



Operations Research

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

In This Issue

To cite this article:

(2026) In This Issue. Operations Research 74(2):v-ix. <https://doi.org/10.1287/opre.2026.inthisissue.v74.n2>

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 © 2026, 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.2026.inthisissue.v74.n2>

Copyright: © 2026 INFORMS

When Market Incompleteness Is Preferable to Market Power

Since the liberalization of power markets in Europe, almost all investments have required some form of subsidy, due in part to the incomplete nature of the market. In that respect, the literature has consistently shown the benefits of risk-sharing instruments to boost investments. Because of entry barriers, however, some producers today might find themselves in a dominant position. In “When Market Incompleteness Is Preferable to Market Power: Insights from Power Markets,” Abada and Ehrenmann investigate the impact of completing a market with financial contracts in the presence of market power. To do so, the authors formulate a set of stochastic equilibrium investment models of risk-averse agents with possibly two types of market failure: incompleteness and market power. They provide existence results and conduct a thorough numerical application inspired by the French case. Contrary to what is reported in the literature, the authors highlight conditions under which social welfare is worse off when partially completing the market with contracts in the presence of a price-making incumbent.

Residential Battery Storage—Reshaping the Way We Do Electricity

Household investments in renewable energy technologies like rooftop solar and battery storage can be driven by motivations beyond mere financial returns. In “Residential Battery Storage—Reshaping the Way We Do Electricity,” Kaps and Netessine develop a structural estimation model that separates observed electricity demand from underlying consumption preferences, enabling the estimation of a nonfinancial utility that households derive from using more self-generated solar power rather than grid-procured electricity. The authors call this utility *nonmarket valuation*; provide evidence that it is driven by sustainability and autarky desires; and link it to the early adoption of residential storage. Applying their model to a novel data set of German households with solar and storage installations, they find that the median household has a nonmarket valuation of 0.29 euros per kilowatt-hour. Additionally, the authors demonstrate that residential storage can have unexpected effects. They show a “rebound effect,” whereby households with storage increase their overall electricity consumption by 4%. Counterintuitively, under Germany’s observed grid mix, residential storage on average increases carbon emissions. However, they also find that demand from the grid decreases by 38% for the average household equipped with solar and storage.

Boosting Profits of Trade-in Programs Through Smarter Pricing

As electronics trade-in programs expand worldwide, firms face complex pricing challenges: quoting offers instantly while managing vast product varieties and uncertain demand. In “Dynamic Pricing for a Multiproduct Consumer Electronics Trade-in Program,” Zhang, Lei, and Zhou address this challenge and present

effective solutions. The study develops pricing policies that determine both trade-in offers and resale prices over time to maximize profits. It first introduces a static pricing approach and proves that it performs nearly optimally; then, it further proposes a dynamic “batched-adjustment” policy that adapts to demand uncertainty and delivers an improved profit performance. Numerical experiments confirm the advantages of these methods. By demonstrating simple yet powerful strategies with provable effectiveness, the research offers firms actionable tools to increase profitability in trade-in operations.

Pricing and Resource Allocation Made Simple for Service Systems

Large-scale service systems—from drone delivery to cloud computing—face the dual challenge of balancing customer delays with revenue maximization. In “Near-Optimal Pricing and Resource Allocation in a Large-Scale Service System,” Wu, Liu, and Sun propose a dual-based pricing and resource allocation policy that is both simple and theoretically powerful. This greedy, one-step heuristic delivers performance guarantees matching the theoretical lower bound. Beyond the stylized model that illustrates its core idea, the study shows the value of dynamic pricing through an insensitivity result: any work-conserving rule can stabilize the system. The policy also proves robust under realistic conditions, including heterogeneous server pools and nonexponential service environments. Perhaps most striking, the authors uncover a “compensation effect”: near-optimal pricing curves need not rise monotonically with congestion. Instead, they may offset customers’ delay disutility to sustain revenue. These insights offer a practical, theory-backed framework for modern service operations.

Local Cost Synergies Matter in Procurement Auctions

Each winter, Minnesota’s Department of Transportation purchases nearly a quarter-million tons of road salt through competitive auctions. In “Local Cost Synergies in Reverse Auctions: An Application to Road Salt Procurement,” Gupta, Schmitt, and Stamatopoulos show that large firms benefit from “local cost synergies” when they win clusters of nearby depots. Because these synergies reduce costs, large firms are able to bid about 9% lower, on average, than they would otherwise. This cost advantage sharply reduces the probability that small firms will win contracts and leads simultaneous auctions to allocate depots less efficiently than possible. The authors conclude that any rigorous examination of procurement auctions must explicitly account for local cost synergies; ignoring them risks producing misleading results.

Post-trade Netting and Contagion

In “Post-trade Netting and Contagion,” Veraart and Zhang analyze how post-trade netting in over-the-counter derivatives markets affects systemic risk. It defines the notion of a post-trade netting (PTN) exercise and shows that two post-trade netting

services that rely on multilateral netting techniques, namely, portfolio rebalancing and portfolio compression, are special cases of a PTN exercise. It formally characterises the relationship between portfolio rebalancing and portfolio compression. The paper then analyses the effects of post-trade netting from a network perspective by considering contagion arising from defaults on variation margin payments. It provides sufficient conditions for post-trade netting to reduce systemic risk and shows that post-trade netting can be harmful. It also explores the implications of post-trade netting when institutions strategically react to liquidity stress by delaying their payments.

Beyond the First Queue: The Value of Information in Multistage Service Systems

In “Foresee the Next Line: Customer Strategies and Information Disclosure in Tandem Queues,” Snitkovsky, Roet-Green, and Ji examine how queue-length information affects customer behavior in multistage services, such as visiting the Apple Store—where you first wait to check in and then wait for a technician. It focuses on a two-stage tandem queue: an admission queue followed by a treatment queue. In the fully observable model, customers see both queues before entering the system, and the equilibrium joining strategy depends on the exact queue lengths rather than just their sum. In the partially observable model—which is more common in practice—customers initially see only the first queue and find the second queue’s length after completing the first stage. In this case, a threshold equilibrium emerges. Comparing models, the partially observable version generally yields higher throughput but lower social welfare. Results combine exact analysis in special regimes with extensive numerical computations.

Optimizing Police Presence to Prevent Crime Before It Happens

In “Proactive Policing: Resource Allocation for Crime Prevention with Deterrence Effect,” He, Li, and Zhao introduce a novel optimization framework that reimagines how police resources should be allocated to prevent crimes rather than react to them. Departing from traditional models that respond after incidents occur, the authors propose a proactive approach grounded in the economics of deterrence effect and rational choice theory. By modeling offenders’ location choices using a multinomial logit framework, the study captures how visible police presence not only suppresses crime locally but also redistributes potential criminal activity across space, a phenomenon known as crime displacement and diffusion of crime control. The authors establish the computational complexity of the problem, develop tractable mixed-integer conic reformulations, and extend the framework to dynamic settings. A data-driven case study in New York City demonstrates the model’s potential to guide smarter, prevention-focused deployment of urban policing resources.

Technical Note—Error Noted in “Robust Contract Designs: Linear Contracts and Moral Hazard” by Yu and Kong (2020)

One of the key claims in the paper “Robust Contract Designs: Linear Contracts and Moral Hazard” by Yu and Kong (2020) is proposition 4. This proposition states that when designing the best robust contract—assuming the agent’s utility function is piecewise linear and concave—the optimal solution involves only progressive fixed payments and linear rewards with progressive commission rates. In “Technical Note—Error Noted in “Robust Contract Designs: Linear Contracts and Moral Hazard” by Yu and Kong (2020),” Yang and Xin provide a clear counterexample backed by a detailed proof to show that this claim is actually incorrect. The authors further use numerical examples to demonstrate

that the contract design suggested by Yu and Kong can perform significantly worse than expected. Lastly, they identify and explain the specific mistakes in their mathematical proofs that led them to this incorrect conclusion.

Technical Note—Smarter Strategies for Learning What People Really Want

Understanding consumer preferences is crucial for designing better products, optimizing assortments, and personalizing recommendations. In “Technical Note—Active Learning for Non-parametric Choice Models,” Susan, Golrezaei, Emamjomeh-Zadeh, and Kempe present an active learning approach that efficiently uncovers the most informative patterns in the choices made by members of a heterogeneous population. Instead of relying solely on historical transaction data, their method strategically selects sets of products to offer and then uses the responses to construct a directed acyclic graph (DAG) representation of preferences. This DAG captures the top k choices, their probabilities, and how they relate to each other as k changes. Experiments on synthetic and real-world data show that the method learns preferences more accurately—and with fewer data points—than leading offline techniques. This work advances both the theory and practice of preference learning, with implications for retail, online platforms, and artificial intelligence agents that need to model human decision making.

A Novel Fully Online Matching Algorithm in Stochastic Environments

In many modern applications, such as ride-sharing, participants arrive and depart dynamically, requiring quick decisions about who to match and when. In “Fully Online Matching with General Stochastic Arrivals and Departures,” Li, Wang, and Yan introduce a fully online matching framework that models these real-world dynamics. Each arrival follows a known identical and independent distribution over agent types, whereas the sojourn time is unknown in advance and follows type-specific distributions with known expectations. To address this, they propose a linear programming-based algorithm that guarantees a competitive ratio of at least 0.192 under mild conditions—more than 50% better than the state-of-the-art result of 0.125. They also establish several hardness results to highlight the intrinsic difficulty of the problem. Numerical experiments further confirm the effectiveness and efficiency of their algorithm, offering new insights for decision making in stochastic and dynamic environments.

Demand in the Shadows: Proxy-Based Solutions for Smarter Pricing

Understanding how price affects customer demand is a cornerstone of data-driven pricing, but traditional approaches often struggle under endogeneity because of the presence of confounding factors. In “Proxy-Aided Demand Learning with an Application to Various Pricing Problems,” Shen and Cui tackle this challenge by leveraging ideas from proximal causal inference. They introduce a framework that incorporates proxy variables—categorized into treatment and outcome types—to enable reliable identification and estimation of customer demand. Central to their method is the use of a bridge function that allows accurate recovery of potential sales at different price points. Besides theoretical and managerial insights, the paper demonstrates practical applications in both static and contextual pricing, with the proposed algorithms applied to a real-world e-commerce data set. Tellingly, the proposed framework offers a promising new direction for practitioners aiming to optimize pricing with confounded data.

SOARing to Optimality: Smarter Matching Algorithms for On-Demand Platforms

In “Feature-Based Dynamic Matching,” Chen, Kanoria, Kumar, and Zhang study dynamic two-sided matching where both customers and service providers are characterized by high-dimensional feature vectors, motivated by platforms like on-demand home services. A key finding is that myopic greedy policies—which match each customer to the best available provider—can be highly suboptimal. The authors introduce SOAR (simulate-optimize-assign-repeat), a forward-looking algorithm that balances immediate match quality against preserving valuable supply for future arrivals. Through a novel analytical framework connecting dynamic matching to empirical optimal transport, they prove SOAR achieves near-optimal regret scaling under various distributional assumptions, establishing fundamental performance limits for feature-based matching problems and as a by-product resolve an open problem from prior work on dynamic spatial matching.

Fixing AI’s “Peer Review Lottery”

Getting a paper into a top artificial intelligence (AI) conference can feel like a lottery, with studies showing reviewer scores are often arbitrary. Now, in “You Are the Best Reviewer of Your Own Papers: The Isotonic Mechanism,” Su introduces a fix set to reform the field. The new “isotonic mechanism” tackles the crisis by asking authors to do the seemingly counter-intuitive: rank their *own* submissions from best to worst. The method’s effectiveness lies in its game-theoretic proof that honesty is actually the author’s best possible strategy. By harnessing this truthful self-assessment, the mechanism calibrates noisy and random reviewer scores, ensuring genuine scientific merit rises to the top. After successful large-scale experiments at major conferences, this mechanism isn’t just a theory. It’s being officially adopted by the International Conference on Machine Learning (ICML) in 2026, promising a fairer, more reliable future for millions of AI researchers.

New Algorithms Advance Revenue Management with Demand Learning

In “A Primal-Dual Approach Toward Resource-Constrained Revenue Management with Demand Learning and Large Action Space,” Miao, Wang, and Zhang introduce efficient algorithms for revenue management when firms face limited resources and uncertain demand. Revenue management, used in industries such as airlines, hotels, and retail, requires dynamic decisions on pricing and product assortments, whereas resources such as seats or inventory cannot be replenished. Traditional approaches often struggle with complexity or weak theoretical guarantees. The authors propose a primal-dual learning framework that combines optimization with machine learning’s upper confidence bound method. Their approach achieves near-optimal regret bounds, remaining computationally efficient even in large or infinite decision spaces. Applications include dynamic assortment selection, network revenue management with generalized linear demand, and joint pricing–assortment optimization. Numerical experiments show the methods consistently outperform benchmarks, offering practical, scalable solutions for data-driven industries.

Synthetic Interventions: Extending Synthetic Controls to Multiple Treatments

Traditional methods for policy evaluation typically focus on a single intervention. Yet many real-world settings feature multiple, often concurrent, interventions—making it crucial to understand their comparative effects. In “Synthetic Interventions: Extending Synthetic Controls to Multiple Treatments,” Agarwal, Shah, and

Shen introduce the synthetic interventions framework, a generalization of the synthetic controls method that accommodates multiple treatments within a unified model. By representing outcomes as a low-rank tensor capturing relationships across time, units, and interventions, their approach enables researchers to estimate counterfactual outcomes under all interventions for each unit. The authors establish consistency of their estimator and show that a bias-corrected version achieves asymptotic normality, permitting valid statistical inference. This work offers a new perspective for evaluating complex, multipolicy environments where traditional causal inference tools fall short.

Rethinking Causal Inference Through Robust Sensitivity Models

In “Sensitivity Analysis Under the f -Sensitivity Model: A Distributional Robustness Perspective,” Jin, Ren, and Zhou introduce a breakthrough in causal inference for observational data. Traditional analyses often fail when hidden confounders distort cause-and-effect relationships, but the newly proposed f -sensitivity model tackles this challenge by measuring the “average” impact of unobserved confounding instead of its worst-case effect. This framework connects causal inference to distributionally robust optimization, providing more realistic and interpretable bounds on treatment effects. With novel estimation and debiasing techniques, the method achieves statistical validity, even under minimal assumptions. The approach offers a flexible, computationally efficient way to test how robust conclusions remain in the presence of uncertainty, marking a significant advance in data-driven decision making across economics, healthcare, and policy analysis.

Layered Confidence for Contextual Pricing

Personalized pricing must learn from censored, binary demand signals while adapting to shifting contexts. In “Minimax Optimality in Contextual Dynamic Pricing with General Valuation Models,” Gong, You, and Zhang introduce an episode-based, parameter-free policy that couples explore-then-UCB with layered data partitioning. By discretizing the relevant noise range into a finite set of candidate prices and using offline regression oracles to estimate valuations, their method forms Azuma-style confidence bounds that avoid elliptical-potential lemma and do not require knowing the Lipschitz constant. The analysis yields minimax-optimal regret, and extends naturally to cases with known noise distributions, observable valuations, or additional smoothness. The authors also strengthen lower bounds and report experiments showing substantial gains over recent baselines. Together, these results provide a practical and theoretically sharp blueprint for contextual dynamic pricing across a broad class of valuation models.

Diagnosing Why Models Fail Under Distribution Shift and What to Do Next

Predictive models often perform worse when deployed in a new target setting, but it is rarely clear *why*. In “Diagnosing Model Performance Under Distribution Shift,” Cai, Namkoong, and Yadlowsky introduce a diagnostic, distribution shift decomposition (DISDE), that attributes the change in performance from the training to target distributions into terms for (i) an increase in harder but previously seen inputs from training, (ii) changes in how outcomes relate to inputs, and (iii) poor performance on new input regions absent from the training data. Applications to employment prediction demonstrate how this decomposition can inform potential modeling improvements, guiding whether to use domain adaptation techniques, adjust model covariates, or collect new samples. Additionally, DISDE is used to help explain

why certain domain adaptation methods fail to improve model performance for satellite image classification.

Post Reinforcement Learning Inference

In “Post Reinforcement Learning Inference,” Syrgkanis and Zhan develop a new inferential framework for data collected via reinforcement learning algorithms, the adaptive systems that update strategies as outcomes unfold. Traditional statistical methods fail in this setting because adaptivity induces time-varying variance and dependence across samples. The authors propose an *adaptively weighted generalized method of moments (AW-GMM)* estimator that stabilizes this variance through data-dependent weights. They prove that the weighted estimator achieves consistency and asymptotic normality, enabling valid hypothesis testing and confidence intervals for policy values and dynamic treatment effects. Their method provides a unified approach for structural estimation and inference under nonstationary, adaptively generated sequence data, with applications to dynamic off-policy evaluation and personalized decision systems.

An Algorithmic Approach to Managing Supply Chain Data Security: The Differentially Private Newsvendor

In “An Algorithmic Approach to Managing Supply Chain Data Security: The Differentially Private Newsvendor,” Chen and Chua examine an emerging challenge in data-driven operations: inventory decisions can unintentionally reveal sensitive demand information. Even when raw demand data are never shared, observers may reverse-engineer the data from small changes in order quantities. This risk is particularly acute in supply chains, where decisions are frequent, data-intensive, and easily observable. To mitigate this threat, the paper develops a suite of differentially private algorithms for the contextual newsvendor problem. These methods introduce carefully calibrated randomness into the newsvendor model, ensuring strong privacy protection under differential privacy while maintaining near-optimal operational performance. The authors also identify key drivers of the cost of privacy—data set size, contextual richness, and product variety—offering actionable guidance for firms seeking secure yet efficient data-driven operations. Finally, they demonstrate that privacy-preserving decisions can distort demand signals and reduce upstream supplier profits, highlighting important supply-chain-wide implications of data protection.

Smart Servers, Smarter Speed Scaling: A Decentralized Algorithm for Data Center Efficiency

A team of researchers from Georgia Tech and the University of Minnesota has introduced a cutting-edge algorithm designed to optimize energy use in large-scale data centers. As detailed in “Distributed Speed Scaling in Large-Scale Service Systems,” Rutten, Zubeldia, and Mukherjee developed a decentralized method allowing each server to adjust its processing speed autonomously without the need for communication or knowledge of system-wide traffic. The algorithm uses idle time as a local signal to guide processing speed, ensuring that all servers converge toward a globally optimal performance rate. This innovation addresses a critical issue in modern computing infrastructure: balancing energy efficiency with performance under uncertainty and scale. The authors demonstrate that their approach not only stabilizes the system but achieves asymptotic optimality as the number of servers increases. The work is poised to significantly reduce energy consumption in data centers, which are projected to account for up to 8% of U.S. electricity use by 2030.

Greater Price Stability Without Sacrificing Optimal Revenue Guarantees

In today’s digital marketplaces, sellers often rely on dynamic pricing—changing prices frequently—to optimize revenue. However, frequent price changes can undermine customer trust. In “Algorithmic Challenges in Ensuring Fairness at the Time of Decision,” Salem, Gupta, and Kamble demonstrate that it is possible to restrict price changes to only decrease over time—through strategic markdowns—while still achieving strong performance guarantees. Their research shows that optimizing pricing trajectories, even under monotonicity constraints, does not compromise optimal regret guarantees for maximizing revenue. This finding challenges the prevailing assumption that taming price experimentation may hinder success. The work opens new avenues for how online platforms can rethink pricing strategies to foster consumer loyalty without sacrificing profitability.

Smarter Load Balancing with Fewer Messages

Modern data centers rely on real-time information to route jobs efficiently—but constant communication between servers and load balancers can overwhelm the network. In “Load Balancing Using Sparse Communication,” Mendelson and Xu present a breakthrough: high-performance load balancing using drastically fewer messages. They introduce a flexible framework based on state approximation and develop new algorithms and communication protocols that maintain near-optimal performance even when communication is sparse. The key insight is that servers can monitor how wrong the load balancer’s estimate is and communicate only when necessary. Their approach reduces communication by over 90% while sacrificing little quality, as proven through both theory and simulation. These results offer a scalable and practical solution for large service systems where bandwidth is precious.

Joint Capacity Allocation and Job Assignment Under Uncertainty

In “Joint Capacity Allocation and Job Assignment Under Uncertainty,” Wang, Lim, and Loke propose a state-dependent “distributive decision rule” for simultaneous capacity allocation and job assignment decisions in a multiperiod stochastic resource allocation context with random supply replenishment, random demand, job waiting and abandonment. The decision rule can be reformulated into a convex optimization problem with polynomial number of constraints and decision variables. The framework can be applied in many service management settings such as ride-sharing fleet repositioning and patient management in healthcare. In simulations, the framework records 1%–15% improvements over alternative paradigms such as fluid approximations and approximate dynamic programming.

Finding Equilibrium in Chaos: A Breakthrough in Game Theory

When players in a game face messy decisions—like yes/no choices, rules layered within rules, or conflicting objectives—traditional algorithms often fail to find stable outcomes. In “The Cut-and-Play Algorithm: Computing Nash Equilibria via Outer Approximations,” Carvalho, Dragotto, Lodi, and Sankaranarayanan introduce a new algorithm, Cut-and-Play, that breaks this barrier. Unlike previous methods, Cut-and-Play handles nonconvex and unbounded decision spaces—the kind often found in real-world markets, public policy, and artificial intelligence systems. It works by iteratively solving simpler approximations of a complex game and then refining them with mathematical “cuts” until a solution is reached. Most strikingly, the algorithm finds equilibria up to 10× faster than existing techniques and is the first

of its kind to offer a general-purpose solution method for this class of problems. The work is a leap forward for both the theory and application of strategic decision making.

A Unified Framework for Analyzing and Optimizing a Class of Convex Fairness Measures

Fairness concerns arise naturally across a wide range of decision-making contexts and application domains. Addressing these concerns requires integrating fairness measures into optimization models; however, quantifying fairness, as well as formulating and solving fairness-promoting optimization problems, remain significant challenges. In “A Unified Framework for Analyzing and Optimizing a Class of Convex Fairness Measures,” Tsang and Shehadeh propose a new framework that unifies different fairness measures into a general, parameterized class of convex fairness measures. They introduce a unified framework for optimization problems with a convex fairness measure objective or constraint, including unified reformulations and solution methods. Additionally, they establish mechanisms for quantifying the impact of employing different convex fairness measures on the optimal solutions to the resulting fairness-promoting optimization problem. Numerical experiments, including applications to resource allocation and facility location, demonstrate the computational efficiency of the unified framework over traditional ones.

Dual-Directed Algorithm Design for Efficient Pure Exploration

Although experimental design often focuses on selecting the single best alternative from a finite set, many pure-exploration problems pursue richer goals. Given a specific goal, adaptive

experimentation aims to achieve it by strategically allocating sampling effort, with the underlying sample complexity characterized by a maximin optimization problem. In “Dual-Directed Algorithm Design for Efficient Pure Exploration,” Qin and You introduce a unified dual-directed framework for efficiently solving general pure-exploration problems, yielding a unified algorithm design principle that extends the top-two approach beyond best-arm identification. Their theoretical analysis proves asymptotic optimality for classical problems, such as Gaussian best-arm identification, thresholding bandits, and epsilon-best-arm identification. Extensive numerical experiments confirm these theoretical insights, showcasing significant improvements over existing methods. This dual-directed framework offers researchers and practitioners a powerful and versatile tool to navigate uncertainty and optimize exploration strategies effectively.

Some Bilinear Optimization Problems Are Surprisingly Easy to Solve

Bilinear terms in a continuous optimization problem are computationally challenging and are typically addressed with global optimization techniques or McCormick-based relaxations. In “Hidden Convexity in a Class of Optimization Problems with Bilinear Terms,” Gorissen, den Hertog, and Reusken discover a reformulation technique that provides an equivalent convex formulation that is solvable in polynomial time for a rich subclass of problems. The subclass is characterized by products of variables where one variable is nonnegative and the other variable interacts only with variables that are multiplied with the same nonnegative variable. Their finding provides a new avenue to efficiently solve inverse optimization problems, nonlinear robust optimization problems via their dual problem, and problems with variable coefficients.