MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Vol. 7, No. 1, Winter 2005, pp. 81–99
ISSN 1523-4614 | EISSN 1526-5498 | 05 | 0701 | 0081
DOI 10.1287/msom.1050.0075

The MSOM Society Student Paper Competition: Extended Abstracts of 2004 Winners

As is our tradition at the journal, we are pleased to publish the extended abstracts from the winners of the 2004 MSOM Student Paper Competition. We do this to celebrate the achievements of these young scholars and provide you with the opportunity to learn about their work in more detail.

The 2004 prize committee was chaired by Professor Phil Kaminsky (University of California, Berkeley). The other committee members were: Naren Agrawal (University of Santa Clara), Hyun-Soo Ahn (University of Michigan), Damian Beil (University of Michigan), Fernando Bernstein (Duke University), Izak Duenyas (University of Michigan), Wedad Elmaghraby (Georgia Institute of Technology), Jeremie Gallien (Massachusetts Institute of Technology), Teck Ho (University of Pennsylvania), Seyed Iravani (Northwestern University), Ananth Iyer (Purdue University), Eric Johnson (Dartmouth College), Roman Kapuscinski (University of Michigan), Pinar Keskinocak (Georgia Institute of Technology), Anton Kleywegt (Georgia Institute of Technology), Özalp Özer (Stanford University), Georgia Perakis (Massachusetts Institute of Technology), Alan Scheller-Wolf (Carnegie Mellon University), Sridhar Seshadri (New York University), Max Shen (University of California, Berkeley), David Simchi-Levi (Massachusetts Institute of Technology), Jay Swaminathan (University of North Carolina), Terry Taylor (Columbia University), Beril Toktay (INSEAD), Scott Webster (Syracuse University), and David Wu (Lehigh University).

The 2004 prizewinners are as follows.

First Place
Retsef Levi, Cornell University
Martin Pál, Rutgers University
Advisors: Robin O. Roundy, David B. Shmoys, and Eva Tardos
"Approximation Algorithms for Stocastic Inventory Control Models"

Second Place
Ravi Subramanian, University of Michigan
Advisors: Sudheer Gupta and Brian Talbot
"Emissions Compliance Strategies: A Permit Auction Model"

Honorable Mention
Xinxin Hu, University of Michigan
Advisors: Izak Duenyas and Roman Kapuscinski
"Optimal Joint Inventory and Transshipment Control Under Uncertain Capacity"

Finalists
Gad Allon, Columbia University
Advisor: Awi Federgruen
"Competition in Service Industries"

Felipe Caro, Massachusetts Institute of Technology
Advisor: Jérémie Gallien
"Dynamic Assortment with Demand Learning for Short Life–Cycle Consumer Goods"

Holly S. Lutze, University of Texas at Dallas
Advisor: Özalp Özer
"Promised Leadtime Contracts and Renegotiation Incentives
Under Asymmetric Information"

The MSOM Society Student Paper Competition: *Extended Abstracts of 2004 Winners*
Manufacturing & Service Operations Management 7(1), pp. 81–99, ©2005 INFORMS

Approximation Algorithms for Stochastic Inventory Control Models

Retsef Levi
School of ORIE, Cornell University, Ithaca, New York 14853, rl227@cornell.edu

Martin Pál
DIMACS Center, Rutgers University, Piscataway, New Jersey 08854-8018, mpal@dimacs.rutgers.edu

Robin O. Roundy (Advisor)
School of ORIE, Cornell University, Ithaca, New York 14853, robin@orie.cornell.edu,

David B. Shmoys (Advisor)
School of ORIE, and Department of Computer Science, Cornell University, Ithaca, New York 14853, shmoys@cs.cornell.edu

In this work we address the longstanding problem of finding computationally efficient and provably good inventory control policies in supply chains with correlated and nonstationary (time-dependent) stochastic demands. This problem arises in many domains and has many practical applications, such as dynamic forecast updates (for some applications see Erkip et al. 1990 and Lee et al. 1999). We consider two classical models, the *periodic-review stochastic inventory control problem* and the *stochastic lot-sizing problem* with correlated and nonstationary demands. Here the correlation is intertemporal, i.e., what we observe in the current period changes our forecast for the demand in future periods. We provide what we believe to be the first computationally efficient policies with constant worst-case performance guarantees; that is, there exists a constant C such that, for any given joint distribution of the demands, the expected cost of the policy is guaranteed to be within C times the expected cost of an optimal policy. More specifically, we provide a worst-case performance guarantee of 2 for the periodic-review stochastic inventory control problem, and a performance guarantee of 3 for the stochastic lot-sizing problem.

The Models

The details of the periodic-review stochastic inventory control problem are as follows. A sequence of random demands for a single commodity at a single location occurs over a finite planning horizon of T discrete periods. The random demands over the T periods can be nonstationary and correlated. The goal is to coordinate a sequence of orders over the planning horizon aiming to satisfy these demands with minimum expected cost. In each period we can order any number of units that are assumed to arrive only after a leadtime of L periods. We consider a traditional cost structure with per-unit ordering, holding, and backlogging penalty costs that are incurred at the end of each period. The cost parameters are time dependent and the only assumption is that we do not have a speculative motivation to hold inventory or have shortages.

In the stochastic lot-sizing problem, we consider, in addition, a fixed ordering cost that is incurred in each period in which an order is placed (regardless of its size), but we assume that there is no leadtime ($L = 0$).

In both models, the goal is to find a policy of orders with minimum expected overall discounted cost over the given planning horizon.

The assumptions that we make on the demand distributions are very mild and generalize all of the currently known approaches in the literature to model correlation and nonstationarity of demands over time (for details about the different approaches, we refer the reader to Iida and Zipkin 2001 and Lingxiu and Lee 2003). As part of the model, we will assume that at the beginning of each period, we are given what we call an *information set* that contains all of the information that is available at the beginning of the time period (e.g., the realized demands so far and external information that becomes available). In addition, we assume that in each period there is a known conditional joint distribution of the future demands that is a function of the observed information set at the beginning of the period (but is independent of the specific inventory policy). The only assumption on the demands is that in each period and for each observed information set, all the future demands are well defined and have finite mean. We consider only policies that are *nonanticipatory*, i.e., in each period they can use only the current information set.

Related Literature

These models have attracted the attention of many researchers over the years, and there exists a huge body of related literature. The dominant paradigm in almost all of the existing literature has been to formulate these models using a dynamic programming framework. This framework has turned out to be very effective in characterizing the optimal policy of the overall system. Surprisingly, the optimal policies for these rather complex models follow simple forms, known as *state-dependent base-stock policies* (see Iida and Zipkin 2001, Lingxiu and Lee 2003, and Zipkin 2000 for details). In each period, there exists an optimal target base-stock level that is determined only by the given (observed) information set. The optimal policy aims to keep the inventory level at each period as close as possible to the target base-stock level. That is, it orders up to the target level whenever the inventory level at the beginning of the period is below that level, and orders nothing otherwise.

Unfortunately, these rather simple forms of policies do not always lead to efficient algorithms for computing the optimal policies. This is especially true in the presence of correlated and nonstationary demands, which cause the state space of the relevant dynamic programs to grow exponentially and explode very fast. This phenomenon is known as *the curse of dimensionality*. Moreover, because of this phenomenon, it seems unlikely that there exists an efficient algorithm to solve these huge dynamic programs.

Muharremoglu and Tsitsiklis (see Muharremoglu and Tsitsiklis 2001) have proposed an alternative approach to the dynamic programming framework. They have observed that this problem can be decoupled into a series of *unit supply-demand subproblems*, where each subproblem corresponds to a single unit of supply and a single unit of demand that are matched together. This novel approach enabled them to substantially simplify some of the dynamic programming-based proofs on the structure of optimal policies, as well as to prove several important new structural results. However, their computational methods are essentially the dynamic programming approach applied to the unit subproblems, and hence they suffer from similar problems in the presence of correlated and nonstationary demand.

As a result of this apparent computational intractability, many researchers have attempted to construct computationally efficient (but suboptimal) heuristics for these problems. However, we are aware of no computationally efficient policies for which there exist constant performance guarantees. For details on some of the proposed heuristics and a discussion of others, see Lingxiu and Lee (2003), Lu et al. (2003), and Iida and Zipkin (2001). One specific class of suboptimal policies that has attracted a lot of attention is the class of *myopic policies*. In a myopic policy, in each period we attempt to minimize the expected cost for that period, ignoring the impact on the cost in future periods. The myopic policy is attractive because it yields a base-stock policy that is easy to compute online, that is, it does not require information on the control policy in future periods. In many cases, the myopic policy seems to perform well (see, for example, Veinott 1965, Iida and Zipkin 2001, Lu et al. 2003). However, in many other cases, especially when the demand can drop significantly from period to period, the myopic policy performs poorly and even arbitrarily badly (see Levi et al. 2004a).

Our Work

Our work is distinct from the existing literature in several significant ways, and is based on the following three novel ideas.

Marginal cost accounting. We introduce a novel approach for cost accounting in stochastic inventory control problems. We use the convention that units in inventory are consumed on a first-ordered-first-consumed basis. This implies that once we place an order of a certain number of units in some period, then the expected ordering and holding cost that these units are going to incur over the rest of the planning horizon is a function only of future demands, not of future orders. Hence, with each period, we can associate the overall expected ordering and holding cost that is going to be incurred over the entire horizon by the units ordered in this period. This new way of marginal cost accounting is significantly different from the dynamic programming approach, which, in each period, accounts only for the costs that are incurred in that period. We believe that this new approach will have more applications in the future in analyzing stochastic inventory control problems.

Cost balancing. The idea of cost balancing was used in the past to construct heuristics with constant performance guarantees for deterministic inventory problems. The most well-known example is the *part-period cost-balancing heuristic* of Silver and Meal (1973) for the lot-sizing problem. We are not aware of any application of these ideas to stochastic inventory control problems. For the periodic-review inventory control problem we propose what we call *a dual-balancing policy*. In this policy we balance, in each period, the expected marginal ordering and holding cost incurred by the units ordered in this period against the expected backlogging penalty cost a leadtime ahead. Conditioned on the observed information set at the beginning of the period, these two costs are both functions of the size of the order placed in the current period. We then order exactly the amount needed to make these two costs equal to each other. For the stochastic lot-sizing problem we propose what we call *a triple-balancing policy* that balances the fixed ordering cost, the marginal holding cost, and the backlogging penalty cost.

Non-base-stock policies. Our policies are not state-dependent base-stock policies. This enables us to use, in each period, the distributional information about the future demands beyond the current period (unlike the myopic policy) without the burden of solving huge dynamic programs. Moreover, our policies can be easily implemented online (i.e., independent of future decisions) and are simple, both conceptually and computationally.

Our Results

Using these ideas, we provide what is called a 2-approximation algorithm for the periodic-review stochastic inventory control problem; that is, the expected cost of our policy is guaranteed to be no more than twice the expected cost of an optimal policy. Our result is valid for all known approaches used to model correlated and nonstationary demands. We also note that this guarantee refers only to the worst-case performance, and it is likely that on average the performance would be significantly better. We then use a standard cost transformation to achieve significantly better guarantees if the ordering cost is the dominant part of the overall cost, as is the case in many real-life situations. Finally, we show how to extend the result to the case where the leadtimes are stochastic under the assumption of noncrossing orders. For the stochastic lot-sizing problem, we provide a 3-approximation algorithm. This is again a worst-case analysis and we would expect the typical performance to be much better. In addition, we present an extended family of myopic policies. In particular, we generate lower bounds on the optimal base-stock levels, which are especially effective given the fact that the classical myopic policy is known to generate respective upper bounds.

Conclusions and Subsequent Work

In this work, we have proposed a new approach for devising provably good policies for stochastic inventory control problems with time-dependent and correlated demands. These policies are simple, both conceptually and computationally. Moreover, this approach leads to a new class of policies for these hard models that combines traditional base-stock policies with the new ideas introduced in this paper. We also believe that the new ideas underlying our policies will have application in additional stochastic inventory control models. In particular, in a subsequent work (see Levi et al. 2004b), we were able to construct a 2-approximation algorithm for the capacitated periodic-review inventory control problem, where there is a capacity constraint on the size of the order in each period.

A full version of this paper is available on request.

Acknowledgments
The research of the first author is partially supported by NSF Grants CCR-9912422, CCR-0430682, DMI-0075627, and by a grant from Motorola. The research of the second author is supported by ONR Grant N00014-98-1-0589. The first author thanks Shane Henderson for stimulating discussions and helpful suggestions that made the paper significantly better.

References

Erkip, N., W. H. Hausman, S. Nahmias. 1990. Optimal centralized ordering policies in multi-echelon inventory systems with correlated demands. *Management Sci.* **36** 381–392.

Iida, T., P. H. Zipkin. 2001. Approximate solutions of a dynamic forecast-inventory model. Working paper, Fuqua School of Business, Duke University, Durham, NC.

Lee, H. L., K. C. So, C. S. Tang. 1999. The value of information sharing in two-level supply chains. *Management Sci.* **46** 626–643.

Levi, R., M. Pál, R. O. Roundy, D. B. Shmoys. 2004a. Approximation algorithms for stochastic inventory control models. Technical Report TR1412, ORIE Department, Cornell University, Ithaca, NY.

Levi, R., R. O. Roundy, D. B. Shmoys, V. A. Troung. 2004b. Approximation algorithms for capacitated stochastic inventory control models. Working paper, School of ORIE, Cornell University, Ithaca, NY.

Lingxiu, D., H. L. Lee. 2003. Optimal policies and approximations for a serial multi-echelon inventory system with time-correlated demand. *Oper. Res.* **51** 969–980.

Lu, X., J. S. Song, A. C. Regan. 2003. Inventory planning with forecast updates: Approximate solutions and cost error bounds. Working paper, Fuqua School of Business, Duke University, Durham, NC.

Muharremoglu, A., J. N. Tsitsiklis. 2001. A single-unit decomposition approach to multiechelon inventory systems. Working paper, School of Business, Columbia University, New York.

Silver, E. A., H. C. Meal. 1973. A heuristic selecting lot-size requirements for the case of a deterministic time varying demand rate and discrete opportunities for replenishment. *Production Inventory Management* **14** 64–74.

Veinott, A. F. 1965. Optimal policy for a multi-product, dynamic, non-stationary inventory problem. *Oper. Res.* **12** 206–222.

Zipkin, P. H. 2000. *Foundations of Inventory Management*. McGraw-Hill Companies, Inc., New York.

Emissions Compliance Strategies: A Permit Auction Model

Ravi Subramanian
Stephen M. Ross School of Business, University of Michigan, 701 Tappan Street, Ann Arbor, Michigan 48109, rsubramz@bus.umich.edu

Advisors: Sudheer Gupta, Brian Talbot
Stephen M. Ross School of Business, University of Michigan, 701 Tappan Street, Ann Arbor, Michigan 48109

Market-based policy instruments are now considered for virtually every environmental problem that is raised, ranging from endangered species preservation to the greenhouse effect and global climate change (Stavins 1998). Emissions-trading programs for specific pollutants such as lead, sulfur dioxide (SO_2), nitrogen oxides (NO_x), carbon dioxide (CO_2), and volatile organic compounds (VOCs) are prime examples. Emissions trading refers to a market-based mechanism for emissions control that allows parties to buy and sell permits for emissions or credits for reductions in emissions of certain pollutants. It differs from a traditional regulatory or command-and-control approach that relies solely on an agency, usually the government, to issue standards and specific directives on the amount by which emitters must reduce their emissions, how they must do so, and the penalties for failure. In an emissions-trading program, emitters are allocated or permitted a limited amount of emissions. The total number of permits corresponds to the overall emissions target of the covered sources. The fact that the target is less than "business as usual" emissions creates permit scarcity, resulting in a market price for permits. Emitters are responsible for ensuring that they hold sufficient permits to offset their emissions; they have the flexibility to cost-effectively administer compliance levers. The concept of emissions trading has grown from a theoretical curiosity into a central idea in environmental regulation. The theory of emissions trading is well developed (Tietenberg 2001), and attention has now shifted from whether tradable emissions schemes should be implemented to how they should be implemented (Muller and Mestelman 1998).

Most research on market-based mechanisms for emissions control has focused on evaluating the efficacy of emissions-trading programs vis-à-vis other approaches a regulator has at its disposal, such as taxes, subsidies, and standards. We instead take the perspective of a profit-maximizing firm and focus on the trade-offs among different strategies that the firm has at its disposal in complying with a pollution control program chosen by the regulator, especially since the crux of any market-based program for pollution control is to allow firms to flexibly and cost-effectively apply compliance levers. To the regulator, knowledge about the interrelationships among firm levers for compliance is crucial given the goals of pollution control and a desired increasing level of stringency in the stipulation of pollution limits. From a firm's perspective, decisions such as the extent of investment in pollution abatement, permit procurement strategy, and production level need to be made given a policy stipulation, the accompanying cost of compliance with the policy, and the goal of profit maximization. To the best of our knowledge, ours is the first framework that provides for an assessment of compliance strategies in a permit-based program for emissions control, for varying policy stringency and firm dirtiness levels.

We model a three-stage game in an oligopoly where firms invest in abatement to reduce emissions, participate in a "share" auction for permits, produce output, and hence, generate emissions in proportion to output produced. In the first stage of the game, each of the $n \geq 2$ firms decides its abatement level μ_i from an investment $\xi\mu_i^2$ in pollution mitigation (e.g., investments that enable the use of cleaner inputs or the scrubbing of flue gases, or investments in process improvements). In the second stage, firms bid for emissions permits in a sealed-bid uniform price share auction. Each firm submits a sealed tender specifying a schedule of prices for varying fractional shares of the total available pollution permits. For the base model, we consider situations in which the bidders are symmetric, and where the optimal strategy is a symmetric subgame-perfect Nash equilibrium. The assumption of symmetry is relaxed in our treatment of a two-firm asymmetric game. The auction results in a permit allocation vector $\beta = (\beta_1, \beta_2, \ldots, \beta_n)$, and a market clearing price e per permit. $B = \sum_{i=1}^{n} \beta_i$, is the total permissible pollution level. The regulator's choice of the target emissions level B is assumed to be exogenous to the model. We focus on auctioning as the primary means for allocating permits, but we do provide a discussion on how our approach can be extended to incorporate grandfathering in addition to auctioning. Each permit allows a firm to emit a unit of the pollutant. In the third stage of the game, each firm i produces output y_i, which results in a pollution level $(\alpha/(1+\mu_i))y_i$, and firms redeem their permits against their pollution levels. The coefficient α represents the state of current technology in terms of the level of emissions per unit of output and is used to quantitatively define how dirty (or clean) the industry or firm is. We denote $\tilde{\alpha}_i := \alpha/(1+\mu_i)$. We assume that the penalty for not having the requisite number of permits to account for emissions is large enough so that noncompliance is deterred. We consider two distinct demand situations in the end-product market. In the first case, each firm faces an independent, inverse demand function $p_i = a - by_i$. In the second case, the n firms compete in a Cournot fashion and face an inverse total demand function $p_i = p = \hat{a} - \hat{b}Y$, where $Y = \sum_{i=1}^{n} y_i$. In practice, unused permits can be banked for future use or trade. Because we treat a single-period problem, we assume a terminal value of u per unused permit. This value could represent either the value of a permit in a secondary market or the net present value of benefits accruing from future use of a permit left over at the end of the first period. The unit cost of production is c, which is assumed to be constant in the base model. However, we extend our results to the case where the cost of production is influenced by abatement efforts.

The reader is directed to Subramanian et al. (2004) for the complete paper. However, to demonstrate the methodology used in the paper, we present an outline of our analysis for the case of independent demands. An example of such a demand scenario would be the U.S. EPA Acid Rain Program for restricting SO_2 emissions from fossil-fuel-fired electric utilities that operate mostly as local monopolies. We proceed conventionally, by backward induction, to ascertain the equilibria in the different stages of the game. The optimization problem of a representative firm in the production subgame, given abatement levels and a permit allocation vector, is:

$$\text{Maximize}_{\{y\}} \quad \Pi = (a-by)y + (\beta - \tilde{\alpha}y)^+ u - \xi\mu^2 - e\beta - cy,$$
$$\text{subject to: } \tilde{\alpha}y \leq \beta, \quad y \geq 0.$$

Because firms are constrained in producing profit-maximizing output levels by the availability of permits, $y^* = \beta/\tilde{\alpha}$ is the equilibrium output level in the production subgame. To determine its bidding strategy in the auction, a firm must anticipate the outcome of the production subgame where one emissions permit is redeemed for each unit of emissions generated in the process of production. We determine a firm's marginal value function for permits endogenously through the shadow price corresponding to the permit-availability constraint in the production subgame. This is in contrast to most of the auctions literature, where the value placed by a bidder on the items being auctioned is typically specified exogenously or is assumed to be drawn from an exogenous distribution. The marginal value function for permits is given by

$$v(\beta) = \frac{(a-c)}{\tilde{\alpha}} - \frac{2b}{\tilde{\alpha}^2}\beta := \sigma - \lambda\beta$$

where σ and λ are constants, given the abatement level chosen in the first stage. The marginal value function is critical in deducing the optimal bidding strategy and equilibrium permit price in the second stage of the game, where firms participate in a uniform price share auction for permits. We show that in the auction subgame it is an optimal strategy for a firm to submit a schedule such that at each price e, the requested number of permits is

$$\beta(e) = \left(\frac{1 - 2e/(n\sigma - \lambda B)}{n-1}\right)B.$$

The resulting equilibrium price is $e^* = \frac{1}{2}(\sigma - \lambda(B/n))$, and the equilibrium number of permits received by each firm is $\beta^* = B/n$. Concomitantly, we generalize Wilson's (1979) analysis of share auctions to the case when the marginal value placed by bidders on the items being auctioned decreases linearly with the share fraction. In Wilson (1979), the value of a share is proportional to the share fraction. However, in practice it is more likely that the marginal value a bidder places on the items being auctioned decreases with the share fraction. Finally, in the first stage of the game, each firm decides its level of investment in pollution abatement, anticipating the equilibrium number of permits that it will secure in the auction and the output that it will be able to produce. After substituting equilibrium results from the production and auction subgames, the profit function is concave in abatement μ. $\mu^* = (B(a-c))/(4\xi\alpha n)$ is the optimum abatement level.

We conduct a similar analysis for the case when firms compete in the output market in a Cournot fashion. In addition, we provide the following extensions in order to test the sensitivity of our results to some of the model assumptions: (i) a variation on the specification of the emissions function where abatement efforts result in a fixed amount of abatement (e.g., as in sequestration using landfills); (ii) the possibility that permits can be grandfathered in addition to being auctioned; (iii) the possibility that investment in abatement can affect the cost of production either positively or negatively; and (iv) firm heterogeneity in dirtiness, costs of abatement, and costs of production.

It might seem at first glance that varying the total number of permits will influence a dirtier industry to a greater extent than a cleaner one to engage in abatement. However, we find that varying the number of available permits decreasingly influences abatement for increasing industry dirtiness levels. Our model also suggests that equilibrium permit price can be low in a relatively dirty industry because the benefit from expanded production capacity decreasingly outweighs the costs, resulting in a decrease in the value placed on permits, a lower permit price in the auction, and lower production levels. Reductions in the number of available permits mean that firms in a relatively dirty industry are progressively driven to lower and lower output levels, and perhaps to extinction. Industry structure also influences equilibrium outcomes. Interestingly, in the case of Cournot competition in the output market, reducing the number of permits offered in the auction can increase equilibrium firm profit in a relatively clean industry. The regulator can then amicably enforce reductions in the number of available permits. This is favorable from the viewpoint of a regulator aiming to increase the level of stringency in pollution norms. Traditionally, firms lobby to roll back pollution standards because of the lingering belief that stricter environmental regulations erode competitiveness (Porter and van der Linde 1995). However, in a competitive setting, an excessive number of permits is detrimental to firm profit in a relatively clean industry, because of large capacities that bring down product price and firm profits. Our model's predictions can also help a regulator attempting to control a set of pollutants via permit auctions to choose the pollutants to be controlled depending on the characteristics of the industries in which the pollutants are generated. By targeting a pollutant in a competitive and relatively clean industry, the regulator can ensure that firm profitability will not be negatively affected as a result of stringency. On the other hand, if the regulator intends to exclude a pollutant to the greatest possible extent from discharge streams in a dirty industry, it can enforce stringent pollution caps, which will drive down output levels. Our results are robust to the model's extensions.

References

Muller, A. R., S. Mestelman. 1998. What have we learned from emissions trading experiments. *Managerial Decision Econom.* **19** 225–238.

Porter, M. E., C. van der Linde. 1995. Green and competitive: Ending the stalemate. *Harvard Bus. Rev.* (September–October) 120–134.

Stavins, R. N. 1998. What can we learn from the grand policy experiment? Lessons from SO_2 allowance trading. *J. Econom. Perspectives* **12**(3) 69–88.

Subramanian, R., S. Gupta, B. Talbot. 2004. Emissions compliance strategies: A permit auction model. Working Paper 200400885, Stephen M. Ross School of Business, University of Michigan, Ann Arbor, MI.

Tietenberg, T. H. 2001. *Emissions Trading Programs*. Ashgate Publishing Ltd., Aldershot, UK.

Wilson, R. 1979. Auctions of shares. *Quart. J. Econom.* **93**(4) 675–689.

Optimal Joint Inventory and Transshipment Control Under Uncertain Capacity

Xinxin Hu
Ross School of Business, University of Michigan, Ann Arbor, Michigan 48109, huxinxin@bus.umich.edu

Advisors: Izak Duenyas, Roman Kapuscinski
Ross School of Business, University of Michigan, Ann Arbor, Michigan 48109

1. Introduction

Consider a firm that produces the same product in multiple locations, but faces demand and capacity uncertainty. The capacity uncertainty is caused by factors such as downtime, quality problems, yield, etc. The firm faces two related decisions: (1) How much should it produce at each location, and (2) how much should it transship from one location to another? Even though the literature on transshipment is rich, it usually ignores the effects of capacity uncertainty, and our aim is to gain insight into how capacity uncertainty affects both these decisions.

The problem we describe is very common in industry. For example, we recently worked with a diesel engine manufacturer that has multiple locations where castings are made. The capacity of the plants making the castings in any week was random due to quality problems and, therefore, the company was exploring transshipment from one location to another to satisfy engine plants' demands for castings. We observed similar issues in the case of a major paper manufacturer that produces paper cups in multiple locations in the United States, as well as a major newspaper-ink manufacturer with over 20 plants in the United States. In all cases, products would be transshipped from one plant to another plant's markets when capacity in a plant was low in a given period. However, we observed that the actual production policies of the plants did not take into account the fact that such transshipments may occur. In some situations, we also observed that plant management was reluctant to transship beyond a certain amount, due to fear that they may face a shortage next period if their inventory levels are down significantly. All of these observations motivated us to explore how optimal transshipment and production decisions should be made jointly and how the level of demand and capacity uncertainty affect the behavior of optimal policies.

2. The Model

We consider two manufacturing facilities, each serving its individual market, through multiple time periods. The facilities face uncertain capacity, characterized by capacity distributions that are independent in time and of each other. The facilities also face demand uncertainty. The stochastic demand distributions are independent in time, but can be correlated for any given period across the two facilities.

In any period, production decisions are made first: The firm decides how much it will *attempt* to produce in each of the facilities that period. Then, the capacities and demands are realized for both facilities. The actual production is the minimum of planned production and the realized capacity. Finally, decisions are made regarding transshipment of inventory between facilities. We assume that demand that is unsatisfied after transshipment is lost. The firm earns linear revenues on satisfied demand and incurs linear production, holding, and transshipment costs. The objective is to maximize the joint discounted profit for both facilities. Let $i, j = 1, 2$ denote the facilities, and we denote c_i, h_i, s_{ij}, and r_i as facility i's variable production cost, variable holding cost, variable transshipment cost to facility j, and its unit revenue. We assume that marginal profit is always higher when using a unit to satisfy demand at the facility where it was produced rather than transshipping to satisfy demand at the other facility, i.e., $r_2 - s_{12} - r_1 \le 0$, $r_1 - s_{21} - r_2 \le 0$.

Consider $N+1$ periods, where the ending period is $N+1$. For facility i and period k we denote:

$x_i^k =$ starting inventory level;
$y_i^k =$ planned production target (starting inventory + planned production);
$T_i^k =$ stochastic capacity with *pdf* f_i^k and *cdf* F_i^k;
$\bar{y}_i^k =$ achieved production target, i.e., $\bar{y}_i^k$ is the realization of $\overline{Y}_i^k = y_i^k \wedge (x_i^k + T_i^k)$;
$D_i^k =$ stochastic demand, with *pdf* q_i^k and *cdf* Q_i^k;
$z_i^k =$ intermediate inventory position after demand is realized but before transshipment, i.e., z_i^k is the realization of $Z_i^k = \overline{Y}_i^k - D_i^k$;
$\hat{z}_i^k =$ inventory position after transshipment, i.e., $z_1^k + z_2^k = \hat{z}_1^k + \hat{z}_2^k$ and $(z_i^k - \hat{z}_i^k)^+$ is the quantity transshipped from facility i to the other one;
$\alpha_k =$ discount rate, $0 \le \alpha_k \le 1$.

Realizations of D_i^k and of T_i^k are denoted by d_i^k and t_i^k.

We analyze the problem using two-stage backward induction and denote $G_*^k(\mathbf{x}^k)$ as the optimal discounted profit-to-go from period k with starting inventory $\mathbf{x}^k$. We use boldface notation to represent two-dimensional vectors and formulate the model as follows:

Stage One:

$$G_*^k(\mathbf{x}^k) = \max_{\mathbf{y}^k \ge \mathbf{x}^k} \mathrm{E}_{\mathbf{T}^k, \mathbf{D}^k}\{-\mathbf{c}(\mathbf{y}^k \wedge (\mathbf{x}^k + \mathbf{T}^k) - \mathbf{x}^k) + \mathbf{r}\mathbf{D}^k + G_v^k(\mathbf{y}^k \wedge (\mathbf{x}^k + \mathbf{T}^k) - \mathbf{D}^k)\}, \quad (1)$$

Stage Two:

$$G_v^k(\mathbf{z}^k) = \max_{\hat{z}_1^k + \hat{z}_2^k = z_1^k + z_2^k} \bar{G}^k(\mathbf{z}^k, \hat{\mathbf{z}}^k) \quad (2)$$

where

$$\bar{G}^k(\mathbf{z}^k, \hat{\mathbf{z}}^k) = -\mathbf{r}(\hat{\mathbf{z}}^k)^- - \mathbf{h}(\hat{\mathbf{z}}^k)^+ - \mathbf{s}(\mathbf{z}^k - \hat{\mathbf{z}}^k)^+ + \alpha_k G_*^{k+1}((\hat{\mathbf{z}}^k)^+) \quad (3)$$

and $G_*^{N+1}(\mathbf{x}^{N+1}) \equiv 0$.

It is important to note that the function maximized in Stage 1 is not concave in $\mathbf{y}^k$. We can prove, however, that the profit function behaves predictably and that $G_*^k(\mathbf{x}^k)$ is concave and submodular.

3. Optimal Policy

3.1. Transshipment Policy

The policy defined below is the basic structure of the optimal transshipment policy for our model:

Definition 1. Consider intermediate inventories (z_1, z_2), with $z = z_1 + z_2$. Define:

(a) state-dependent rationing policy for facility i, SR $i(\chi_i(z))$, as follows: Facility i transships $(z_i - \chi_i(z))^+$ to facility $3-i$.

(b) floor-rationing policy for facility i, FR $i(\underline{\chi}_i, \bar{\chi}_i(z))$, as a state-dependent rationing policy SR $i(\chi_i(z))$, where $\underline{\chi}_i$ is a constant, and $\chi_i(z) = \underline{\chi}_i$ for $z < \underline{\chi}_i$, and $\chi_i(z) = \bar{\chi}_i(z)$ for $z \ge \underline{\chi}_i$. We refer to $\underline{\chi}_i$ as facility i's floor.

Theorem 1 (Optimal Transshipment for Multi-. period Problem, $h_1 \ge h_2$). *In period k, let the total intermediate inventories be $z^k = z_1^k + z_2^k$. The optimal transshipment policy is defined by floor-rationing policies*, FR1

$(\underline{\chi}_1^k, \bar{\chi}_1^k(z^k))$ *for Facility* 1, *and* FR2$(\underline{\chi}_2^k, z^k)$ *for Facility* 2, *where*

(1) *For* $z^k > \underline{\chi}_1^k$, $\underline{\chi}_1^k \leq \bar{\chi}_1^k(z^k) \leq z^k$; *particularly, if* $h_1 = h_2$, $\bar{\chi}_1^k(z^k) = z^k$.

(2) $0 \leq \partial\bar{\chi}_1^k/\partial z^k \leq 1$.

(3) $\underline{\chi}_i^k$ *is nonincreasing in the current period* h_i *and* r_{3-i}, *nondecreasing in the current period* $s_{i,3-i}$, *and independent of current period* h_{3-i}, r_i, *and* $s_{3-i,i}$.

(4) $\bar{\chi}_1^k(z^k)$ *is nonincreasing in the current period* h_1 *and nondecreasing in the current period* h_2 *and* s_{12}.

Theorem 1 implies that it may be profitable to ship inventory that is not immediately needed from higher-holding-cost Facility 1 to Facility 2. At the same time, both facilities may also have an incentive to ration their inventory. What makes this problem interesting is that, depending on the intermediate inventory (after production and demands are realized), a range of behaviors may be optimal: rationing, transshipping full needed amounts, and transshipping inventory even though none is immediately needed by the other facility.

3.2. Production Policy

We are interested in how our results depend on the capacity distribution. We define *uncertain* capacity as the most general case where capacity is stochastic. We define *certain-limited* capacity as the case where capacity in every period is deterministic, but finite. Finally, the third case is where we assume no capacity limitation, i.e., *infinite* capacity.

Case 1. *Two Facilities with Uncertain Capacities.* Consider the general case where both facilities may have uncertain capacities, i.e., $\Pr(T_i^k = \infty) < 1$.

Theorem 2. *The optimal production policy at facility* i *is a function of the other facility's inventory* x_{3-i}^k *and is defined by two thresholds* $\underline{x}_i^k \leq \bar{x}_i^k$, (*also functions of* x_{3-i}^k)[1]

$$\bar{x}_i^k = \bar{x}_i^k(x_{3-i}^k) = (\inf\{x_i^k: F_i^k(y_i^{*k} - x_i^k) = 0\}) \wedge (\inf\{x_i^k: y_{3-i}^{*k} = x_{3-i}^k\})$$

$$\underline{x}_i^k = \underline{x}_i^k(x_{3-i}^k) = \sup\{x_i^k: F_{3-i}^k(y_{3-i}^{*k} - x_{3-i}^k) = 1\} \wedge \bar{x}_i^k(x_{3-i}^k),$$

two order-up-to levels $\underline{y}_i^{*k} \leq \bar{y}_i^{*k}$, *and one function* $y_i^{*k}(\mathbf{x}^k)$ *such that* (1) *if* $x_i^k < \underline{x}_i^k$, *then produce up to* $\underline{y}_i^{*k}$; (2) *if* $x_i^k > \bar{x}_i^k$, *then produce up to* $\bar{y}_i^{*k}$; *and* (3) *if* $\underline{x}_i^k \leq x_i^k \leq \bar{x}_i^k$, *then produce to* $y_i^{*k}(\mathbf{x}^k)$.

Furthermore, all thresholds, up-to levels, and $y_i^{*k}(\mathbf{x}^k)$ *are nonincreasing in* x_{3-i}^k, $y_i^{*k}(\mathbf{x}^k)$ *is nondecreasing in* x_i^k, *but* $y_i^{*k}(\mathbf{x}^k) - x_i^k$ *is nonincreasing in* x_i^k.

[1] $\inf\varnothing = \sup\mathbb{R} = \infty$ and $\sup\varnothing = \inf\mathbb{R} = 0$.

Case 2. *One Facility with Uncertain Capacity.*

Theorem 3. *Let Facility* 1 *have uncertain capacity and Facility* 2 *have infinite capacity.*

(1) *Facility* 1*'s optimal policy is exactly as defined in Theorem* 2 *with* $\underline{x}_1^{*k} = 0$ *and* $\underline{y}_1^{*k} = 0$. *Furthermore, when* $x_2^k < y_2^{*k}(x_1^k)$, *the production target of Facility* 1 *is independent of Facility* 2*'s starting inventory* x_2^k.

(2) *Facility* 2 *produces up to* $y_2^{*k}(x_1^k)$, *and* $y_2^{*k}(x_1^k)$ *decreases in Facility* 1*'s inventory* x_1^k.

Case 3. *Two Facilities with Certain-Limited Capacities.* When two facilities have certain-limited capacities C_i^k $(i = 1, 2)$, the optimal policy is defined by modified order-up-to levels that are functions of the other facility's starting inventory, i.e., facility i produces $\min((y_k^{*i}(x_{3-i}^k) - x_i^k)^+, C_i^k)$.

Case 4. *Two Facilities with Infinite Capacities.* In this case, each facility uses a base-stock policy and can ignore the other facility's inventory level in determining its production quantity, as long as both starting levels are below the base stocks.

4. Sensitivity of Optimal Policy

A number of parameters influence the optimal policies in the expected direction. For instance, when one location's capacity stochastically increases, that facility's optimal inventory target increases and the other facility's target decreases. Similarly, the effects of production and transshipment costs are also predictable. If a facility's production or transshipment cost is increased, the facility's inventory target is decreased while the other's is increased. Less intuitive are the effects of changing demand and the changes in holding costs and revenues. (We only outline the main behavior here, while the paper describes them in detail.)

4.1. Sensitivity to Demand

If Facility 1's demand stochastically increases in the current period k, then the sum of the production targets at the two facilities is increased. Furthermore, if Facility 1 has infinite or certain-limited capacity, its production target is increased.

It is important, however, to note that the above statement fails when Facility 1 has uncertain capacity. Consider the situation where Facility 2 has stochastically larger capacity than Facility 1. Facing stochastically larger demand, it may be optimal for Facility 1 to rely on Facility 2's capacity to help deal with production uncertainty. This may result in pushing down Facility 1's target and pushing up Facility 2's target.

4.2. Sensitivity to Holding Cost and Revenue Coefficients

(1) Sensitivity to Holding Cost. As Facility 1's current-period holding cost increases, the sum of the production targets at both facilities is decreased. If Facility 1 has infinite or certain-limited capacity, then its production target decreases; otherwise, increasing its holding cost may lead to a higher production target at Facility 1. This behavior can be explained by the fact that due to uncertain capacity, increased target production level at Facility 1 may result in lower total expected remaining inventory in the system.

(2) Sensitivity to Revenue. Intuitively, a facility's inventory target should be increasing in its revenue, but uncertainty in capacity complicates this effect. Similarly to above, the inventory target of a facility with infinite or certain-limited capacity increases in its revenue. The sum of the inventory targets is always increased with increasing revenue. A facility with uncertain capacity, however, may decrease its inventory target.

5. Discussion and Extensions

Our model considers optimal production and transshipment control in a centralized system with stationary linear cost and revenue coefficients and lost sales. As compared to other papers, our model is much more general. We impose only one assumption on the facilities' revenues, $r_2 - s_{12} - r_1 \leq 0$ and $r_1 - s_{21} - r_2 \leq 0$, and even this assumption does not influence the structure of our optimal policy. In fact, the same structure also holds with a model that allows backlogging, nonstationary, and some nonlinear cost-revenue coefficients. The model can also be generalized to allow Markov-modulated capacity and demand processes.

6. Conclusions

In this paper, we fully characterize the optimal production and transshipment policies for a firm that produces in two locations and faces capacity uncertainty. Due to uncertain capacity, the firms ration their inventory that is available for transshipment. We describe the rationing transshipment policy and also production policy, which is state-dependent produce-up-to threshold. We also provide sensitivity of the optimal production and transshipment policies to problem parameters and, in particular, explain how uncertain capacity can lead to counterintuitive behavior, such as produce-up-to limits decreasing for locations that face stochastically higher demand. We finally explore, through a numerical study, when applying the optimal policy is most likely to yield significant benefits compared to simple policies.

Competition in Service Industries

Gad Allon
Graduate School of Business, Columbia University, New York, New York 10027, ga2003@columbia.edu

Advisor: Awi Federgruen
Anderson Graduate School of Management, University of California, Los Angeles, Los Angeles, California 90095

We analyze a general market for an industry of competing service facilities. Firms differentiate themselves by their price levels and the waiting time their customers experience, as well as different attributes not determined directly through competition. A given firm's demand volume may depend on all prices and all (steady-state) waiting-time standards in the industry. The latter may be specified by the *expected* steady-state waiting time or a given (e.g., 95th) percentile of the waiting-time distribution. In some settings, the waiting-time standard is explicitly announced, possibly with monetary compensation offered if a customer's waiting time exceeds the standard. In other cases, it is the waiting-time performance as observed by the clientele or reported by independent organizations. Either way, each firm commits itself to the chosen standard by adopting a sufficiently large capacity level. Different types of competition and equilibrium behavior arise, depending on the industry dynamics through which the firms make their strategic choices. In one case, the firms make all choices simultaneously: *simultaneous competition* (SC). Alternatively, firms may *initially* choose their waiting-time standards, selecting their prices in a *second* stage: *service-level-first competition* (SF). As a third alternative, the sequence of strategic choices may be reversed: *price-first competition* (PF).

Numerous service industries use waiting-time standards as an explicitly advertised competitive instrument. Dominos has offered free-of-charge delivery if pizza delivery were to take more than 30 minutes. Restaurant chains such as Black Angus offer free lunches if lunch is not served within 10 minutes. Banks like Wells Fargo award $5 when a customer waits more than five minutes in line. Various call or contact centers promise that the customer will be helped within one hour, possibly by a call back. Supermarket chains like Lucky launched a "three is a crowd" campaign, guaranteeing that no checkout counter line would have more than three customers waiting. Ameritrade made major inroads into the online discounted brokerage market, waiving commissions for certain types of trades if service were to take more than 10 seconds. As a final example, airlines advertise waiting-time characteristics such as "on-time arrival percentage," while independent government agencies (e.g., the Aviation Consumer Protection Division of the DOT), as well as Internet travel services (e.g., Expedia) report, on a flight-by-flight basis, the average delay and percentage of flights arriving within 15 minutes of schedule. The literature has shown that "on-time arrival percentages" increase significantly with the number of competing carriers on the flight link.

Customers select a specific firm by trading off *three* categories of service attributes: (1) the *price*, (2) the *waiting-time standard*, and (3) all *other attributes*. For example, for competing mail services, the "other attributes" include the convenience of the pick-up process, the ease at which deliveries can be traced, and the likelihood of the packages being damaged. In the restaurant and fast-food industry, the location, ambiance, and the quality of the food are important components of "other attributes," and for Internet service providers, the frequency of service interruption and the quality of the staff. Prior service competition models assume that the first two attributes (i.e., price and waiting time) can be aggregated into a so-called *full price*, usually defined as the direct price plus a multiple of the expected waiting time. This is tantamount to assuming that all customers assign a *specific* cost value to their waiting time *and* that the cost of waiting is simply proportional to the total waiting time.

While consistent with *classical* economic theory, many studies in the *modern* psychology, economics, marketing, and operations literature have demonstrated that both assumptions are often violated. The full-price assumptions reduce the customers' choice to a trade-off between the full price and the "other attributes." Many prior models also assume that *all* customers select a firm with the *lowest* full price, albeit that different customers may be attracted to different firms because of differences in their waiting-time cost rate. This of course amounts to assuming that the firms' services are perfect substitutes, i.e., *no* attributes other than price and waiting time matter, reducing the customers' multidimensional trade-off process to the *full price* as the *single* criterion.

Defining a firm's *service level* as the difference between a given upper-bound benchmark for the waiting-time standard and the *actual* waiting-time standard, we represent a firm's demand rate as a function of *all* prices and service levels in the industry. (We focus primarily on a separable specification that, in addition, is linear in the price vector.) This class of demand models represents *general* trade-offs between the above three categories of attributes. Price and waiting time are treated as truly independent attributes in that, in general, a change in a firm's waiting time (distribution) can *not* be compensated for by a price change that leaves all firms' demand volumes unchanged. We model each firm as an $(M/M/1)$ queueing facility, which receives a given firm-specific price and incurs a given cost per customer served. Each firm incurs a cost per unit of time proportional to its adopted capacity level, determined to satisfy the waiting-time standard under the expected demand rate.

We characterize the equilibrium behavior in the above three possible ways in which prices and service levels may be selected, i.e., SC, PF, and SF. We show that in all three settings an equilibrium pair of price and service-level vectors exists, in full generality, provided the waiting-time benchmark is not excessively large. We also develop efficient procedures to compute the equilibria in the various competition models.

These existence results are in stark contrast to the known behavior in existing service competition models. For example, the seminal model due to Luski (1976) and Levhari and Luski (1978) confines itself to two service providers and assumes all customers choose their provider strictly on the basis of the full price, i.e., the price plus the expected waiting time multiplied with a customer-specific cost rate. Customers' cost rates are independent and identically distributed (i.i.d.). With service rates exogenously given, the competition between the two firms is confined to their price choices only. Whether or not an equilibrium exists in this elementary model remained an open question until it was answered in the *affirmative* by Chen and Wan (2000) for the case in which the firms' service rates are identical, while under nonidentical service rates an example is given where *no* (pure) Nash equilibrium exists. The same example shows that the equilibrium behavior is very *unstable*: As the total market size varies from 1.2 to 1.3, the industry moves from a unique equilibrium, to no equilibrium, to an infinite number of equilibria.

Cachon and Harker (2002), again for the case of two service providers, allows each firm's demand rate to be specified as a function of both firms' full-price values; in this model, customers do not necessarily patronize the lowest full-price provider (i.e., other attributes matter). When the demand rate functions are linear, the known equilibrium results merely exclude the existence of multiple equilibria, and this only when the demand rates are sufficiently large. When the demand rate functions are (truncated) logit functions, the authors examine a specific *symmetric* numerical instance. Varying a single cost rate parameter, the industry moves from a situation with a unique equilibrium under which both firms share the market, to one without any equilibrium, and next to a situation with two equilibria, one with Firm 1 and the other with Firm 2 as the monopoly provider.

To further appreciate the existence results for an equilibrium in the three competition models, note that they apply to an *arbitrary* number of competing service providers. Also, in the SC model, the noncooperative game involves *essentially multidimensional*[1] strategy spaces. Finally, in the PF and SF models, the existence results pertain to *two-stage* games. In the process of analyzing these two-stage games, we characterize the price (service-level) equilibrium that arises under a given vector of service levels (prices) and show how the former varies as a function of the latter. These second-stage "price-only" and "service-only" competition models are of interest by themselves when one of the two strategic variables is specified in a way different than through noncooperative competition.

We cannot guarantee that the equilibrium is unique. In general, the existence of multiple equilibria is unsettling, as it is hard to predict which of the equilibria is adopted by an industry. We show, however, that in our model the set of equilibria always has a *componentwise-largest* and a *componentwise-smallest* pair of equilibrium vectors. In other words, there exists an equilibrium such that each firm's price, as well as its service level, is *higher*, and there exists an equilibrium such that these are *lower* than his price and service level under any other Nash equilibrium. Most importantly, the componentwise-largest pair of price and service-level vectors is preferred by all of the firms. Finally, the schemes used to compute an equilibrium can also be applied to verify numerically whether multiple equilibria exist. Evaluating thousands of instances across a broad spectrum of parameters, we have never encountered a case with multiple equilibria.

[1] A firm's strategy space is *essentially* multidimensional, if each of the strategy variables (e.g., price and service level) impacts on *all* firms' profit functions and these strategy variables cannot be replaced by a single aggregate variable (e.g., the full price).

The set of equilibria is *identical* under the SC and the PF models. Moreover, each firm's equilibrium service level in *any* such equilibrium is *uniquely* determined as a function of *that* firm's characteristics only, and it is a dominant choice for this firm, i.e., with fixed prices, the equilibrium service level is the firm's optimal choice, regardless of what service levels are adopted by its competitors. In contrast, the equilibrium in the SF model differs from that in the other two competition models. Here, a firm's equilibrium service level *does* depend, in general, on the characteristics of the competitors. Assuming the SF model has a unique equilibrium, we derive a simple sufficient condition under which each firm adopts a *higher price* and a *higher service level* while enjoying a *higher* demand volume, compared to the other types of competition. In the presence of multiple equilibria, the same uniform ranking applies to the componentwise-smallest equilibria. Thus, if firms choose and announce their service levels before choosing their price, this will result in higher but more expensive service by all competitors. Because all firms' demand volumes increase as well, this type of competition appears to benefit the consumer. It also suggests that value is added to the consumer when government agencies, industry consortia, or independent organizations periodically report on service levels.

References

Cachon, P. C., P. T. Harker. 2002. Competition and outsourcing with scale economies. *Management Sci.* **48**(10) 1314–1333.

Chen, H., Y.-W. Wan. 2003. Price competition of make-to-order firms. *IIE Trans.* **35**(9) 817–832.

Levhari, D., I. Luski. 1978. Duopoly pricing and waiting lines. *Eur. Econom. Rev.* **11** 17–35.

Luski, I. 1976. On partial equilibrium in a queueing system with two servers. *Rev. Econom. Stud.* **43** 519–525.

Dynamic Assortment with Demand Learning for Short Life–Cycle Consumer Goods

Felipe Caro
Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02142, fcaro@sloan.mit.edu

Advisor: Jérémie Gallien
Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02142, jgallien@mit.edu

1. Introduction

Long development, procurement, and production leadtimes resulting in part from a widespread reliance on overseas suppliers have traditionally constrained fashion retailers to make supply and assortment decisions well in advance of the selling season, when only limited and uncertain demand information is available. With little ability to modify product assortments and order quantities after the season starts and demand forecasts can be refined, many retailers are seemingly cursed with simultaneously missing sales for want of popular products, while having to use markdowns in order to sell the many unpopular products still accumulating in their stores (see Fisher et al. 2000).

Recently, however, a few innovative firms, including Spain-based Zara, Mango, and Japan-based World Co. (sometimes referred to as "fast-fashion" companies), have gone substantially further, implementing product development processes and supply chain architectures that allow them to make *most* product design and assortment decisions *during* the selling season. Remarkably, their higher flexibility and responsiveness is partly achieved through an increased reliance on more costly local production relative to the supply networks of more traditional retailers.

At the operational level, leveraging the ability to introduce and test new products once the season has started motivates a new and important decision problem, which seems key to the success of these fast-fashion companies: Given the constantly evolving demand information available, which products should be included in the assortment at each point in time?

The problem just described seems challenging, in part because it relates to the classical trade-off known as exploration versus exploitation: In each period the retailer must choose between including products in the assortment that he has a "good sense" are profitable (exploitation) or products for which he would like to gather more demand information (exploration). That is, he must decide between being "greedy" based on his current information or trying to learn more about product demand (which might be more profitable in the future). In that respect, the dynamic assortment problem can be seen as a variant of the multiarmed bandit problem with finite horizon and several plays per stage. Each arm represents a product, and pulling an arm is equivalent to including the respective product in the assortment.

2. Model Definition

2.1. Supply and Demand

Consider a retailer selling products in a store during a limited selling season. The set of all products that the retailer may potentially sell is denoted by $\mathcal{S} = \{1, 2, \ldots, S\}$; this set includes both the products already available when the season starts and all the variants and new products that may be designed during the season. The net margin r_s of product $s \in \mathcal{S}$ is assumed to be exogenously given, positive, and constant. We assume that the selling season can be divided into T periods and that at the beginning of each of these periods the product assortment in the store may be revised; time is counted backward and denoted by the index t.

The store's limited shelf space is captured by the constraint that the assortment in each period may include at most N different products out of the S available; we are thus implicitly assuming that all products require the same shelf space. We also assume a perfect inventory replenishment process during each assortment period, so that there are no stockouts or lost sales. Consequently, in our model, realized sales equal total demand, and for each product we focus on assortment inclusion or exclusion as opposed to order quantity. Holding costs are ignored in our formulation.

The demand for each product in the assortment is exogenous and stationary, but stochastic, and we do not capture substitution effects. Specifically, we assume that customers willing to buy one unit of each product s in the assortment arrive at the store according to a Poisson process with an unknown but constant rate γ_s. That is, the underlying arrival rate γ_s is assumed to remain constant throughout the entire season, but the resulting actual demand for product s may only be observed in the periods when that product is included in the assortment. In addition, the arrival processes corresponding to different products are assumed to be independent.

We adopt a standard Gamma-Poisson Bayesian learning mechanism. The underlying demand rate γ_s for each product s is initially unknown to the retailer; however, he starts each period with a prior belief on the value of that parameter represented by a Gamma distribution with shape parameter m_s and scale parameter α_s (m_s and α_s must be positive, and m_s is assumed to be integer). Redefining time units if necessary, we can assume with no loss of generality that the length of each assortment period is 1; the predictive demand distribution under that belief for product s in the upcoming period is then given by a negative binomial distribution with parameters m_s and $\alpha_s(\alpha_s + 1)^{-1}$. If now product s is included in the assortment and n_s actual sales are observed in that period, it follows from Bayes's rule that the posterior distribution of γ_s has a Gamma distribution with shape parameter $(m_s + n_s)$ and scale parameter $(\alpha_s + 1)$.

2.2. Dynamic Programming Formulation

Given the discrete and sequential character of our problem, the natural solution approach is dynamic programming (DP); the state at time t is given in our model by the parameter vector $\mathbf{I}^t = (\mathbf{m}, \boldsymbol{\alpha})$, which summarizes all relevant information, including past assortments and observed sales. For ease of notation, we omit the dependence of $\mathbf{m}$ and $\boldsymbol{\alpha}$ on t. In each period, the decision to include product s in the assortment or not can be represented by a binary variable $u_s \in \{0, 1\}$, where $u_s = 1$ means that product s is included.

The optimal profit-to-go function $J_t^*(\mathbf{m}, \boldsymbol{\alpha})$ given state $(\mathbf{m}, \boldsymbol{\alpha})$ and t remaining periods must then satisfy the following Bellman equation:

$$J_t^*(\mathbf{m}, \boldsymbol{\alpha}) = \max_{\substack{\mathbf{u}\in\{0,1\}^S:\\ \sum_{s=1}^S u_s \le N}} \sum_{s=1}^S r_s \frac{m_s}{\alpha_s} u_s + \mathbb{E}_{\mathbf{n}}[J_{t-1}^*(\mathbf{m}+\mathbf{n}\cdot\mathbf{u}, \boldsymbol{\alpha}+\mathbf{u})],$$

where $\mathbf{v}\cdot\mathbf{u}$ represents the componentwise product of two vectors, and the terminal condition is $J_0^*(\mathbf{m}, \boldsymbol{\alpha}) = 0$ for all states.

Note that the only link between consecutive periods in this model is the information acquired about demand, and that different products are only coupled at a given period through the shelf space constraint $\sum_{s=1}^S u_s \le N$; this type of problem is known as a *weakly coupled* DP.

3. Analysis

3.1. The Dual Dynamic Program

The analysis of the model is based on Lagrangian relaxation and the decomposition of weakly coupled dynamic programs (see, for instance, Bertsimas and Mersereau 2004 and the references therein). Specifically, we relax the shelf space constraint, which leads to the definition of *dual policies* that can be shown to be useful in finding near-optimal *primal* policies and upper bounds for the optimal profit-to-go. Let $\lambda_t(\mathbf{m}, \boldsymbol{\alpha})$ denote any function associated with period t that maps the state space into the set of nonnegative real values; we define a *dual policy* to be a vector of functions $\boldsymbol{\lambda}_t = (\lambda_t(\cdot), \lambda_{t-1}(\cdot), \ldots, \lambda_1(\cdot))$.

For any dual policy $\boldsymbol{\lambda}_\mathbf{t}$ and any initial state $(\mathbf{m}, \boldsymbol{\alpha})$, the corresponding profit-to-go is obtained by solving the *dual dynamic program* given by:

$$H_t^{\boldsymbol{\lambda}_\mathbf{t}}(\mathbf{m}, \boldsymbol{\alpha}) = N\lambda_t(\mathbf{m}, \boldsymbol{\alpha}) + \max_{\mathbf{u}\in\{0,1\}^S} \sum_{s=1}^S \Big(r_s \frac{m_s}{\alpha_s} - \lambda_t(\mathbf{m}, \boldsymbol{\alpha})\Big) u_s + \mathbb{E}_{\mathbf{n}}[H_{t-1}^{\boldsymbol{\lambda}_{\mathbf{t-1}}}(\mathbf{m}+\mathbf{n}\cdot\mathbf{u}, \boldsymbol{\alpha}+\mathbf{u})],$$

with $H_0^{\boldsymbol{\lambda}_\mathbf{0}}(\mathbf{m}, \boldsymbol{\alpha}) = 0 \;\forall\, (\mathbf{m}, \boldsymbol{\alpha})$.

In words, a dual policy gives the price of a unit of shelf space for each period and each possible state. As expected, weak duality holds, and for any dual policy and initial state we have that $J_t^*(\mathbf{m}, \boldsymbol{\alpha}) \le H_t^{\boldsymbol{\lambda}_t}(\mathbf{m}, \boldsymbol{\alpha})$. By considering *open-loop* dual policies (i.e., a constant shadow price per period), one can calculate an upper bound for $J_t^*(\mathbf{m}, \boldsymbol{\alpha})$ using standard convex nondifferentiable optimization methods.

3.2. The Index Policy

It is well known that index policies are not optimal for our version of the multiarmed bandit problem (see Berry and Fristedt 1985), however they are still appealing given their simple structure. Through a sequence of intuitive approximations to the dual DP we derive a heuristic index policy for the dynamic assortment problem. The suggested rule is to include the N products with the highest indices in the assortment, where the index for product s at period t is given by the following formula:

$$\eta_{t,s} \approx r_s \mathbb{E}[\gamma_s] + z_t \frac{r_s \mathbb{V}[\gamma_s]}{\sqrt{\mathbb{V}[\gamma_s] + \mathbb{E}[\gamma_s]}}. \tag{1}$$

The factor z_t is the unique solution to the equation $(t-1)\cdot\Psi(z_t) = z_t$, where $\Psi(z) = \int_z^\infty (x-z)\phi(x)\,dx$ is the loss function of a standard normal. The values z_t, which are independent of the problem data, are increasing and concave in t.

The index $\eta_{t,s}$ represents the highest price at which one should be willing to rent some shelf space to display (and sell) product s there; it is thus a measure of the desirability of including each individual product in the assortment, and from that standpoint the rationale behind the suggested index policy is to fill all shelf space with the most desirable products. Note that the first term in the index expression (1) favors exploitation, and the second term favors exploration because it is increasing in both the variance of γ_s and the number of remaining periods (through z_t). Intuitively, when uncertainty about demand for a product s (captured by $\mathbb{V}[\gamma_s]$) is high, there is more benefit to learn from including s in the assortment because of the upside potential from future sales. However, one should increasingly favor exploitation over exploration as the remaining planning horizon (and opportunity for leveraging exploration) shortens, which is captured by the decrease with t of the multiplicative factor z_t. The summation of a variance and an expectation in the second term of (1) is not a mistake, but rather a consequence of the period length being equal to 1.

Finally, when assessing the performance of the index policy defined above, our primary benchmark is the *greedy policy*, which consists of selecting in each period the N products with the highest immediate expected profit $r_s\mathbb{E}[\gamma_s]$. Note that the greedy policy still involves learning despite its myopic nature, but the impact of assortment decisions on future learning is ignored. As a result, several authors also refer to it as *passive learning*.

4. Conclusions

We have developed a discrete-time DP model for the dynamic assortment problem faced by a fast-fashion retailer refining his estimate of consumer demand for his products over time. The main assumptions made were: (i) independent products, (ii) no lost sales, (iii) constant demand rates, and (iv) immediate assortment implementation. Under these assumptions, we have formulated this problem as a multiarmed bandit with finite horizon and multiple plays per stage. Using the Lagrangian decomposition of weakly coupled DPs, we have derived a closed-form index policy and we have derived an upper bound for the optimal profit-to-go, which allows us to assess the suboptimality gap of the suggested policy.

A simulation study indicates that the index policy always performs at least as well as the greedy policy (or passive learning), and significantly outperforms it in scenarios with diffuse or biased prior demand information. Also, numerical computations of the bound mentioned above suggest that the index policy is close to optimal. In general, the improvement of the suggested index policy on the greedy rule increases with the planning horizon length and the variance of the initial priors.

Although the major assumptions of our model may be particularly strong in some environments, our approach was partly motivated by the belief that the closed-form policy that they allow to derive constitutes a useful starting point for designing heuristics or developing extensions in more complex environments.

References

Berry, D. A., B. Fristedt. 1985. *Bandit Problems, Sequential Allocation of Experiments*. Chapman and Hall, New York.

Bertsimas, D., A. Mersereau. 2004. A learning approach to customized marketing. Working paper, University of Chicago, Chicago, IL.

Fisher, M. L., A. Raman, A. S. McClelland. 2000. Rocket science retailing is almost here—Are you ready? *Harvard Bus. Rev.* **78**(4) 115–124.

Promised Leadtime Contracts and Renegotiation Incentives Under Asymmetric Information

Holly S. Lutze
School of Management, University of Texas at Dallas, Richardson, Texas 75083-0688, hlutze@utdallas.edu

Advisor: Özalp Özer
Department of Management Science and Engineering, Stanford University, Stanford, California 94305-4026

Firms often establish supply chain relationships through contracts that provide rules for interaction. These contracts help align incentives for decision making and establish how partners will share both the benefits of interaction and the risks from uncertain supply or demand. We propose a new multiperiod contract form, the promised leadtime contract. The contract reduces supplier risk from future demand uncertainty, and it eliminates buyer risk from uncertain inventory availability. The supplier agrees to ship buyer orders in full after a promised leadtime, and the buyer pays the supplier for this privilege. The supplier and buyer may each carry inventory, depending on the agreed on promised leadtime and their respective production and processing leadtimes. The buyer satisfies consumer demand through on-hand inventory.

Promised leadtime contracts permit us to blend periodic-review inventory control theory with contracting for optimal inventory risk sharing between two firms. Using adverse selection theory, we also address the challenge of asymmetric information when designing an optimal promised leadtime contract.

Our study of promised leadtime contracts appears timely. According to *Automotive Industries*, the senior vice president of major automotive supplier Federal Mogul Corporation reports, "The key to supply chain is time" (Haight 2003). Having upstream supply chain partners that can deliver on time every time is a primary concern for many firms (Moyer and Burnson 2003, Hauser 2003). Analytics engineer Charles Ng of Vivecon Corporation, a company offering products and services for supply chain risk management, reports that contracts with leadtimes are widely used in industry, perhaps 90% of the time (Ng 2004).

A promised leadtime contract consists of a *promised leadtime* and a *corresponding payment*. The supplier guarantees that she will ship buyer orders in full after the promised leadtime. To do so, if necessary, the supplier acquires emergency units from an alternative source and pays a penalty each period until she resupplies this source. In exchange for the promised leadtime guarantee, the buyer pays the supplier the corresponding payment in each period.

Both the supplier and the buyer under a promised leadtime contract face a periodic review, finite-horizon inventory control problem with stationary unit costs for inventory ordering, holding, and backordering. Consumer demand is nonnegative and has a known stationary distribution with logconcave density (Bagnoli and Bergstrom 2005). Due to demand uncertainty and positive leadtimes, both the supplier and the buyer carry inventory. The results of Veinott (1965) ensure that a myopic base stock policy is optimal for the periodic review, stationary, finite-horizon inventory control problem of each firm. Hence, we model the impact on a firm of a promised leadtime contract with the expected inventory cost per period under the optimal base-stock policy. Using these expected inventory costs, the supplier chooses the optimal multiperiod promised leadtime contract.

A promised leadtime shifts responsibility for demand uncertainty from the supplier to the buyer. When the promised leadtime is zero, the buyer demands immediate shipment of all orders, so the supplier holds inventory to satisfy uncertain demand. As the promised leadtime increases, the supplier receives buyer order information earlier and is able to hold relatively less inventory in anticipation of demand. On the other hand, the buyer must place orders that anticipate demand further into the future. Promised leadtimes exceeding total supplier production time suggest a make-to-order supplier operation.

In our model, the supplier proposes a promised leadtime contract to the buyer. Based on budget restrictions, the buyer has a maximum acceptable expected inventory cost per period, his expected inventory cost under his existing procurement strategy. Given supply chain expected inventory cost functions, the supplier proposes the acceptable promised leadtime contract that minimizes her own expected inventory cost.

The buyer's optimal service level, the probability that he does not experience a stockout during his replenishment leadtime, poses an interesting problem for the supplier. A buyer satisfying consumers in a high-service market incurs higher inventory-holding and shortage costs. Shorter promised leadtimes reduce buyer uncertainty when ordering from the supplier and, hence, the amount of inventory necessary to provide the same service level to consumers. We demonstrate that the buyer's minimum expected inventory cost in each period increases with both promised leadtime and service level. In contrast, the supplier's minimum expected inventory cost decreases when she has more time to respond to an order from the buyer.

When the buyer's service level is common knowledge, the supplier offers a single first-best contract that is efficient for the supply chain and, hence, generates the optimal inventory risk-sharing strategy. Given the efficient promised leadtime, the supplier then offers the buyer his highest acceptable corresponding payment while retaining all supply chain surplus from the promised leadtime contract. The supplier shoulders more inventory risk when doing business with a high-service buyer, so a high-service buyer pays more to shift this additional inventory risk to the supplier.

While the supplier may know the market-specific factors influencing the buyer's optimal service level, she is unlikely to know confidential company-specific factors related to competitive strategy, position in the marketplace, and assessment of inventory risks. Asymmetric service-level information means the supplier does not know the value to the buyer of a shorter promised leadtime and, hence, cannot determine the first-best promised leadtime contract. If asked for this information, the buyer has an incentive to exaggerate his service level, thereby shortening the promised leadtime and reducing his expected inventory cost per period under the first-best contract.

To produce optimal supply chain performance *and* minimize her share of inventory risk, the supplier in our model designs a mechanism composed of different contracts for different possible buyer service levels. At minimum possible cost to the supplier, the mechanism creates an incentive for the buyer to choose the contract designed for his true service level, thus revealing this information. Assuming the supplier holds prior beliefs about the likelihood that the buyer provides a particular service level, we prove the *single-crossing property* for the buyer's expected inventory cost and characterize the properties of the optimal contract mechanism as an application of adverse selection theory (Salanie 1997).

Under asymmetric service-level information, only the promised leadtime for the lowest possible buyer service level is necessarily efficient. The optimal promised leadtime for a higher buyer service level is shorter than the first-best promised leadtime. A supplier without full information offers information rent for all but the highest buyer service level. In our model, information rent is a reduction in the buyer's share of inventory-related cost to encourage the buyer to truthfully report his service level. Because the buyer's expected inventory cost increases with service level, his information rent (his necessary encouragement) decreases, and the highest possible service level yields his maximum acceptable cost. As the likelihood of a high buyer service level increases, the expected loss to information rent decreases.

We show that the supplier may not offer a promised leadtime contract acceptable to a high-service buyer when efficiency loss and expected information rent from doing so increases her expected inventory cost per period. This policy, known as a *shut-down solution* (Laffont and Martimort 2002), essentially means that high-service buyers represent a market segment that is not profitable for the supplier to pursue, so she concentrates on more profitable low-service markets.

As long as noncontracted parameters of the business environment do not change, a multiperiod promised leadtime contract is *renegotiation-proof* (Laffont and Tirole 1990). However, if any parameter impacting inventory costs changes, the supply chain may benefit from renegotiating the contract before it expires. We assume that the supplier and buyer agree to Pareto-improving contract renegotiation when both firms benefit. When the buyer reduces his processing leadtime before a multiperiod contract expires, we show that the supplier must provide incentives for the buyer to truthfully report the change and agree to an optimal renegotiated contract. We determine optimal renegotiation incentives that take the form of transfer payments.

References

Bagnoli, M., T. Bergstrom. 2005. Log-concave probability and its applications. *Econom. Theory* **26**(2) 445–469.

Haight, B. 2003. The perils of supply chain management. *Automotive Indust.* (August).

Hauser, L. M. 2003. Risk-adjusted supply chain management. *Supply Chain Management Rev.* **7**(6) 64–71.

Laffont, J.-J., D. Martimort. 2002. *The Theory of Incentives: The Principal-Agent Model.* Princeton University Press, Princeton, NJ.

Laffont, J.-J., J. Tirole. 1990. Adverse selection and renegotiation in procurement. *Rev. Econom. Stud.* **57** 597–625.

Moyer, L., P. Burnson. 2003. Pepsi-Penske partnership satisfies consumer demand. *World Trade* **16**(9) 42–44.

Ng, C. 2004. Personal communication. (November 30).

Salanie, B. 1997. *The Economics of Contracts: A Primer.* MIT Press, Cambridge, MA.

Veinott, A. 1965. Optimal policy for a multi-product, dynamic, nonstationary inventory problem. *Management Sci.* **12**(3) 206–222.