Operations Research

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

In This Issue

To cite this article:
(2010) In This Issue. Operations Research 58(3):iii-vi. https://doi.org/10.1287/opre.1100.0840

OPERATIONS RESEARCH
Vol. 58, No. 3, May–June 2010, pp. iii–vi
ISSN 0030-364X | EISSN 1526-5463 | 10 | 5803 | 0iii

DOI 10.1287/opre.1100.0840

In This Issue

Do Forward Markets Reduce Market Power, or Can Companies Plan Their Capacities to Blunt Competition?

The impact of contracts is a recurrent theme in industries where firms have market power. Economists working on the restructuring of the electricity market have emphasized the potential of contracts for mitigating market power. More generally, other economists have insisted on their role for foreclosing the market. Finally, another stream of the literature values the role of contracts for incentivizing investments by reducing demand risk. In "On the Impact of Forward Markets on Investments in Oligopolistic Markets with Reference to Electricity," F. Murphy and Y. Smeers show that the situation is far from clear and that the impact of contracts is hard to predict. They consider two firms with market power that operate different technologies and need to invest. They show that contracts have no effect whatsoever when demand is certain: investments and the exercise of market power remain as without contracts. In contrast, the impact of contracts becomes completely unpredictable when demand is uncertain. Contracts can increase or decrease the incentive to invest and reinforce or mitigate market power.

Modeling the Economic Effects of Different Allocation Schemes for Pollutant Cap-and-Trade Systems

The U.S. Congress is presently considering implementing a tradable allowance scheme to limit carbon dioxide emissions that arise from burning fossil fuels. Such schemes work by first distributing a fixed number of allowances among sources of emissions. The number of permits is strictly limited so that a target reduction in total emissions is achieved. Then those sources are allowed to trade, so that sources with high control costs would buy allowances from those with low costs. A market price would then emerge for pollution. Such systems are already in place in the European Union and several U.S. states. The market value of the allowances under some Congressional bills would be over $100 billion per year, according to the Congressional Budget Office. Hence, how Congress chooses to initially distribute the allowances will create large "winners" and "losers." Furthermore, if the distribution of allowances depends on present or recent output, investment, or other decisions, companies may be motivated to change those decisions in a way that increases the net cost to society of meeting the carbon reduction target. In "Long-Run Equilibrium Modeling of Emissions Allowance Allocation Systems in Electric Power Markets," J. Zhao, B. F. Hobbs, and J.-S. Pang develop a model for analyzing the inefficiencies in electricity generation that could result from alternative allowance distribution systems. Their nonlinear complementarity model considers long-run equilibrium investments in electricity generation technologies that would result under alternative distribution rules for power markets characterized by time varying demand and multiple technology types. Existence of equilibria is shown under mild conditions. Example solutions show that allocating allowances to new capacity based on the amount and/or type of fuel consumed (as implemented in the phase I of the European system) can have several negative economic impacts. These include possibly large distortions in capacity investment relative to the least-cost means of meeting the given carbon target, as well as inversions in the operating order of power plants and inflation of consumer costs. The distortions can be smaller for tighter carbon restrictions and are somewhat mitigated if there is also a market for electricity capacity or if there exist minimum-run restrictions on coal plants. Distortions are also minimized if allowances are allocated to plants in proportion to energy sales rather than installed capacity.

Limit Order Books as Queuing Systems

With the advent of electronic trading platforms, many stock markets have been shifting to an order-driven trading system where incoming buy and sell orders are aggregated in a limit order book and executed at the best available prices. Limit-order books are in many ways analogous to queuing systems, and insights from queuing can be helpful in modeling order book dynamics. Building on these ideas, R. Cont, S. Stoikov, and R. Talreja propose in "A Stochastic Model for Order Book Dynamics," a continuous-time model for the dynamics of a limit order book. The paper exploits queuing analogies to analytically compute quantities of interest in high-frequency trading and shows how the model can be estimated and applied in a practical context using high frequency data from the Tokyo Stock Exchange.

Real Option Valuation of Natural Gas Storage

The real option valuation of contracts on the capacity of natural gas storage facilities is an important practical problem. Due to its difficulty, practitioners typically use heuristics

for this purpose, whose effectiveness is currently unknown. In "An Approximate Dynamic Programming Approach to Benchmark Practice-Based Heuristics for Natural Gas Storage Valuation," G. Lai, F. Margot, and N. Secomandi develop a novel upper bound to benchmark common heuristics used in practice and a novel heuristic, both based on approximate dynamic programming. Their computational results, based on real data and realistic instances, show that with periodic reoptimization all these heuristics perform very well relative to their proposed upper bound. Besides natural gas, their methods and results have relevance for the valuation of the real option to store other commodities.

A Robust Approach to Inventory Management

Effective management of inventory in the face of demand uncertainty is an important but challenging problem. As such, research on inventory management usually comes in the more descriptive flavor where optimal policies are carefully studied under specific assumptions. In the real world, it is impossible to elicit the exact distributions of the uncertain demands. In "Robust Approximation to Multiperiod Inventory Management," C.-T. See and M. Sim introduce a new approach for tackling the inventory management problem. The approach allows the modeler to consider uncertain demands across time periods whose distributions are only prescribed by their descriptive statistics such as means, variances, and directional deviations, among others. Unlike earlier works, the solution to the inventory control problem can be obtained by solving a mathematical programming problem in the form of tractable second order cone program. They also demonstrate the effectiveness of the approach through extensive computational studies.

Accounting for Distribution and Moment Uncertainty in a Stochastic Optimization Problem

Although one never knows what the future holds, stochastic optimization models can provide key insights about the structure of good decisions. However, when, as is often the case in practice, a decision maker's information about the stochastic nature of the problem is limited, he may resist applying the decisions proposed by the model. This is, for example, the case in portfolio optimization problems, where one cannot simply assume to know the model that drives the evolution of financial assets' values but instead must rely on imperfect information about this process, e.g., historical data. In "Distributionally Robust Optimization Under Moment Uncertainty with Application to Data-Driven Problems," E. Delage and Y. Ye present a distributional set that has statistical guarantees of containing the distribution from which the observed historical samples were drawn. Moreover, they show that finding the min-max (or robust) decision with respect to this set is tractable for a large class of problems and, in particular, leads in practice to better performing portfolios on the "true" distribution driving the evolution of financial assets valuation.

Utility-Maximizing Resource Control

The utility-maximizing control refers to the real-time allocation of each server's capacity among the job classes in a stochastic network. It is motivated by certain resource allocation protocols in Internet, e.g., the TCP, and also by managing resources in assemble-to-order systems and airline revenue management problems. It has been a challenge to characterize any performance measures (e.g., delay, congestion, etc.) of the network under such resource controls, let alone optimize the performance. In their paper, "Utility-Maximizing Resource Control: Diffusion Limit and Asymptotic Optimality for a Two-Bottleneck Model," H.-Q. Ye and D. D. Yao study a certain instance of the network model and establish its diffusion limit. They have also identified a cost objective associated with the utility function and show that it is minimized at the diffusion limit by the utility-maximizing allocation within a broad class of "fair" allocation schemes.

Fairness Towards Employees in Large-Scale Service Organizations

Justice towards employees is an aspect that has been mostly ignored in the operations management literature. However, it has been shown that when employees do not feel there is a sense of justice, or fairness, they are more likely to leave, leading to higher employee turnover, and, therefore, higher operational costs. In "Fair Dynamic Routing in Large-Scale Heterogeneous-Server Systems," M. Armony and A. R. Ward focus on fairness towards servers (employees) in large-scale service systems and its trade-off with customer satisfaction. Specifically, they simultaneously minimize customer waiting time while guaranteeing that system idleness is equally distributed among the servers.

Investment Policy with an Option to Exit

Firms often invest in projects with a declining profit stream. For example, even in the face of deteriorating and highly volatile demand, firms often invest in, rather than discard, aging technologies. Reinvestment in a deteriorating project requires special attention because there is an embedded option to abandon the project sometime after the reinvestment, and the impact of uncertainty and the embedded option has not been fully studied. In particular, conventional real options models have shown that the threshold of investment increases in the volatility of the underlying asset value, but it is not clear whether this result always holds when there is an embedded option. To address this question in "Invest or Exit? Optimal Decisions in the Face of a Declining Profit Stream," H. D. Kwon uses the methods of stochastic analysis to study an analytically tractable model of investment and exit decisions of a firm with a declining profit stream. When the profit boost upon investment is sufficiently large, this paper finds a novel result: the investment

threshold decreases in volatility. This is in contrast to the results of conventional real options models without embedded options.

An Effective Method for Comparing Approximate Solution Sets for Multiple Criteria Combinatorial Optimization Problems

Multiple objective a posteriori heuristics, which provide a set of nondominated solutions for the decision maker rather than a single final solution, have received significant attention in the recent years due to their potential to solve real-world multiple objective combinatorial optimization (MOCO) problems. However, evaluating the quality of approximate solutions obtained through the use of these heuristics is generally difficult since sets of alternate solutions should be evaluated and compared, which necessitates the use of an appropriate comparison measure. Previous research on the integrated preference functional (IPF), which assigns an expected utility-based scalar value to a given discrete set of nondominated points, is shown to provide a robust quantitative measure for comparing different solution sets for MOCO problems. In "Quantitative Comparison of Approximate Solution Sets for Multicriteria Optimization Problems with Weighted Tchebycheff Preference Function," B. Bozkurt, J. W. Fowler, E. S. Gel, B. Kim, M. Köksalan, and J. Wallenius extend previous work on the IPF, so that the weighted Tchebycheff function can be used as the underlying implicit value function. This extension is useful because modeling the decision maker's value function with the weighted Tchebycheff function has the property of capturing all nondominated points, rather than being restricted to supported nondominated points. Accounting for the impact of unsupported nondominated points is especially important as it has been shown for several MOCO problems that the majority of nondominated points are unsupported. The authors present an exact calculation method for the IPF measure for an arbitrary number of objectives and make use of two- and three-objective numerical examples to illustrate the desirable properties of the IPF with the weighted Tchebycheff function.

Cost Allocation in a Supply Chain

The cost allocation problem is important in defining profitability of various entities in the supply chain and in the way each of them assesses its value. Moreover, the cost allocation scheme may have a significant impact on long-term strategic decisions, and therefore must be fair in the sense that no entity would feel that it can do better by breaking away and forming other cooperations. In "Cooperation in Service Systems," S. Anily and M. Haviv consider the question of how service providers should split among themselves the system-wide gain due to pooling their resources. By looking at subsets of retailers as coalitions and using notions from game theory, the authors show that the grand coalition is the optimal formation of any bargaining process. They define the cost allocation question as a cooperative transferable utility game and investigate its core. They show that even though the game is neither additive nor concave, its core is not empty. Special attention is given to nonnegative core cost allocations. These allocation schemes guarantee that no subset of servers will ever quit the grand coalition, ensuring that the system is stable.

Optimal Solution to a Problem in IMRT Treatment Planning

Intensity modulated radiation therapy (IMRT) is an advanced form of radiation therapy for combating cancer. In IMRT, a set of beams of different shapes is delivered from multiple angles in order to deliver a desired dose to clinical targets while sparing the surrounding functional and healthy tissue. The equipment that dynamically forms the apertures that shape these beams is called a multileaf collimator (MLC). The total amount of IMRT treatment time is a function of the number of aperture setups that the MLC must perform and the amount of radiation delivered using each aperture. Minimizing this function is a very difficult combinatorial optimization problem that is akin to a minimum-cardinality matrix decomposition problem. In "Optimal Multileaf Collimator Leaf Sequencing in IMRT Treatment Planning," Z. C. Taşkın, J. C. Smith, H. E. Romeijn, and J. F. Dempsey provide a hybrid integer programming/constraint programming decomposition algorithm that is the first approach capable of solving this problem within clinically reasonable computational resources. The efficiency of the algorithm is demonstrated on a set of 25 clinical instances.

Using Real-Time Information to Manage Assemble-to-Order Systems

Enabled by advanced information technology, assemble-to-order (ATO) manufacturing systems are becoming more prevalent. In such systems, components and subassemblies are kept in stock, but final products are assembled only after customer orders are confirmed. Most of the literature on continuous-time ATO systems assumes that component stocks are allocated to product demands on the FCFS basis. While this approach is used in many practices, it is known not to be optimal. Some other alternatives are also used in practice but are much harder to analyze. In "No-Holdback Allocation Rules for Continuous-Time Assemble-to-Order Systems," Y. Lu, J.-S. Song, and Y. Zhao analyze one of these alternatives, which uses real-time stock availability information for stock allocation. They conduct a comparison of this alternative with the FCFS rule. More importantly, their analytical approach presents the opportunity to analyze other non-FCFS allocation rules that use real-time information to advance our understanding of effective operating policies for ATO systems.

Gradient Estimation for Complex Maintenance Systems

Individuals and companies depend increasingly on complex systems for their daily functioning. The potential costs of breakdowns of these systems are quite high, so the maintenance of systems is crucial in order to increase their availability. Therefore, during the last decades, there has been a growing interest in the modeling and optimization of maintaining systems consisting of multiple components. In "Gradient Estimation for Multicomponent Maintenance Systems with Age-Replacement Policy," B. Heidergott and T. Farenhorst-Yuan consider a multicomponent maintenance system with an F-failure group age-replacement policy: it keeps failed components idling until F components are failed, and then replaces all failed components together with the nonfailed components whose age has passed the critical threshold age Θ_n for components of type n. With each maintenance action, costs are associated. Various unbiased estimators are established for the gradient of the average cost with respect to the age thresholds Θ_n. The performance of the gradient estimators when applied in a stochastic optimization procedure for finding the optimal vector of age threshold is compared with that of gradient-free methods.

Deliberate Stockouts to Stimulate Consumer Interest

While stockouts have traditionally been viewed as evidence of inventory mismanagement, some manufacturers recently have used them as an effective tool for luring consumers who would not make a purchase otherwise. This practice, which is implemented by manufacturers of fashion and luxury products, combines innovative marketing with deliberate understocking to create a sense of exclusivity of a product. In "How Much Demand Should Be Fulfilled?" A. Sapra, V.-A. Truong, and R. Q. Zhang investigate how a company should manage its inventory when demand can be manipulated by restricting supply. They model the problem as a finite horizon, periodic inventory system where demand in each period is a decreasing function of the net ending inventory in the previous period. They show that understocking is indeed optimal in various scenarios. They then propose a novel strategy, referred to as the inventory-withholding strategy, to further explore the benefits of understocking by withholding some inventory from consumers.

Inventory Management with Two Sources of Supply

Today's supply chains span multiple continents, for example, a U.S. retailer could procure a product from a supplier in Asia. Such sourcing arrangements arise primarily from lower manufacturing costs in Asia. On the other hand, the long lead times in such a supply chain result in increased spending in inventory for the same level of customer service (availability of inventory to customers). One strategy that allows retailers to benefit from the lower manufacturing costs in Asia without spending excessively on inventory is to have a secondary or emergency supplier "on-shore," who offers shorter lead times in return for higher unit costs. In "New Policies for the Stochastic Inventory Control Problem with Two Supply Sources," A. Sheopuri, G. Janakiraman, and S. Seshadri provide algorithms for such a firm, which uses two suppliers with different costs and lead times, to manage its inventory.

Rescheduling

Often in manufacturing practice, a schedule and corresponding resource allocations over time to support it have been planned, and then a disruption occurs. Disruptions occur for various reasons, and each has a different effect on the schedule. It is important not only to reschedule in order to consider the new information, but also to minimize the costly disruptive effect on the rescheduling resources. This trade-off is key in making effective rescheduling decisions. In "Rescheduling for Job Unavailability," N. G. Hall and C. N. Potts develop a variety of optimal and approximate procedures for rescheduling when the delay to the arrival of components or raw materials prevents some jobs from starting at their planned times. These procedures assume a constraint on the disruption caused to any job in the original schedule. The authors' heuristics with guaranteed performance bounds and their approximation schemes are among the first to be developed for rescheduling problems.

Approximating the Stationary Distribution of Complex Markovian Systems

Nowadays, people calling a hotline are routed to call centers where they usually have to wait until being served. Callers are typically impatient and will hang up if they have to wait for too long, and call again later. Including impatience and retrial of customers into a queueing model leads to the class of so-called retrial queues. Unfortunately, so far it was not possible to derive the stationary distribution for this type of queues in closed form, hindering the analysis of their long-run behavior. In "Series Expansions for Continuous-Time Markov Processes," B. Heidergott, A. Hordijk, and N. Leder provide update formulas that allow for the approximation of the stationary distribution of a Markov process by means of another, simpler, Markov process. They apply their approach to the performance analysis of a retrial queue modeling a call center; the retrial queue is approximated by the corresponding model without retrials. They complement their theory by presenting conditions that assure convergence of their approximation at a geometric rate. Since their approach does not require uniformizability of the respective processes, it is applicable to a broader range of systems than were comparable theories presented previously.