



## Operations Research

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

### In This Issue

To cite this article:

(2020) In This Issue. Operations Research 68(5):iii-vi. <https://doi.org/10.1287/opre.2020.2047>

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](mailto: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 © 2020, 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.2020.2047>

Copyright: © 2020 INFORMS

### Monitoring in Dynamic Contract Design

Adverse events are harmful to a firm or to the society. In many occasions, better effort in safeguarding a system can reduce the chance of such events. Consider the scenario in which a company (i.e., “principal”) hires a subcontractor (i.e., “agent”) to fulfill the duty of a safeguard. The agent’s effort is unobservable, and the principal may, from time to time, conduct on-site monitoring. However, monitoring is often costly to the principal. In a dynamic environment in which the principal can use both monetary payments and monitoring to induce effort, how to optimally schedule them is a challenging and important problem. In “Optimal Monitoring Schedule in Dynamic Contracts,” Chen, Sun, and Xiao provide theoretical guidance on designing the optimal monitoring and payment schedules that always induce full effort from an agent. The authors formulate the contract design problem as a stochastic optimal control model and provide a complete characterization of the optimal solution. Their analysis suggests that the optimal dynamic contracts are simple to describe, easy to compute and implement, and intuitive to explain.

### Improving Ambulance Response Times in Developing Urban Centers

The lack of emergency medical transportation is viewed as the main barrier to the access and availability of emergency medical care in low- and middle-income countries (LMICs). Designing emergency response systems for urban centers in LMICs presents several unique challenges not encountered in high-income countries. For example, it is not the cultural norm and often not physically possible because of congestion for motorists to yield for emergency vehicles. In “Ambulance Emergency Response Optimization in Developing Countries,” Boutilier and Chan develop a prediction-optimization framework tailored for emergency response optimization in developing urban centers and provide an in-depth investigation into various policy questions. To do this, the authors traveled to Dhaka, Bangladesh, one of the largest cities in the world, to conduct field research and collect data. The data were initially used to develop machine learning models that estimate the spatial-temporal demand and travel times through the city—a first in a developing urban center. The predictions were then leveraged to create uncertainty sets as part of a robust-optimization model that provides a unified framework for emergency response optimization under travel time and demand uncertainty. Their policy experiments found that the performance of the current

system can be realized with one third of the current resources and that a fleet of small motorcycle-style ambulances can capture three times more demand and reduce response times by up to 35%.

### Data-Driven Ordering and Pricing

When a new product has just been introduced or the economy has just entered a new phase, a firm is often at a loss as to what the underlying demand pattern has become, let alone how best to respond to it. A natural idea is to manage ordering and pricing activities while simultaneously learning about the demand pattern. In “Dynamic Inventory and Price Controls Involving Unknown Demand on Discrete Nonperishable Items,” Katehakis, Yang, and Zhou formalize this idea. They design control policies that adapt with demand observations made over time. A policy’s regret is the lag of its performance behind that of an all-knowing one tailor-made for a specific demand pattern. For the proposed policies, the worst regrets over large ambiguity sets of unknown demand patterns are found to grow reasonably slowly over time.

### New Prescriptive Analytics Models to Implement and Benchmark Agile and Waterfall New Product Development (NPD) Approaches

When engaging in the development of new products, the primary objective of start-up companies is to generate a specified return level quickly and with high confidence. Achieving this goal is complicated because of uncertainties in projects’ returns and durations. In “Technical Note—Waterfall and Agile Product Development Approaches: Disjunctive Stochastic Programming Formulations,” Kettunen and Lejeune develop new disjunctive chance-constrained programming models that capture this goal. The first static model reflects the traditional waterfall product development process, whereas the second one is dynamic and depicts the agile product development process. Kettunen and Lejeune design a novel reformulation method and a decomposition algorithm and use them on a new product development problem encountered by a U.S.-based software start-up company. The results reveal that high confidence in reaching a certain return can be achieved by investing in projects with a longer development time and higher risk. Additionally, overlooking the capability to make dynamic decisions, as allowed by the agile approach, leads to overestimating the time needed to obtain the targeted return.

### Will Service Providers Prioritize Under Competition?

Research in service operations management generally considers prioritization beneficial for service providers (SPs), as they can better differentiate their customers. The healthcare, entertainment, restaurant, and airline industries use prioritization in some form for their service offerings. However, what happens when SPs compete? How will they provide service? In “Technical Note—Pricing and Prioritization in a Duopoly with Self-Selecting, Heterogeneous, Time-Sensitive Customers Under Low Utilization,” Sainathan answers these questions by analyzing a three-stage game between two SPs. In the first stage, they decide whether to prioritize; in the second stage, they set their prices; and in the third stage, customers self-select and purchase their best service options. The author identifies six possible equilibria: *high price*, *low price*, *regular*, *extreme*, *reduced class*, and *differentiated*. He also derives conditions for each of them to be the overall equilibrium. In particular, the author shows there are many situations in which at least one SP chooses not to prioritize.

### Optimizing Flight Schedules in a Network of Airports via a New Stochastic Integer Programming Algorithm

To mitigate flight delays, air traffic managers have access to two main levers: ground-holding programs (operational interventions at the tactical level) and demand management (scheduling interventions at the strategic level). In “A Stochastic Integer Programming Approach to Air Traffic Scheduling and Operations,” Wang and Jacquillat optimize these two sets of decisions in a network of airports, under weather-related uncertainty. The problem is formulated as a two-stage stochastic program with integer recourse. To solve it, the authors propose a decomposition algorithm based on new optimality cuts (leveraging the dual relaxation of the second-stage problem) and accelerate it by using neighborhood constraints (switching from exploration to exploitation in later iterations). Results suggest that limited scheduling interventions can lead to large delay reductions, and that significant benefits can be achieved through scale integration (by optimizing scheduling and operating decisions in the entire network) and scope integration (by optimizing scheduling and operating decisions together).

### Adaptive Matching Under Uncertainties

Two-sided online markets often propose matches based on imperfect knowledge of the characteristics of the two parties to be matched. Such uncertainty may result in inferior matches and may in turn incur negative externalities: A matched resource becomes unavailable to a more suitable party, at least for a while. For example, in labour platforms, if an expert is matched to a task which does not meet his or her expertise, then the other tasks which meet the expert's expertise may suffer. Similarly, in hospitality platforms, an economical accommodation becomes unavailable to a financially constrained customer if it is matched to a flexible customer. Thus, to optimize the long-term efficiency, a platform needs to optimize exploration-exploitation trade-

off while accounting for the resource usage. In “Adaptive Matching for Expert Systems with Uncertain Task Types,” Shah, Gulikers, Massoulié, and Vojnović address this challenge by developing a new model and optimal algorithm for a task-expert matching system, where a task is matched to an expert using not only the prior information about the task but also the feedback obtained from the past matches.

### Production Scheduling for Strategic Open Pit Mine Planning: A Mixed-Integer Programming Approach

Production scheduling is a large-scale optimization problem that must be solved on a yearly basis by every open pit mining project throughout the world. Surprisingly, however, this problem has only recently started to receive much attention from the operations research community. In “Production Scheduling for Strategic Open Pit Mine Planning: A Mixed-Integer Programming Approach,” Rivera, Espinoza, Goycoolea, Moreno, and Muñoz propose an integer programming methodology for tackling this problem that combines new classes of preprocessing schemes, cutting planes, heuristics, and branching mechanisms. This methodology is shown to compute near-optimal solutions on a number of real-world planning problems whose complexity is beyond the capabilities of preexisting approaches.

### Data-Based Dynamic Pricing and Inventory Control with Censored Demand and Limited Price Changes

Pricing and inventory replenishment are important operations decisions for firms such as retailers. To make these decisions effectively, a firm needs to know the demand distribution and its dependency on selling price, which is usually estimated using sales data at various testing price levels. Although more testing prices can lead to a better estimation of the demand–price relationship, frequent price changes are costly and come with adverse effect such as customers' negative perception. In “Technical Note—Data-Based Dynamic Pricing and Inventory Control with Censored Demand and Limited Price Changes,” by Chen, Chao, and Wang, data-driven algorithms are developed that learn the demand structure with constraints on the number of price changes. These algorithms are shown to converge to the optimal clairvoyant solution, and the convergence rates are the best possible in terms of profit loss.

### Robustness of Linear Contracts Under Moral Hazard

Linear contracts and their variants are quite popular in practice, for example, salesforce incentives and chief executive officer compensation. However, agency theory typically stipulates complex contract forms. In “Robust Contract Designs: Linear Contracts and Moral Hazard,” Yu and Kong provide an alternative explanation for the popularity of linear contracts: The robustness to model uncertainty renders the linear or generalized linear forms of the contracts under moral hazard. They adopt the worst-case decision criterion, and robust incentive compatibility

to ensure that the agent always behaves. The results are robust to general effort-contingent distributions and the risk-averse agent. These findings also shed light on how to design robust contracts when firms are facing model uncertainty or incomplete model information.

### A New Algorithm for Static Bin Packing

Static bin packing is the problem of partitioning a set of items (with scalar sizes) into identical bins of a given capacity, under the constraint that the total size of no partition exceeds the bin capacity. In the online variant, the list of items is revealed one at a time—an item must be assigned irrevocably to a partition at the time it is revealed. In “Interior-Point Based Online Stochastic Bin Packing,” Gupta and Radovanović look at the online variant under the assumption that the item sizes are independent samples from a common unknown distribution. The authors show that a greedy heuristic that assigns items so as to minimize a carefully chosen penalized Lagrangian of the offline math program is asymptotically optimal. Since the proposed heuristic is greedy in nature, it extends readily to nonstationary item size sequences for which the authors prove a novel family of performance bounds.

### Personalizing Treatment with Non-Stationary Multiarmed Bandits

In many sequential decision-making settings where there is uncertainty about the reward of each action, frequent selection of specific actions may reduce expected reward while choosing less frequently selected actions could lead to an increase. These effects are commonly observed in settings ranging from personalized healthcare interventions and targeted online advertising. In “Nonstationary Bandits with Habituation and Recovery Dynamics,” Mintz, Aswani, Kaminsky, Flowers, and Fukuoka address this problem. The authors propose a new class of models called ROGUE (reducing or gaining unknown efficacy) multiarmed bandits. They present a maximum likelihood approach to estimate the parameters of these models, and we show that these estimates can be used to construct upper confidence bound algorithms and epsilon-greedy algorithms for optimizing these models with strong theoretical guarantees. The authors conclude with a simulation study to show that these algorithms perform better than current nonstationary bandit algorithms in terms of both cumulative regret and average reward.

### Learning High-Dimensional Sparse Linear Models at Scale

In several scientific and industrial applications, it is desirable to build compact, interpretable learning models where the output depends on a small number of input features. Recent work has shown that such best-subset selection-type problems can be solved with modern mixed integer optimization solvers. Despite their promise, such solvers often come at a steep computational price when compared with open-source, efficient specialized solvers based on convex optimization and greedy heuristics. In

“Fast Best-Subset Selection: Coordinate Descent and Local Combinatorial Optimization Algorithms,” Hazimeh and Mazumder push the frontiers of computation for best-subset-type problems. Their algorithms deliver near-optimal solutions for problems with up to a million features—in times comparable with the fast convex solvers. The author’s work suggests that principled optimization methods play a key role in devising tools central to interpretable machine learning, which can help in gaining a deeper understanding of their statistical properties.

### Online Learning with Multiperiod Lookaheads

Sequential decision-making problems with high measurement noise usually involve nonconcave value of information, that is, the value of information is not concave in the amount of information collected. In such cases, existing single-period lookahead policies may be ineffective. In “Optimal Online Learning for Nonlinear Belief Models Using Discrete Priors,” Han and Powell propose multiperiod lookahead policies that are effective in measuring the value of information, and yet computationally tractable. They also establish consistency for the proposed policies, a rare result for online learning. Finally, they demonstrate the effectiveness of the approach in three applications: a health setting where one makes medical decisions to maximize healthcare response, a dynamic pricing setting where one makes pricing decisions to maximize the cumulative revenue, and a clinical pharmacology setting where one makes dosage controls to minimize the deviation between actual and target effects.

### Central Limit Theorems for Estimated Functions at Estimated Points

The need to estimate a function value from sample data at a point that is itself estimated from the same data set arises in many application settings. Such applications include value-at-risk, conditional value-at-risk, and other so-called distortion risk measures. In “Technical Note—Central Limit Theorems for Estimated Functions at Estimated Points,” Glynn, Fan, Fu, Hu, and Peng provide a simple proof for a central limit theorem for such estimators, and provide an accompanying batching/sectioning methodology that can be used to construct large-sample confidence intervals in the presence of such estimators.

### Decision Making When Things Are Only a Matter of Time

Suppose that “risk” does not concern “what” may happen but “when” something may happen. In “Decision Making When Things Are Only a Matter of Time,” Ebert analyzes time risk preferences: that is, preferences toward the risk of something happening sooner or later. The author defines and characterizes “prudence” and “temperance” for discount functions in analogy to the corresponding concepts for utility functions. Prudent discounting and temperate discounting matter for time risk preferences and thus, decisions in which an important event is only a matter of time. For example, the disutility from climate change—which is

only a matter of time—is underestimated when ignoring the uncertainty about when its undesired consequences come into effect.

### Time Inconsistency of Optimal Policies of Distributionally Robust Inventory Models

In “Technical Note—Time Inconsistency of Optimal Policies of Distributionally Robust Inventory Models,” Shapiro and Xin extend previous studies of time inconsistency to risk averse (distributionally robust) inventory models and show that time inconsistency is not unique to robust multistage decision making, but may happen for a large class of risk averse/distributionally robust settings. In particular, the authors demonstrate that if the respective risk measures are not strictly monotone, then there may exist infinitely many optimal policies which are not base-stock and not time consistent. This is in a sharp contrast with the risk neutral formulation of the inventory model where all optimal policies are base-stock and time consistent.

### Combinatorial Bandits: Learning Through the Optimality Cover Problem

When moving from the traditional to combinatorial multiarmed bandit setting, addressing the classical exploration versus exploitation trade-off is a challenging task. In “Learning in Combinatorial Optimization: What and How to Explore,” Modaresi, Sauré, and Vielma show that the combinatorial setting has salient features that distinguish it from the traditional bandit. In particular, combinatorial structure induces correlation between cost of different solutions, thus raising the questions of what parameters to

estimate and how to collect and combine information. The authors answer such questions by developing a novel optimization problem called the *lower-bound problem* (LBP). They establish a fundamental limit on asymptotic performance of any admissible policy and propose near-optimal LBP-based policies. Because LBP is likely intractable in practice, they propose policies that instead solve a proxy for LBP, which they call the *optimality cover problem* (OCP). They provide strong evidence of practical tractability of OCP and illustrate the markedly superior performance of OCP-based policies numerically.

### From Tree Ensemble Models to Decisions

Predictive models based on ensembles of trees, such as random forests and gradient boosted trees, are widely used in machine learning and data science. In many applications, the features that these models use are controllable and can be regarded as decision variables. This leads to a natural prescriptive analytics problem: How should these features be set, so as to maximize the value predicted by the tree ensemble model? In “Optimization of Tree Ensembles” Mišić proposes a mixed-integer optimization (MIO) model of this problem, proposes a hierarchy of approximations to this formulation based on truncating the trees at a particular depth, and develops two specialized constraint generation methods for solving the problem at scale. Using real data sets, including two detailed case studies in drug design and customized pricing, the author shows how this approach can efficiently solve large-scale problem instances to full or near optimality and outperforms solutions obtained by heuristic approaches.