outcomes and provide effective column and constraint generation algorithms to solve these problems. Although full core stability quickly becomes intractable, the authors show that realistic problem sizes can be solved if the designer limits attention to deviations of small coalitions. This n -coalition stability is a practical approach to tame the computational complexity of the general problem and at the same time provides a reasonable level of stability for markets in the field where buyers have budget constraints.

Improving Newborn Screening for Genetic Diseases

Screening newborns for life-threatening genetic diseases is an important public health initiative. Cystic fibrosis is one of the most prevalent diseases in this context. As part of the cystic fibrosis screening process, all states in the United States use

multiple tests, including genetic tests that detect a subset of the more than 300 genetic variants (specific mutations) that cause cystic fibrosis. In “Optimal Genetic Screening for Cystic Fibrosis,” El-Hajj, D.R. Bish, and E.K. Bish develop a decision support model to select which genetic variants to screen for, considering the trade-off between classification accuracy and testing cost, and the technological constraints that limit the number of variants selected. Because variant prevalence rates are highly uncertain, a robust optimization framework is developed. Further, two commonly used cystic fibrosis screening processes are analytically compared, and conditions under which each process dominates are established. A case study based on published data are provided.

Fairness in Stochastic Resource Allocation

In settings where a platform must allocate finite supplies of goods to buyers, balancing overall platform revenues with the fairness of the individual allocations to platform participants is paramount to the well-functioning of the platform. This is made even more difficult by the fact that the supply of goods is in practice stochastic and difficult to forecast, such as in the case of online ad allocation, where the platform manages a supply of impressions that varies over time. In “Fair Resource Allocation in a Volatile Marketplace,” Bateni, Chen, Ciocan, and Mirrokni design a fair allocation scheme that works in the presence of supply uncertainty. Algorithmically, the scheme repeatedly solves for Fisher market equilibria in a model predictive control fashion and is proved to admit constant factor guarantees versus the offline optimal. In addition, the scheme is tested on a sequence of real ad datasets, showing strong empirical performance.

Customer Choice Models vs. Machine Learning: Finding Optimal Product Displays on Alibaba

In “Customer Choice Models vs. Machine Learning: Finding Optimal Product Displays on Alibaba,” Feldman, Zhang, Liu, and Zhang compare the performance of two approaches for finding the optimal set of products to display to customers landing on Alibaba’s two online marketplaces, Tmall and Taobao. Both approaches were placed online simultaneously and tested on real customers for one week. The first approach the authors test is Alibaba’s current practice. This procedure embeds thousands of product and customer features within a sophisticated machine-learning algorithm that is used to estimate the purchase probabilities of each product for the customer at hand. Their second approach uses a featurized multinomial logit (MNL) model to predict purchase probabilities for each arriving customer. In this way, they use less sophisticated machinery to estimate purchase probabilities, but they employ a model that was built to capture customer purchasing behavior and, more specifically, substitution patterns. Their experiments show that despite the lower prediction power of our MNL-based approach, it generates significantly higher revenue per visit compared with the current machine-learning algorithm with the same set of features.

Abstraction Methods for Large Markets

In “Computing Large Market Equilibria Using Abstractions,” Kroer, Peysakhovich, Sodomka and Stier-Moses compare the performance of two approaches for finding the optimal set of products to display to customers landing on Alibaba’s two online marketplaces, Tmall and Taobao. Both approaches were placed online simultaneously and tested on real customers for one week. The first approach we test is Alibaba’s current practice. This procedure embeds thousands of product and customer features within a sophisticated machine-learning algorithm

that is used to estimate the purchase probabilities of each product for the customer at hand. Their second approach uses a factorized multinomial logit (MNL) model to predict purchase probabilities for each arriving customer. In this way, the authors use less sophisticated machinery to estimate purchase probabilities, but they employ a model that was built to capture customer purchasing behavior and, more specifically, substitution patterns. Their experiments show that despite the lower prediction power of our MNL-based approach, it generates significantly higher revenue per visit compared with the current machine-learning algorithm with the same set of features.

Optimal Market Integration Decisions by Policy Makers: Modeling and Analysis of Agriculture Market Data

Policymakers often seek to integrate markets as a way to maximize social welfare. In “Optimal Market-Integration Decisions by Policymakers: Modeling and Analysis of Agriculture Market Data,” Gupta and Bansal consider the spectrum of all possible integration policies, from full isolation to complete integration, and characterize the socially optimal market integration, under general demands. They identify market conditions under which social surplus is indeed maximized at partial market integration. For the linear price-responsive demand model that is used extensively in the operations management literature, these conditions are identified as thresholds on (i) the relative size of the markets being integrated, and (ii) the relative price sensitivity of consumers in these markets. The authors then apply the model to the commercial seed market in the European Union (EU). Their analysis shows that socially optimal market integration for these countries provides a further improvement in the social surplus for the EU by 2.80%, relative to complete integration. Results show that policymakers should exercise caution in determining the extent to which markets are integrated.

Designing Fair and Efficient Matching Service Systems

The first-come, first-served (FCFS) service discipline provides a simple, and more importantly, fair method to allocate limited-service capacity to customers and to manage queues. At the same time, it can limit the ability of service providers to optimize this allocation based on individual customers’ preferences and servers’ characteristics. To address this limitation, providers can segment demand into classes and serve them using different subsets of servers on an FCFS basis. The hope is that, by appropriately selecting the subset of servers that may serve each customer class, the service provider can improve service quality (i.e., matching customers to their most preferred servers) without increasing (significantly) waiting times. In “On the Optimal Design of a Bipartite Matching Queueing System,” Afèche, Caldentey, and Gupta propose a queueing framework (inspired by the seminal work of Edward Kaplan on public housing) to solve this matching design problem. The proposed methodology allows service providers to construct a Pareto frontier of service menus that efficiently trade-off the quality of the matching and customers’ waiting times.

Nash Social Welfare with Trading Posts

In “Nash Social Welfare Approximation for Strategic Agents,” Brânzei, Gkatzelis, and Mehta study the problem of allocating divisible resources to agents with different preferences. They analyze a market game known as Trading Post, first considered by Shapley and Shubik, where each agent gets a budget of virtual currency to bid on goods: after bids are placed, goods are allocated to players in proportion to their bids. In this setting,

the agents choose their bids strategically, aiming to maximize their utility, and this gives rise to a game. The authors study the equilibrium allocations of this game, measuring the quality of an allocation via the Nash social welfare, the geometric mean of utilities (a measure of aggregate welfare that respects individual needs). They show that any Nash equilibrium of Trading Post approximates the optimal Nash welfare within a factor of two for all concave valuations, and the mechanism is essentially optimal for Leontief valuations.

Dynamic Resource Allocation with Capacity Constraints

In “A Restless Bandit Model for Resource Allocation, Competition and Reservation,” Fu, Moran, and Taylor study a resource allocation problem with varying requests and with resources of limited capacity shared by multiple requests. This problem is modeled as a set of heterogeneous restless multi-armed bandit problems (RMABPs) connected by constraints imposed by resource capacity. Following Whittle’s idea of relaxing the constraints and Weber and Weiss’s proof of asymptotic optimality, the authors propose an index policy and establish conditions for it to be asymptotically optimal in a regime where both arrival rates and capacities increase. In particular, they provide a simple sufficient condition for asymptotic optimality of the policy and, in complete generality, propose a method that generates a set of candidate policies for which asymptotic optimality can be checked. Via numerical experiments, they demonstrate the effectiveness of these results even in the pre-limit case.

Solving Large-Scale Ranking and Selection Problems in Parallel Computing Environments

With the rapid development of computing technology, using parallel computing to solve large-scale ranking and selection (R&S) problems has emerged as an important research topic. Inspired by the knockout-tournament structures of tennis Grand Slam tournaments, in “Knockout-Tournament Procedures for Large-Scale Ranking and Selection in Parallel Computing Environments,” Zhong and Hong propose new R&S procedures to solve large-scale problems in parallel computing environments. They show that their procedures can theoretically achieve the lowest growth rate on the expected total sample size with respect to the number of alternatives and thus, are optimal in rate. Moreover, common random numbers can be easily adopted in their procedure to further reduce the sample size. Meanwhile, the comparison time in their procedures is negligible compared with the simulation time, and their procedures barely request for communications among the processors.

Data-Driven Optimization Using Reproducing Kernel Hilbert Spaces

A fundamental problem in operations research is to minimize a cost function with uncertain parameters given some data of historical realizations of the parameters along with auxiliary data of realizations of other variables used to predict the parameters. Existing proposals to solve this problem are based on “local” machine learning techniques, such as nearest neighbors, that suffer from a curse of dimensionality, requiring an exponentially increasing amount of data to maintain their performance as the number of auxiliary variables increases. In “Data-Driven Optimization: A Reproducing Kernel Hilbert Space Approach,” Bertsimas and Koduri use the theory of reproducing Kernel Hilbert spaces to develop two new “global” methods for solving the problem. They provide theoretical finite sample guarantees and prove asymptotic optimality for both methods. In computational experiments, they find that one of the new

methods does overcome the curse of dimensionality, requiring much less data to achieve the same performance as existing methods even in small dimensions.

Regulating Network Pricing Games by Using Price Caps

Network pricing games provide a framework for modeling real-world settings with two types of strategic agents: operators and users of a network. A landmark example is road operators that set tolls for the usage of their road, whereas traffic users select routes that are cheap yet not too congested. Competition between the operators induces several issues in both the operators' and users' market. In "Network Pricing: How to Induce Optimal Flows Under Strategic Link Operators" by Correa, Guzmán, Lianes, Nikolova, and Schröder, all these issues are resolved by setting appropriate upper bounds (price caps) for each of the network operators. It is proven that great tolls (as price caps) exist and can be computed by means of a linear program.

Wasserstein Inverse Covariance Shrinkage Estimator

The optimal solutions of many decision problems such as the Markowitz portfolio allocation and the linear discriminant analysis depend on the inverse covariance matrix of a Gaussian random vector. In "Distributionally Robust Inverse Covariance Estimation: The Wasserstein Shrinkage Estimator," Nguyen, Kuhn, and Mohajeri Esfahani propose a distributionally robust inverse covariance estimator, obtained by robustifying the Gaussian maximum likelihood problem with a Wasserstein ambiguity set. In the absence of any prior structural information, the estimation problem has an analytical solution that is naturally interpreted as a nonlinear shrinkage estimator. Besides being invertible and well-conditioned, the new shrinkage estimator is rotation equivariant and preserves the order of the eigenvalues of the sample covariance matrix. If there are sparsity constraints, which are typically encountered in Gaussian graphical models, the estimation problem can be solved using a sequential quadratic approximation algorithm.

Steady State in Equilibrium for Flows Over Time

Understanding the equilibrium behavior of dynamic traffic models, where traffic flows vary over time, is of great practical as well as theoretical interest. In "Long-Term Behavior of Dynamic Equilibria in Fluid Queuing Networks," Cominetti, Correa, and Olver consider a basic but important dynamic traffic model which has received significant attention within both operations research and transportation economics. They study the long-term behavior of equilibrium solutions and prove that eventually the system will reach a steady state situation which can be characterized as the solution of a certain linear program.

Dynamic Scheduling to Differentiate Delay-Based Service Levels in Multiclass Service Systems

Tail probability of delay (TPoD), defined as the probability that the customer delay exceeds a customary delay target, is widely used as a performance metric in many real-world service systems. Examples include the 80–20 rule in customer contact centers and the Canadian triage and acuity scale (CTAS) guideline that classifies patients in the emergency department into five acuity levels. In those settings, how to ensure that diverse customer needs are met based on prescribed TPoD targets via effective capacity planning and dynamic resource allocation has been deemed notoriously difficult. In response to such a

challenge, in "Scheduling to Differentiate Service in a Multiclass Service System," Liu, Sun, and Hovey study a practical queueing system having multiple customer classes, nonstationary customer arrivals, and customer abandonment. They devise a new class of staffing (number of servers) and scheduling (assigning newly idle servers to a waiting customer from one of the classes) policies that can help achieve class-differentiated service levels measured by TPoD. This newly proposed class of control rules not only exhibits nice separation of scales under appropriate heavy-traffic scaling, but also gives rise to novel heavy-traffic stochastic-process limits for delay-related performance measures. The effectiveness of their approach is substantiated by heavy-traffic asymptotic stability theorems and extensive numerical studies in which important managerial insights are also generated.

A Unifying Token-Based Framework for Product-Form Queue Models

In the recent literature, parallel server queueing models with different customer types that have a product-form stationary distribution have received considerable interest. For these models, the queue's stationary distribution can be expressed as a product of terms, each of which corresponds to a customer in the system. Examples of such models arise naturally in applications such as matching systems and redundant scheduling in computer clusters. In "A Token-Based Central Queue with Order-Independent Service Rates," Ayesta, Bodas, Dorsman, and Verloop propose a token-based queueing model that subsumes many of these models, while it retains a product-form stationary distribution. Furthermore, the authors derive novel expressions for a variety of performance measures for each of the underlying models.

The Limitations and Opportunities for Using Correlation in Mechanism Design

Traditionally, much of the focus of the mechanism/auction design community has been on revenue optimal mechanisms for settings where bidders' private valuations over outcomes can be reasonably thought of as independent of each other. This has been the case even though there is good reason to believe that valuations are often correlated and there are theoretical results suggesting that mechanisms designed with this correlation in mind can generate much higher revenue. In "Mechanism Design for Correlated Valuations: Efficient Methods for Revenue Maximization," Albert, Conitzer, Lopomo, and Stone look at the setting where there is correlation, but the exact distribution is unknown and must be estimated from samples. They show that in this setting, the previous extremely strong theoretical results around the usefulness of correlation are now very sensitive to the degree of correlation in the underlying distribution and the number of samples that the mechanism designer has access to. However, the authors also show that if correlation is sufficient, we can construct mechanisms, using a computationally efficient procedure, that significantly outperform traditional mechanism design paradigms.

The Uniqueness of a Mean Field Equilibrium

The lack of a unique equilibrium has significantly limited game theory models' predictive power and their ability to derive robust comparative statics results. Mean field equilibrium has received significant attention in the last decade as a solution concept for dynamic stochastic games because of its computational tractability and behavioral appeal resulting from the reduced assumptions on players' rationality. In "Mean Field Equilibrium: Uniqueness, Existence, and Comparative Statics," Light

and Weintraub provide conditions that ensure that a mean field equilibrium is unique and derive general comparative statics results in the context of discrete-time mean field games. The paper's existence, uniqueness, and comparative statics results are applied to various models from the economics and operations research literature, including quality ladder models, capacity competition models, advertising competition models, dynamic reputation models, and heterogeneous agent macroeconomic models.

Learning to Approximate Industrial Problems by Operations Research Classic Problems

Operations research (OR) practitioners are accustomed to dealing with variants of classic OR problems. Indeed, an industrial problem often looks like a traveling salesman problem, a vehicle routing problem, a shortest path problem, etc., but has an additional constraint or a different objective that prevent the use of the powerful algorithms produced by decades of research on the classic OR problems. This situation can be frustrating, notably when we realize that the classic problem catches most of the structure of the variant. In "Learning to Approximate Industrial Problems by Operations Research Classic Problems," Parmentier introduces a machine learning approach to use the algorithms for the classic OR problems on the variant.

The idea is to leverage structured learning to obtain a mapping that approximates an instance of the variant by an instance of the classic problem.

Two-Stage Sample Robust Optimization

In "Technical Note—Two-Stage Sample Robust Optimization," Bertsimas, Shtern, and Sturt investigate a simple approximation scheme, based on overlapping linear decision rules, for solving data-driven two-stage distributionally robust optimization problems with the type-infinity Wasserstein ambiguity set. Their main result establishes that this approximation scheme is asymptotically optimal for two-stage stochastic linear optimization problems; that is, under mild assumptions, the optimal cost and optimal first-stage decisions obtained by approximating the robust optimization problem converge to those of the underlying stochastic problem as the number of data points grows to infinity. These guarantees notably apply to two-stage stochastic problems that do not have relatively complete recourse, which arise frequently in applications. In this context, the authors show through numerical experiments that the approximation scheme is practically tractable and produces decisions that significantly outperform those obtained from state-of-the-art data-driven alternatives.