



Operations Research

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

In This Issue

To cite this article:

(2024) In This Issue. Operations Research 72(3):iii-vi. <https://doi.org/10.1287/opre.2024.inthisissue.v72.n3>

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 © 2024, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

In This Issue

<https://doi.org/10.1287/opre.2024.inthisissue.v72.n3>

Copyright: © 2024 INFORMS

Countercyclical Investment Patterns with Endogenously Determined Credit

Anticipated returns are sufficiently high to offset the elevated risks during downturns, presenting favorable investment opportunities. For instance, Warren Buffett accumulated more than \$10 billion in profits amidst the 2008 financial crisis, as highlighted in his *Wall Street Journal* interview from October 2013: “*In terms of simple profitability, an average investor could have performed just as well investing in the stock market if they bought during the panic period.*” In “Endogenous Credit, Business Cycle, and Portfolio Selection,” Choi, Koo, Lim, and Yoo reveal a divergence in investment behavior between affluent and less affluent individuals during economic downturns. Whereas wealthier individuals tend to increase their exposure to risky investments, less affluent individuals tend to reduce theirs. The authors employ rigorous modeling and solution analysis to establish endogenous borrowing constraints in general equilibrium, elucidating the observed investment patterns.

How to Price a Product Line When Customer Preferences Change over Time

Quality-differentiated products can help sellers increase their profits through market segmentation. However, in many business applications, such as online search and consumer lending, customer preferences evolve over time, making it difficult for sellers to use market segmentation. In “Selling Quality-Differentiated Products in a Markovian Market with Unknown Transition Probabilities,” Keskin and Li analyze dynamic pricing of a vertically differentiated product line when customer preferences for quality can shift over time. They show that data-driven learning is essential when operating in a changing market with unknown customer heterogeneity. The authors also develop a bounded learning policy that implements near-optimal data-driven learning in a Markov-modulated demand environment.

Navigating the Complex Web of Incentives in Asset Management: A Study on Investor, Partner, and Manager Dynamics

In “Dynamic Contracting in Asset Management Under the Investor-Partner-Manager Relationship,” Keppo, Touzi, and Zuo delve into the intricate world of incentives in asset management. The focus is on the complex interplay of actions and relationships among three key players: an investor, an investment company partner, and a fund manager. They uniquely explore the challenges of incentive contracts in an environment where the actions of these individuals are not fully observable to each other. It introduces a hierarchical contracting framework to understand the dynamics under the hidden actions. For instance, the research reveals how the optimal action of the manager is influenced by the cost of the partner’s actions and extends the analysis to scenarios involving multiple managers. This work sheds light on the often hidden intricacies of asset management, providing valuable

perspectives for the different players’ collaborative efforts for optimal outcomes in this industry.

Efficient Algorithm for Base Stock Policy Optimization in Single-Product Assemble-to-Order Systems with Stochastic Lead Times

Single-product assemble-to-order systems can be commonly observed in industrial settings. For example, firms targeting a niche market often choose to offer a single product; moreover, sole-product rollout is usually a preferable business strategy for manufacturers promoting a brand-new product or trying to penetrate a new market. However, maintaining an effective inventory management for those systems is a notoriously difficult problem, especially when the order lead times exhibit stochastic patterns. In “Single-Product Assemble-to-Order Systems with Exogenous Lead Times,” Muharremoglu, Yang, and Geng investigate the case with a special type of stochastic lead time and utilize a certain analysis technique to develop an efficient algorithm for the performance evaluation of base stock policies. In addition, efficiently computable upper and lower bounds are proved based on the key idea of the algorithm. Extensive numerical studies show that the heuristics and the approximation methods developed from those bounds perform well in base stock optimization.

Charactering a Cournot Equilibrium in an Asymmetric Multi-Market Setting

Charactering Cournot equilibria in settings with asymmetric firms is a challenging problem. In “Technical Note—Multimarket Cournot Equilibria with Heterogeneous Resource-Constrained Firms,” Caldenty and Haugh consider competition in a multimarket setting where firms are subject to varying budget or capacity constraints. They explicitly characterize the unique equilibrium in this environment by introducing the concepts of augmented and cutoff budgets for firms and markets. In particular, the authors show that a firm operates in a particular market if only if its augmented budget surpasses the market’s cutoff budget. They also investigate how the equilibrium properties change as the number of firms N varies while maintaining a fixed aggregate budget. The authors show via a numerical study that increasing N leads to a higher total output across all markets but that this monotonicity need not hold individual at the level of an individual market level. They also show that, although the firms’ cumulative payoff decreases in N , the consumer surplus and social surplus increase in N .

Risk Sensitivity and Firm Power: Price Competition with Mean²-Variance Profit Objective Under Multinomial Logit Demand

Firms with varying degrees of risk sensitivity perceive the same stochastic profit prospect differently. It is important to understand the equilibrium pricing behaviors under risk aversion and their implications on firms’ survival and market health. In “Technical Note—Risk Sensitivity and Firm Power: Price Competition with

Mean²-Variance Profit Objective Under Multinomial Logit Demand,” Li and Webster identify a power index, which is the ratio of product attractiveness to risk sensitivity, and show that the set of profitable firms is power index ordered. Firms with a favorable power index value earn a positive profit in a price competition, and others are driven to zero profit. Interestingly, although high risk aversion handicaps a firm, this paper shows that moderate risk aversion may give a firm an advantage over an otherwise equivalent competitor that is less risk sensitive. The authors also establish that in an equilibrium with market entry and exit, the power index is generalized to the ratio of product attractiveness and an entry cost-adjusted risk sensitivity measure.

The Power of Traffic Tolls

In “In Congestion Games, Taxes Achieve Optimal Approximation” Paccagnan and Gairing investigate the power and limitations of taxes as interventions in congestion games. Perhaps surprisingly, they find that efficiently computed taxes can achieve the same performance as the best polynomial time algorithm, even when the algorithm has full control over the agents’ actions. The authors establish three main results to support this claim. First, they prove that minimizing congestion is NP-hard to approximate within a given factor based on the latency functions. Second, they demonstrate how convex optimization tools can be used to design optimal taxes. Third, they provide a polynomial time algorithm with an approximation factor matching the hardness barrier. The upshot of their contribution can be summarized as follows: Judiciously designed taxes achieve optimal approximation, and no other tractable intervention can improve upon this result.

Incentivized Exploration in Reinforcement Learning

How do you incentivize self-interested agents to *explore* when they prefer to *exploit*? In “Exploration and Incentives in Reinforcement Learning,” Simchowicz and Slivkins consider complex exploration problems, where each agent faces the same (but unknown) Markov decision process (MDP). In contrast with traditional formulations of reinforcement learning (RL), agents control the choice of policies, whereas an algorithm can only issue recommendations. However, the algorithm controls the flow of information, and can incentivize the agents to explore via information asymmetry. The authors design an algorithm which explores all reachable states in the MDP. They achieve provable guarantees similar to those for incentivizing exploration in static, stateless exploration problems studied previously.

Robust Dynamic Assortment Optimization

Assortment optimization with choice model estimation and learning has been studied extensively in the data-driven revenue management literature. Existing methods and analysis, however, do not take into consideration the fact that some customers arriving at certain time periods might exhibit outlier purchasing behaviors. In “Robust Dynamic Assortment Optimization in the Presence of Outlier Customers,” Chen, Krishnamurthy, and Wang study dynamic assortment optimization in the presence of outlier customers modeled by an ϵ -contamination model. The impact of outlier customers on the revenue performance of an algorithm is analyzed and discussed.

Low-Dimensional Structure of Data Can Solve the Adversarial Robustness-Accuracy Conflict for Machine Learning Systems

Modern machine learning systems have demonstrated breakthrough performance in a multitude of applications. However,

they are known to be highly vulnerable to small perturbations to the input data, known as adversarial attacks. There are many well-documented examples of such behavior, for example small perturbations of an image, which is imperceptible to a human, can significantly degrade performance of modern classifiers. Adversarial training has been put forward as a way to improve robustness of learning algorithms to adversarial attacks. However, this benefit often comes at the cost of decreasing accuracy on natural unperturbed inputs, pointing to a potential conflict between adversarial robustness and standard accuracy. In “Adversarial Robustness for Latent Models: Revisiting the Robust-Standard Accuracies Tradeoff,” Javanmard and Mehrabi develop a theory to show that when the data enjoys low-dimensional structure, then it is possible to train models that are nearly optimal with respect to both, the standard and robust accuracies.

Optimal Dynamic Control of an Epidemic

In “Optimal Dynamic Control of an Epidemic,” Kruse and Strack analyze how to optimally engage in social distancing in order to minimize the spread of an infectious disease. They identify conditions under which any optimal policy first engages in increasingly more social distancing and subsequently decreases its intensity. The authors show that an optimal policy might substantially delay measures that decrease the transmission rate to create herd immunity and that engaging in social distancing suboptimally early can increase the number of fatalities. Finally, they find that optimal social distancing can be an effective measure to reduce the death rate of a disease.

New Algorithm Enables Efficient Decentralized Learning in Bipartite Queueing Systems

Bipartite queueing systems, where agents with individual job queues request service from a pool of heterogeneous servers, are standard models for service applications like data networks and call centers. Traditionally, a central controller schedules agent requests with full knowledge of system parameters. However, emerging applications require decentralized operation without this central coordination or complete system information. This presents challenges as agents lack the global knowledge needed to efficiently route jobs. Recent research into efficient decentralized learning algorithms for such systems faces limitations in non-optimal throughput, demanding computations, or degrading efficiency with exponential queue growth. In contrast, in “Efficient Decentralized Multi-agent Learning in Asymmetric Bipartite Queueing Systems,” Freund, Lykouris, and Weng introduce an algorithm that enables queues to efficiently learn decentralized scheduling policies while ensuring throughput optimality. The approach is computationally lightweight, achieving queue length bounds that scale polynomially rather than exponentially in system size. Experiments demonstrate faster convergence and robustness of our algorithm compared with prior decentralized algorithms.

Offline Evaluation of Dynamic Decision Making Policies under Unmeasured Confounding

In applications of offline reinforcement learning to observational data, such as in healthcare or education, a general concern is that observed actions might be affected by unobserved factors, inducing confounding and biasing estimates derived assuming a perfect Markov decision process (MDP) model. In “Proximal Reinforcement Learning: Efficient Off-Policy Evaluation in Partially Observed Markov Decision Processes,” Bennett and Kallus tackle this by considering off-policy evaluation in a partially observed MDP (POMDP). Specifically, they consider estimating the value of a given target policy in an unknown POMDP, given observations

of trajectories generated by a different and unknown policy, which may depend on the unobserved states. The authors consider both when the target policy value can be identified the observed data and, given identification, how best to estimate it. Both these problems are addressed by extending the framework of proximal causal inference to POMDP settings, using sequences of so-called bridge functions. This results in a novel framework for off-policy evaluation in POMDPs that they term proximal reinforcement learning, which they validate in various empirical settings.

Boosting Employment of Resettled Refugees

Whether a resettled refugee finds employment in the United States depends in no small part on which host community they are first welcomed to. Every week, resettlement agencies are assigned a group of refugees who they are required to place in communities around the country. In “Dynamic Placement in Refugee Resettlement,” Ahani, Gözl, Procaccia, Teytelboym, and Trapp develop an allocation system that recommends where to place an incoming refugee family with the aim of boosting the overall employment success. Should capacities in high-employment areas be used up as quickly as possible, or does it make sense to hold back for a perfect match? The simple algorithm, based on two-stage stochastic programming, achieves over 98% of the hindsight-optimal employment, compared with under 90% for the greedy-like approaches that were previously used in practice. Their algorithm is now part of the Annie™ MOORE optimization software used by a leading American refugee resettlement agency.

What Is the Optimal Way to Collect Data from Privacy-Concerned Users?

The data for many machine learning tasks are owned by individuals who are typically concerned about privacy. In “Optimal and Differentially Private Data Acquisition: Central and Local Mechanisms,” Fallah, Makhdoumi, Malekian, and Ozdaglar study the optimal design of a data acquisition mechanism aimed at learning the mean of a population. This data acquisition scheme includes the design of a payment rule to compensate users for their privacy loss. It also involves selecting an estimator that minimizes estimation error while simultaneously providing privacy guarantees to users in line with their privacy preferences. The authors formulate this problem as a Bayesian mechanism design problem and propose approximately optimal data acquisition mechanisms.

Bounded Regret for LP-Based Revenue-Management Problems

In “Technical Note—An Improved Analysis of LP-Based Control for Revenue Management,” Chen, Li, and Ye study a class of quantity-based network revenue-management problems. The authors consider a stochastic setting where all the orders are i.i.d. sampled and the customers are of finite type. They focus on the classic LP-based adaptive algorithm and consider regret as the performance measure. They found that when the underlying LP is non-degenerate, the algorithm achieves a problem-dependent regret upper bound that is independent of the horizon/number of time periods T ; when the underlying LP is degenerate, the algorithm achieves a tight regret upper bound that scales on the order of $T \log(T)$ and matches the lower bound up to a logarithmic order.

A Unified View of Online Matching and Resource Allocation Problems

Whereas small regret online algorithms for applications as diverse as network revenue management (NRM), assemble-to-order

(ATO) systems, and online stochastic bin packing (SBP) are known in the literature, the design of the existing algorithms are tailored to the specific application and often use the strategy of resolving a planning linear program. In “Technical Note—Greedy Algorithm for Multiway Matching with Bounded Regret,” Gupta proposes a unified model for studying such online matching/allocation problems. In the unified model, resources of three types—off-line (e.g., inventory in NRM), online-queueable (e.g., orders or resources in ATO systems), or online-nonqueueable (e.g., requests in NRM, items in SBP)—must be combined to create feasible configurations. Leveraging the unified framework, the author gives one simple greedy algorithm that gives small regret (bounded or logarithmic in time horizon) for these diverse applications under a mild nondegeneracy condition on the off-line planning problem.

Adaptivity in Stochastic Submodular Cover

Solutions to stochastic optimization problems are typically sequential decision processes that make decisions one by one, waiting for (and using) the feedback from each decision. Whereas such “adaptive” solutions achieve the best objective, they can be very time-consuming because of the need to wait for feedback after each decision. A natural question is are there solutions that only adapt (i.e., wait for feedback) a few times whereas still being competitive with the fully adaptive optimal solution? In “The Power of Adaptivity for Stochastic Submodular Cover,” Ghuge, Gupta, and Nagarajan resolve this question in the context of stochastic submodular cover, which is a fundamental stochastic covering problem. They provide algorithms that achieve a smooth trade-off between the number of adaptive “rounds” and the solution quality. The authors also demonstrate via experiments on real-world and synthetic data sets that, even for problems with more than 1,000 decisions, about six rounds of adaptivity suffice to obtain solutions nearly as good as fully adaptive solutions.

Wasserstein Distributionally Robust Optimization and Variation Regularization

In “Wasserstein Distributionally Robust Optimization and Variation Regularization,” Gao, Chen, and Kleywegt build a bridge between two areas in optimization and machine learning by establishing a general connection between Wasserstein distributional robustness and variation regularization. It helps to demystify the empirical success of Wasserstein distributionally robust optimization and devise new regularization schemes for machine learning.

Integer Factorization: Why Two-Item Joint Replenishment Is Hard

Joint replenishment problems constitute an important class of models in inventory management. They exhibit aspects of possible coordination among multiple products to save costs. Their computational complexity had been open even if there are just two products that need to be synced. In “Integer Factorization: Why Two-Item Joint Replenishment Is Hard,” Schulz and Telha present a simple framework based on integer factorization to establish the computational hardness of two variants of the joint replenishment problem with two items. Whereas difficult to solve in practice and not believed to be solvable in polynomial time, integer factorization is not as difficult as NP-complete problems. The authors show that a similar technique can be used to show even the NP-completeness of one variant of the joint replenishment problem (again with just two items).

Revenue Management of Service Systems under Incomplete Information

Revenue management with reusable resources finds many important applications in today's economy, such as cloud computing services, car/bicycle rental services, ride-hailing services, hotel management, project team management, and call center services. The existing literature predominantly assumes that the stochastic demand and service processes are given as an input to the models, and the pricing decisions are made with full knowledge of the distributional information. However, in practice, the decision maker may not know how demand or service rates react to price changes. Thus, the decision maker needs to learn the underlying mapping between prices and rates from past observations, while maximizing the total expected revenue on the fly. In "Online Learning and Pricing for Service Systems with Reusable Resources," Jia, Shi, and Shen develop a series of online learning algorithms for revenue management problems with reusable resources and showed that they admit an optimal regret bound.

A New Concept to Study Substitute Structures in Economics and Operations Models

In "S-Convexity and Gross Substitutability," Chen and Li propose a novel concept of S-convex functions defined on continuous spaces, which extends a key concept of M-natural-convex functions in discrete convex analysis. They develop a host of fundamental properties and characterizations of S-convex functions. In a parametric maximization model with a box constraint, they show that the set of optimal solutions is nonincreasing in the parameters if the objective function is S-concave and prove the necessity of S-concavity under some conditions. The monotonicity result finds notable inventory models. Interestingly, the authors show that S-concavity is the correct notion characterizing gross substitutability, an important concept in economics for markets with divisible goods.

Markov Decision Process Tayloring for Approximation Design

Optimal control problems are difficult to solve for problems on large state spaces, calling for the development of approximate solution methods. In "A Low-Rank Approximation for MDPs via Moment Coupling," Zhang and Gurvich introduce a novel framework to approximate Markov decision processes (MDPs) that stands on two pillars: (i) state aggregation, as the algorithmic infrastructure, and (ii) central-limit-theorem-type approximations, as the mathematical underpinning. The theoretical guarantees are grounded in the approximation of the Bellman equation by a

partial differential equation (PDE) where, in the spirit of the central limit theorem, the transition matrix of the controlled Markov chain is reduced to its local first and second moments. Instead of solving the PDE, the algorithm introduced in the paper constructs a "sister" (controlled) Markov chain whose two local transition moments are approximately identical with those of the focal chain. Because of this moment matching, the original chain and its sister are coupled through the PDE, facilitating optimality guarantees. Embedded into standard soft aggregation, moment matching provides a disciplined mechanism to tune the aggregation and disaggregation probabilities.

Balancing the Abandonment and Cancellation of a Ride-Matching Market

In "On-Demand Ride-Matching in a Spatial Model with Abandonment and Cancellation," Wang, H. Zhang, and J. Zhang propose a spatial model to approximate pickup times based on the number of waiting passengers and idle drivers. They analyze the dynamics of passengers and drivers in a queueing model in which the platform can control the matching process by setting a threshold on the expected pickup time. Applying fluid approximations, we obtain accurate performance evaluations and an elegant optimality condition based on which they propose a policy that adapts to time-varying demand.

Data-Driven Optimization with Distributionally Robust Second Order Stochastic Dominance Constraints

In "Data-Driven Optimization with Distributionally Robust Second Order Stochastic Dominance Constraints," Peng and Delage present the first comprehensive study of a data-driven formulation of the distributionally robust second order stochastic dominance constrained problem (DRSSDCP) that hinges on using a type-1 Wasserstein ambiguity set. It is, furthermore, for the first time shown to be axiomatically motivated in an environment with distribution ambiguity. They formulate the DRSSDCP as a multistage robust optimization problem and further propose a tractable conservative approximation that exploits finite adaptability and a scenario-based lower bounding problem. The authors then propose the first exact optimization algorithm for this DRSSDCP. They illustrate how the data-driven DRSSDCP can be applied in practice on resource-allocation problems with both synthetic and real data. Their empirical results show that, with a proper adjustment of the size of the Wasserstein ball, DRSSDCP can reach acceptable out-of-sample feasibility yet still generating strictly better performance than what is achieved by the reference strategy.