Operations Research

In This Issue

To cite this article:
(2026) In This Issue. Operations Research 74(5):v-ix. https://doi.org/10.1287/opre.2026.inthisissue.v74.n5

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

OPERATIONS RESEARCH
Vol. 74, No. 5, September–October 2026, pp. v–ix
ISSN 0030-364X (print), ISSN 1526-5463 (online)

In This Issue

https://doi.org/10.1287/opre.2026.inthisissue.v74.n5

Preserving Multimodular Structure in Substitution Models

Substitution is common in operations problems, where one product or resource may be used to satisfy demand intended for another. In models with multiple resources and demand classes, this flexibility makes it harder to establish structural properties of optimal dynamic decisions. In "Preservation of Multimodularity: New Results and Applications," Wang, Xiao, and Xu develop two new preservation results for multimodularity, a key property for deriving monotonicity and bounded sensitivity of optimal decisions. The first applies to one-way substitution, such as upgrading, and the second extends to two-way substitution under suitable cost and priority conditions. A central insight is that when allocation priorities are determined by marginal cost or value inequalities, both multimodularity and these inequalities are preserved after optimization. Applications include dynamic capacity management, remanufacturing inventory systems with lost sales, buyer-item matching on resale platforms, and multilocation transshipment.

Learning to Manage Multiproduct Inventory from Incomplete Demand Signals

Managing inventory across different products within limited warehouse space is a central challenge in retail and supply chain operations, especially when demand is unknown and lost-sales data are missing. This problem becomes even harder as product assortments grow larger. In "An Online Mirror Descent Learning Algorithm for Multiproduct Inventory Systems," Guo, Shi, Yang, and Zacharias develop the online mirror descent learning algorithm (OMELET), a scalable online learning algorithm that dynamically adjusts replenishment decisions using observed sales data. The method builds on mirror descent with cyclic updates to efficiently handle high-dimensional product menus. The authors show that the algorithm's regret grows only logarithmically with the number of products, a significant theoretical improvement over existing approaches. Extensive numerical experiments using empirical data confirm that OMELET not only outperforms existing state-of-the-art methods but also, translates into practical, actionable insights for managers making real-time inventory decisions.

Improved Inventory Management to Reduce Malaria Rates in the Congo

Humanitarian operations in some of the world's most remote regions often require ingenuity. Still, much can be applied from the extensive research on successfully managing operations in more common and often stable settings. Occasionally, this convergence of existing models with the unique considerations of these settings arrives at a void in the literature. In "Last-Mile Delivery of Malaria Prevention Products in the DRC: An Inventory Management Model Under Supply Chain Disruptions," Ata, Lehman, and Montgomery found themselves at such a point when they partnered with the Lake Tanganyika Floating Health Clinic to deliver malaria prevention products across 41 villages in the far East of the Congo. To allow continuous access to these products in a region that experiences seasonal flooding and frequent armed conflict, the authors introduce a model that incorporates random availability of high-capacity primary delivery methods with agile and costlier secondary delivery methods during disruptions. The optimal policy when the fixed cost varies between the disrupted and normal states is not previously addressed in Markovian inventory model results. The authors prove that an (s, S) policy remains optimal during normal operations even when facing these specific cost dynamics. They conduct numerical analysis demonstrating that the use of an (s, S) policy in normal operating periods and an order-up-to policy in the disrupted state is often optimal although they provide counterexamples of instances when it is not. Reducing costs in such settings allows for more potentially life-saving interventions to reach those in need.

Smarter Community Health Visits Improve Diabetes Outcomes

Diabetes takes a severe toll in developing countries, where high blood glucose contributes more than half of premature deaths. Community health workers (CHWs) offer a culturally tailored lifeline, but deploying them efficiently requires balancing the screening of new patients against managing those already in treatment. In "Planning a Community Approach to Diabetes Care in Low- and Middle-Income Countries Using Optimization," Adams, Boutilier, Deo, and Mintz introduce an innovative optimization framework that personalizes CHW visit plans to maximize community-wide glycemic control. Uniquely, the model explicitly factors in patients' motivational states—predicting their likelihood of enrolling in or dropping out of care—to guide intervention strategies and reduce attrition. Applied to operational data from urban slums in India, the approach delivers remarkable results; optimized visit plans reduced fasting blood glucose levels by up to 25% compared with the best baseline methods using identical capacity. The model also proved to be robust under imperfect information, offering a powerful, practical tool for global health resource allocation.

Searching with Advice You Cannot Fully Trust

Searching for a hidden target is a classic problem in operations research, arising in contexts such as search-and-rescue missions, exploration for natural resources, and robotic navigation. In many real-world scenarios, the searcher is not completely in the dark: it may have some prior information suggesting where the hider might be located. Such predictions can be helpful—but they may also be wrong. In "Search Games with Predictions," Angelopoulos, Lidbetter, and Panagiotou present the first study of search games in a fully randomized setting where both the Searcher and the Hider employ mixed strategies, and the prediction may be erroneous. The authors identify strategies that optimally balance two competing objectives: the search performance when the prediction is correct, and the performance when there are no assumptions on the quality of the prediction. The authors demonstrate this approach on several classic search problems such as box search with or without overlook, line search, and star search. The results quantify both the benefits and limitations of predictive information and provide a framework that does not apply only to search games, but to two-player zero-sum games more broadly.

Strategic Defense Against a Resource Gathering Attacker

In many real-world security settings, defenders must monitor a space using limited resources while attackers try to exploit gaps in coverage without being detected. The space may be a shop, museum, or airport, where the defenders use hidden cameras to detect thieves or smugglers, who wish to steal items or conceal contraband. To capture this interaction, Lidbetter and Lin present an attacker-defender game in "A Booby Trap Game." The defender places several hidden traps within a search space, while the attacker selects a subset of that space to exploit. The attacker's reward is proportional to the size of the selected subset if the subset does not contain any trap; otherwise, the attacker gets nothing. The authors present optimal (max-min/min-max), or near optimal, randomized strategies for both players for various spaces—including circles, line segments, two-connected networks, trees, and arbitrary Lebesgue measurable, path-connected subsets of Euclidean space.

Computing Equilibrium with Heterogeneous User Preferences

In "An Efficient Algorithm for Continuous Bi-Criteria Traffic Assignment," Xie, Wang, Li, and Nie address a long-standing challenge in transportation network modeling: how travelers choose routes in congested networks when balancing multiple route attributes, some of which are naturally continuous. Their new algorithm decomposes this complex problem into a sequence of intuitive route-level adjustments based on implicit boundaries of continuous attributes. This enables the model to capture detailed route-level decisions that have been difficult to compute at scale in such settings. The results show that the new method significantly outperforms existing approaches, often delivering speedups of an order of magnitude. Beyond computational gains, the study also establishes rigorous convergence guarantees and demonstrates the advantages of working directly with a continuous formulation as opposed to conventional discretization-based methods. The work opens the door to more realistic and scalable modeling of transportation systems with behaviorally heterogeneous users.

Robust Reinforcement Learning Without Compromising Data Efficiency

In "The Curious Price of Distributional Robustness in Reinforcement Learning with a Generative Model," Shi, Li, Wei, Chen, Geist, and Chi investigate model robustness in reinforcement learning (RL) to reduce the sim-to-real gap in practice. We adopt the framework of distributionally robust Markov decision processes (RMDPs), which aims to learn a policy that optimizes worst-case performance over a prescribed uncertainty set around a nominal Markov decision process (MDP). Despite recent efforts, the sample complexity of RMDPs has remained largely unresolved, leaving open whether distributional robustness has any statistical consequences when benchmarked against standard RL. Assuming access to a generative model of the nominal MDP, the paper provides a near-optimal characterization of the sample complexity of RMDPs across the full range of uncertainty levels under two common uncertainty sets specified by either total variation (TV) distance or chi-squared divergence. Somewhat surprisingly, the results reveal that RMDPs are not necessarily easier or harder to learn than standard MDPs. The statistical consequences of the robustness requirement depend heavily on the size and shape of the uncertainty set, requiring less data in the TV case but more in the chi-squared case compared with standard MDPs.

A Smarter Way to Plan for Stockouts

Online retail platforms increasingly face a basic operational problem: a product that appears available when a customer places an order may be out of stock when the order is fulfilled. In "Assortment Optimization with Replacement Options for Retail Platforms with Stockout Risk," Mitrofanov, Topaloglu, and Wang show how platforms can respond more effectively by optimizing not only the set of products shown to customers, but also the replacement options offered when stockouts occur. The authors study both non-adaptive and adaptive approaches, where the latter tailors replacement choices to the customer's initially selected item. They show that both problems are computationally challenging, but develop approximation algorithms with strong performance guarantees. Using Instacart data, they find that explicitly modeling replacement options improves expected revenue, and that adaptive replacement assortments provide additional gains, especially in categories with higher stockout risk.

Pricing When Buyers Learn from Each Other's Reviews

In online marketplaces, buyers increasingly rely on reviews from people such as themselves, filtering by attributes such as skin or body type, to estimate how much a product is worth to them. But when a buyer belongs to a less common group, there may be too few relevant reviews to inspire confidence, making that buyer reluctant to purchase at a high price. This creates a fundamental trade-off for sellers: higher prices increase revenue per sale, but they can also suppress early purchases and slow the accumulation of the reviews that future buyers need. In "Leveraging Reviews: Learning to Price with Buyer and Seller Uncertainty," Guo, Haghtalab, Kandasamy, and Vitercik study this two-sided learning problem, in which both the seller and the buyers face uncertainty. The authors present a no-regret pricing algorithm that strategically starts with lower prices to encourage purchases and generate informative reviews and then adapts over time to charge higher prices as those reviews accumulate. The authors also provides matching lower bounds, showing that the algorithm is essentially optimal up to lower order terms.

Learning to Route Orders with Risk in Mind

How should a trader split an order across multiple venues when both returns and risk matter? In "Risk-Aware Linear Bandits: Theory and Applications in Smart Order Routing," Ji, Xu, and Zhu develop a new bandit framework for this problem by combining mean variance optimization with linear bandit structure. Their model is motivated by smart order routing, where traders must learn from partial feedback while facing a very large action space. The authors propose two algorithms, RISE and RISE++, and show that both achieve near-optimal regret guarantees while avoiding the strong dependence on the number of actions that limits existing risk-aware bandit methods. They also validate the linear approximation empirically using the NASDAQ ITCH data set and demonstrate strong numerical performance in both synthetic and market-based experiments. The paper opens a new path for risk-sensitive online learning in financial decision making.

New Framework for Learning Against Long-Lived, Forward-Looking Agents

Repeated Stackelberg games are a canonical model for strategic principal-agent interactions. Learning in these games is well studied against myopic agents who greedily maximize their per-round payoff. However, complications arise with nonmyopic agents because they may strategically deviate from best responding to mislead the principal. In "Learning in Stackelberg Games with Non-myopic Agents," Haghtalab, Lykouris, Nietert, and Wei provide a general framework that reduces learning in the presence of nonmyopic agents to robust bandit optimization against myopic agents. This leads to a challenge of designing *minimally reactive* bandit algorithms, which balance the statistical efficiency of the principal's learning algorithm against its effectiveness at inducing near-best responses. The authors tackle this challenge across problem domains, including *security games*, *dynamic pricing*, and *strategic classification*. Along the way, they uncover a structural property for learning in security games, enabling them to improve the state-of-the-art query complexity with n targets from $O(n^3)$ to a near-optimal $\tilde{O}(n)$.

Maximizing Retail Revenue When Hiding Products Isn't an Option

Retailers often face a difficult balancing act: how to dynamically price a limited inventory of substitutable goods when business rules dictate that all available items must remain on display. Finding computationally tractable algorithms for pricing under this strict "show-all constraint" has remained an open problem. In "Pricing a Finite Inventory of Substitutable Products with Show-All Constraint," Goyal, Iyengar, and Udwani solve this open problem within the large inventory regime. The authors successfully derive an asymptotically optimal pricing policy for the multinomial logit (MNL) choice model, proving that optimal product prices must monotonically increase over time. Furthermore, they provide a generalized approximation framework applicable to any weakly rational and monotone choice model. These findings offer practical algorithmic strategies for maximizing revenue while maintaining comprehensive product visibility.

Learning to Compete under Unknown Demand

In many markets, firms must set prices while competing with others and learning demand from limited data, yet they rarely know how demand responds to prices or how competitors will react. Most existing models assume full knowledge of demand or access to competitors' information, which limits practical relevance. In "LEGO: Optimal Online Learning Under Sequential Price Competition," Li, Shi, and Mehrotra study a dynamic pricing problem in which multiple sellers repeatedly compete while observing only their own demand. The authors propose a decentralized learning algorithm that combines exploration with gradient-based price updates and does not require information sharing among competitors. They show that each seller's revenue loss grows at the minimal rate and that prices converge to the Nash equilibrium over time. Their results identify uncertainty in individual price sensitivity as the main challenge and provide a practical framework for data-driven pricing under competition.

Better Scheduling When the Cost of Delay Grows over Time

Imagine a stream of jobs, in which each job costs us some money (a holding cost) for every hour that it is not complete. Furthermore, the holding cost of each job can increase over time and possibly even jump up at deadlines associated with the job. To minimize the rate at which we hemorrhage money, it is common to deploy well-known strategies such as the celebrated generalized *c-mu* rule, which prioritizes scheduling jobs based on their instantaneous holding cost. However, this approach is limited: it ignores how costs rise in the future; for example, a job with a deadline might be ignored until after the deadline is missed. In "Improving Upon the Generalized $c\mu$ Rule: A Whittle Approach," Li, Gurushankar, Harchol-Balter, and Scheller-Wolf tackle this time-varying holding cost problem by translating it into a restless multiarmed bandit, which they use to derive a Whittle index. This leads to a new priority rule that accounts for how both holding costs and workloads evolve in the future. Their rule is simple to implement, yet it is remarkably robust, consistently outperforming all known heuristics across a wide range of settings.

Beyond Callbacks: A Simple and Easily Implementable Solution for Clinic Appointment Scheduling

Outpatient clinics face a difficult scheduling problem: Patient requests arrive unpredictably, whereas cancellations and no-shows make future demand uncertain. A longer scheduling window gives clinics more flexibility to absorb fluctuations in arrivals, but it also increases the risk of cancellations and no-shows. In "Appointment Scheduling of Outpatient Clinical Services Under Uncertain Patient Flows," Chen, Feng, Shanthikumar, and Yang develop a new policy, called the scheduling diagram, that explicitly captures this tradeoff in general settings. The scheduling diagram policy is a simple rule that specifies in advance the sequence of appointment slots to be filled within the booking window. It is easy to implement, allowing staff to assign appointments immediately as patients arrive. It is also easy to compute because it is constructed by ranking the marginal values of scalar functions and does not require repeated reoptimization. The authors further show that a hybrid approach, using a two-day schedule for a short window and an improved probabilistic policy a longer window, achieves strong and stable performance across a wide range of scenarios. The key message is that a simple interpretable scheduling diagram policy can be both practical and highly effective in managing uncertain patient flows.

When Privacy Concerns Undermine the Wisdom of Crowds

Organizations often rely on committees, internal experts, and employee reports to make high-stakes decisions. But experts may withhold or distort private information when truthful reporting could expose them to retaliation or other personal costs. In "Privacy, Voting, and the Wisdom of Crowds," Astashkina, Momot, and Salikhov study how privacy concerns affect collective decision making. The authors show that privacy concerns can break the usual wisdom-of-crowds logic: adding more experts may eventually reduce, rather than improve, information aggregation because each expert becomes less pivotal and more tempted to hide their signal. The authors characterize mechanisms that combine randomized response—garbling votes to provide plausible deniability—with accuracy-based rewards. These tools can restore truthful voting and, under appropriate conditions, allow organizations to recover near-perfect learning from expert groups.

Piecewise Affine Decision Rule Methods Enable the Smarter and More Efficient Decision Making for Contextual Stochastic Programming

In contextual stochastic programming, the efficacy of data-driven decision rule (DR) methods often faces the trade-off between approximation accuracy of DR hypothesis class and computational efficiency of learning the DR. To this end, in "Data-Driven Piecewise Affine Decision Rules for Stochastic Programming with Covariate Information," Zhang, Liu, and Zhao employ a piecewise affine decision rule (PADR) characterized by the difference of max-affine formula. Their research shows the consistency guarantees of the PADR method with the excess risk bound without the assumption of convexity or linearity. By exploiting the compact max-affine structure of PADR, they develop a novel majorization-minimization algorithm to efficiently solve the highly nonconvex and nonsmooth PADR-based ERM problem. Numerical experiments show that PADR significantly lowers costs, decreases computation time, and is robust to feature dimensions and nonlinearity of the underlying dependency. These results offers a theoretically consistent and practically scalable PADR approach for researchers and practitioners to solve contextual stochastic programming with both high reliability and computational efficiency.

Sparse PCA with Multiple Components

As the dimension of data sets increases, analysts often rely on principal component analysis (PCA) to summarize data via a small number of informative principal components (PCs). Sparse PCA makes those directions easier to interpret by using only selected variables, but computing several sparse components at once creates a problem: standard one-at-a-time methods lose the orthogonality that makes PCA useful in practice. In "Sparse PCA with Multiple Components," Cory-Wright and Pauphilet develop optimization-based methods that choose multiple sparse principal components simultaneously. Their approaches combine semidefinite relaxations, Lagrangian decompositions, and a new combinatorial upper bound to produce sparse, orthogonal components along with certificates of near optimality. Across real and synthetic data sets, the methods deliver high-quality solutions at practical scales, with average bound gaps around 3% for real-world instances with hundreds or thousands of features. The work gives practitioners a principled way to obtain interpretable low-dimensional representations without sacrificing the structural guarantees of PCA.

Curing Consumer Fatigue: A Data-Driven Approach to Dynamic Assortment

Have you ever grown tired of seeing the same lunch options at your corporate cafeteria every day? This phenomenon, known as customer satiation, poses a major challenge for businesses serving repeat customers, such as dining services and online flash-sale platforms. When customer utility depends on historical offerings, traditional assortment planning methods often fail, sacrificing either overall revenue or product variety. In "Assortment Optimization Under History-Dependent Effects," He, Zhang, and Zheng develop a new optimization framework to capture how past assortments influence future customer preferences. By reformulating this challenging nonlinear problem into tractable mixed-integer programs, they derive optimal assortment policies over time. Using real-world corporate cafeteria data, the study demonstrates revenue improvements of up to 10.4% over static assortment policies while maintaining product variety and reducing customer fatigue.

Unique Steady States Without Contraction in Nonlinear Stochastic Models

Many models in operations and economics describe systems whose future behavior depends not only on the current state but also on a function of the current distribution of states. In "Invariant Distributions in Nonlinear Markov Chains with Aggregators: Theory, Computation, and Applications," Light studies when such systems have a unique invariant distribution. The author develops flexible monotonicity-based conditions that can be tailored to different models to establish uniqueness. The author also shows that standard contraction arguments may fail in natural settings with strategic behavior, aggregate feedback, or interacting agents. The author provides existence results and a simple computational method that finds an invariant distribution by solving easier subproblems. The framework applies to strategic queues, inventory systems, nonlinear equations, and wealth distributions in dynamic economies, helping identify and compute unique steady states in stochastic systems where aggregate conditions shape individual dynamics.

Online Selection with Uncertain Disruption

Many digital and service platforms must decide whether to accept a current request or wait for a potentially better one. This trade-off becomes more challenging when serving a request may unexpectedly disrupt future operations. In "Online Selection with Uncertain Disruption," Xu, Kerimov, and Perez-Salazar introduce a model for such settings, where a decision maker observes values sequentially and each accepted value may trigger a disruption that stops the process. The authors develop simple threshold-based algorithms and evaluates them through competitive analysis against a clairvoyant benchmark that knows all values but still faces disruption uncertainty. The authors show that a nonadaptive single-threshold algorithm achieves the tight competitive ratio $1-1/e$, whereas adaptive threshold policies improve the asymptotic guarantee to approximately 0.745. The results connect disruption-aware online selection with classical prophet inequalities and quantify the value of adaptivity under operational uncertainty.

Off-the-Shelf Solvers Take on the Curse of Dimensionality

Many real-world decision problems—from managing power plants and microgrids to coordinating fleets of robots, recommender systems, and traffic networks—involve dynamics and rewards that decompose into smaller, loosely interacting components. Factored Markov decision processes (FMDPs) exploit this structure to escape the curse of dimensionality, but they have traditionally been solved with bespoke algorithms that require painstaking implementation and substantial fine-tuning. In "An MILP-Based Solution Scheme for Factored Markov Decision Processes," Liu, Wiesemann, and Yue introduce a new mathematical programming perspective on FMDPs that swaps hand-crafted machinery for off-the-shelf optimization solvers. By encoding sparse dependencies through mixed-integer linear programming–representable feature vectors, their formulation exploits factored structure in both states and actions, unifies and extends existing treatments of context-specific independence, and pairs an approximate linear program with a novel cutting-plane algorithm. Companion feature-learning and dynamic-basis routines further enhance scalability, and computational experiments across stylized and operational case studies illustrate the method's promise.

Fair Allocation with Perishable Resources

Food banks, blood banks, and vaccine programs routinely face a hard operational question: How can scarce supplies be distributed fairly and efficiently when tomorrow's need is uncertain and today's inventory may expire? In "Online Fair Allocation of Perishable Resources," Banerjee, Hssaine, and Sinclair study this problem in a setting where a decision maker must allocate a limited budget of perishable resources over time to stochastically arriving demand. The goal is to construct allocations that balance two competing objectives: fairness, where individuals receive similar allocations over time, and efficiency, where the available supply of goods is not wasted. The paper shows that perishability fundamentally alters the achievable fairness-efficiency Pareto frontier, relative to settings where resources do not expire. The authors first derive lower bounds characterizing the unavoidable perishing-induced loss incurred by any algorithm; this loss drives the modified frontier. They then design an algorithm that provably matches this Pareto frontier, up to polylogarithmic factors. The algorithm uses forecasts of future demand and perishing to adaptively choose between two guardrail quantities that control allocation decisions. Simulations calibrated to real-world data show strong performance and demonstrate that perishing-agnostic methods can perform poorly.

Choosing Product Assortments When Customers Buy More Than One Item

Retailers often need to decide which products to offer, but many existing assortment-optimization models assume that each customer buys at most one item. In "Multipurchase Assortment

Optimization Under a General Random Utility Model," Abdallah, Braverman, and Gu study a more realistic setting in which customers may purchase multiple items from an offered assortment while also considering outside options such as competing retailers, alternative channels, or not buying. The authors develop a general random utility model for this setting and show how the optimization problem can be approximated by a tractable surrogate problem. The surrogate is motivated by a large-offering asymptotic regime, can be solved efficiently, and performs well in numerical experiments, even for moderate-sized problems. The paper also provides theoretical bounds on the approximation error and develops an estimation framework that remains valid when outside-option choices are not directly observed.

Designing Fair and Monotone Lotteries for Proportional Systems

Proportional apportionment is the fundamental task of translating vote or population shares into an integer allocation of seats. The work monotone randomized apportionment studies randomized apportionment rules that are ex ante proportional; each party receives, in expectation, its prescribed quota, whereas the realized allocation remains integral and respects the target house size. However, the correlations between the seats awarded to different parties may exhibit bizarre nonmonotonicities. When parties or voters care about joint events, such as whether a coalition of parties reaches a majority, these nonmonotonicities can cause paradoxes, including incentives for strategic voting. In "Monotone Randomized Apportionment," Correa, Gölz, Schmidt-Kraepelin, Tucker-Foltz, and Verdugo identify Sampford's method, a classical unequal-probability sampling scheme, as a principled randomized rounding procedure that rules out such nonmonotonicity paradoxes. In doing so, this work connects apportionment, dependent randomized rounding, and sampling theory, providing a rigorous probabilistic framework for fair seat allocation.

Faster CoVaR Estimation via a Decoupled Approach

CoVaR is a widely used measure of systemic financial risk, capturing the risk of a portfolio conditional on another portfolio being under distress. In "Efficient Nested Estimation of CoVaR: A Decoupled Approach," Lin, Song, and Hong address two central challenges in CoVaR estimation: The conditioning event has zero probability, and portfolio losses often require costly simulation-based repricing. Their key insight is that, in nested simulation, the computational cost of outer simulations is negligible compared with that of inner simulations. The proposed approach decouples the nested estimation task by first learning approximations of portfolio loss functions using smoothing techniques and then evaluating the learned functions over a large outer-level sample to handle the zero-probability event. The framework is plug-and-play: Different smoothing techniques can be incorporated, and when their L^∞ approximation rates are available, the authors' theory directly yields the convergence rate of the resulting CoVaR estimator. The authors also establish smoothness of portfolio loss functions, a key foundation for efficient learning. Under suitable conditions, the estimator achieves the favorable near-square-root convergence rate $O_P(\Gamma^{-1/2})$, where Γ is the inner-level simulation budget.

How to Identify Problematic Value Judgements in Models of Efficiency and Decision Analysis?

Data envelopment analysis (DEA) is an optimization-based methodology used for the assessment of efficiency of organizations. Traditional DEA models are based on the vectors of inputs (resources) and outputs (products or services) of the observed organizations. DEA models are often enhanced by additional value judgements stated as tradeoffs between inputs and outputs or as dual weight restrictions. In "Proper and Consistent Production Tradeoffs in Models of Data Envelopment Analysis," Podinovski and Papaioannou showed that such additional value judgments may be contradictory (not proper) and develop analytical and computational tests for their identification. Any value judgments that are not proper indicate an error in their assessment. The developed approach is equally applicable in multicriteria decision analysis in which the tradeoffs between criteria are specified imprecisely by a set of linear inequalities stated in terms of criterion weights. Not proper tradeoffs result in the weights of some criteria (identifiable by the developed computational approach) being equal to zero, which effectively excludes such criteria from the analysis.

Certifying Submodularity with Algebraic Techniques

Submodularity is a structural property encoding the notion of diminishing returns, that is, the benefits one gets from an additional element decrease when many elements have been chosen. This property appears in many applications, from operations research to machine learning and economics. Unfortunately, testing whether a set function is submodular is computationally intractable for set functions of degree 4 or higher. In "Sum of Squares Submodularity," Deza and Hall introduce the notion of t-sum of squares submodularity, a hierarchy of algebraic certificates that provides tractable sufficient conditions for submodularity. For each fixed level t, membership in the hierarchy can be checked through semidefinite programming. The authors develop equivalent characterizations of the hierarchy, identifies operations that preserve it, and clarifies when it coincides with submodularity. They also demonstrate practical value in three settings: submodular regression, bounding submodularity ratios for approximate maximization, and constructing improved difference of submodular decompositions.

Erratum to "Online Allocation and Pricing: Constant Regret via Bellman Inequalities"

Theorem 3 of the paper "Online Allocation and Pricing: Constant Regret via Bellman Inequalities" [Vera A, Banerjee S, Gurvich I (2021) *Operations Research* 69(3):821–840] states a constant regret result for a menu-pricing problem. In this Erratum, Vera, Banerjee, and Gurvich correct the theorem's proof.