This article was downloaded by: [216.73.216.192] On: 13 July 2026, At: 17:43
Publisher: Institute for Operations Research and the Management Sciences (INFORMS)
INFORMS is located in Maryland, USA

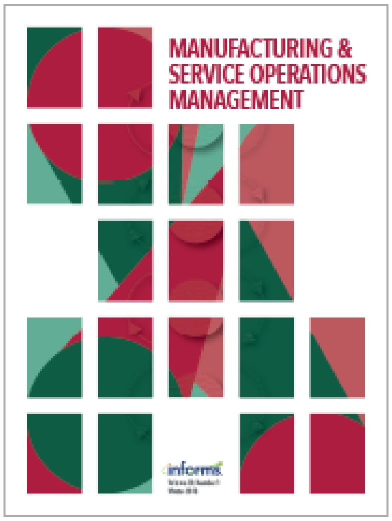

Manufacturing & Service Operations Management

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

The MSOM Society Student Paper Competition: Extended Abstracts of 2000 Winners

Jérémie Gallien, Serguei Netessine, Maher Lahmar, Hakan Ergan, Sergei V. Savin, Michael J. Fry,

To cite this article:
Jérémie Gallien, Serguei Netessine, Maher Lahmar, Hakan Ergan, Sergei V. Savin, Michael J. Fry, (2001) The MSOM Society Student Paper Competition: Extended Abstracts of 2000 Winners. Manufacturing & Service Operations Management 3(1):1-1. https://doi.org/10.1287/msom.3.1.1

Full terms and conditions of use: https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

© 2001 INFORMS

Please scroll down for article—it is on subsequent pages

With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.
For more information on INFORMS, its publications, membership, or meetings visit http://www.informs.org

Downloaded from informs.org by [216.73.216.192] on 13 July 2026, at 17:43 . For personal use only, all rights reserved.

The MSOM Society Student Paper Competition: Extended Abstracts of 2000 Winners

Each year the Manufacturing and Service Operations Management (MSOM) Society of INFORMS conducts a student paper competition. In Volume 2, Issue 2 of *M&SOM,* we published the extended abstracts of the 1999 winners in the hopes that this could become an annual event. Our hopes have become a reality.

Sridhar Seshadri and Ravi Anupindi, New York University, cochaired the 2000 competition. The judges for the final round were Garrett van Ryzin, Columbia University; Yehuda Bassok, University of Southern California; and Michael Pinedo, New York University.

The first-prize winner received $400, while second prize received $200. Prizes were awarded at the INFORMS meeting in San Antonio, Texas in November 2000. All finalists received a $50.00 discount coupon redeemable at an ''e-tailer.''

The winners and their faculty mentors were:

First Prize
Jérémie Gallien, Massachusetts Institute of Technology
Faculty Mentor: Lawrence M. Wein, MIT
''Design and Analysis of a Smart Market for Industrial Procurement''

Second Prize
Serguei Netessine, University of Rochester
Faculty Mentors: Gregory Dobson and Robert A. Shumsky, University of Rochester
''Flexible Service Capacity: Optimal Investment and the Impact of Demand Correlation''

M&SOM would like to congratulate the winners as well as all five finalists. The extended abstracts of these papers follow.

1523-4614/01/0301/0001$05.00
1526-5498 electronic ISSN

Downloaded from informs.org by [216.73.216.192] on 13 July 2026, at 17:43 . For personal use only, all rights reserved.

Design and Analysis of a Smart Market for Industrial Procurement

Jérémie Gallien
Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139 and Neoptis, Inc., New York, New York 10025

Faculty Mentor: Lawrence M. Wein
Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

Online business-to-business auctions are becoming increasingly widespread because their demand properties are superior to that of fixed-priced mechanisms, while the Internet keeps their transaction costs low. However, online implementation of traditional auction mechanisms seems poorly adapted to the selection of industrial suppliers in procurement markets, as this process typically involves capacity constraints, transportation costs, incumbent supplier switching costs, quality requirements, etc. Motivated by various sources suggesting that the impact of suppliers' capacity constraints on the bidding mechanism used by existing procurement auction service providers such as FreeMarkets, Inc. may be particularly significant (Kinney 1999, Jahnke 1998, see also Rangan 1999 for background), our objective in this paper is to design and analyze a stylized online multi-item procurement auction mechanism adapted to supply environments with production capacity constraints.

The exact market environment we consider is a monopsony, where a buyer wants to purchase some specified quantity of several component types. Facing this opportunity, several suppliers are ready to compete on selling prices for the award of the corresponding procurement contracts. In addition to the buyer's quantity requirements and her low-cost objective, the other consideration relevant to the award decision is that each supplier has a limited production capacity across all component types, which we capture through a linear model. Likewise, the suppliers are assumed to have a separable, linear utility function based on individual linear production costs.

In this environment, our proposed allocation mechanism relies on the estimation by the auctioneer of these capacity constraints *before* the auction begins. The mechanism is then structured as a sequence of bidding rounds, where in each round each supplier bids on his unit selling prices for all component types simultaneously, and bidding is subjected to a nonreneging rule and a common-multiple rule. At the end of each round, the mechanism solves a linear program to compute the provisional quantity allocation of each component type to each supplier minimizing the buyer's total procurement cost, under both the quantity requirement constraints and the suppliers' capacity constraints. Suppliers are then informed about their own individual provisional allocation, to which they can respond at the next round by modifying their offer. This mechanism thus gives rise to a dynamic price formation process, which completes when a criteria based on bidding activity is met, at which point the latest provisional allocations become firm. To limit possible collusive behaviors, our mechanism does not communicate any competing bids to suppliers, nor the number, identities, and capacity constraints of their competitors. Using the information set described so far exclusively, potentially losing suppliers would lack sufficient information to decide on which component type to lower their bids, and by how much. As a result, they would most likely engage in exploratory bidding behaviors, hampering the efficiency of the price-formation process. To prevent this phenomenon, our mechanism gives each supplier the possibility to have his *myopic best response* (MBR) bid calculated at every round by a private bidding-suggestion device. More explicitly, this suggested bid to a given supplier is the one that maximizes his potential payoff function at the next round, as-

1523-4614/01/0301/0002$05.00
1526-5498 electronic ISSN

Downloaded from informs.org by [216.73.216.192] on 13 July 2026, at 17:43 . For personal use only, all rights reserved.

suming that all his competitors' bids remain unchanged.

The basic premise underlying our methodology is that bidders will be less than fully rational when confronted with the dynamic game of incomplete information just described. While this approach considerably simplifies the analysis, we share with other researchers (Samuelson 1997) the belief that it is also motivated by behavioral considerations, mainly the overwhelming complexity associated with the bidders' decision problems under any fully rational solution concept. It also seems to be gaining some popularity in the literature for the analysis of comparable problems (e.g., Demange et al. 1986, Wellman et al. 1999, Milgrom 1999, Parkes and Ungar 2000).

To first establish some convergence bounds, we only rely initially on the *weak behavioral assumption* that a supplier receiving, at some round, a private-allocation feedback of zero for all component types will lower at least one of his bids, unless doing so would make his corresponding margin nonpositive. This allows us to prove our first proposition that the highest selected bid on each component type at the end of the auction cannot be larger than some order statistic of the production costs for that component type (modulo the minimum bid decrement). This result generalizes to the case of a multi-item and optimization-based procurement auction the well known observation that in a single-item ascending English auction the price eventually obtained by the seller will be equal to the second highest valuation among bidders.

Although our first proposition is an ex post result (i.e., relying on an observation of the final allocation made a posteriori), it can be used to prove an ex ante result in the case where the suppliers have the same production technology: Our second proposition also provides an upper bound on the winning bids at the end of the auction, but it shows much more explicitly the dependency on the primary data of the problem. In particular, the larger the buyer's required quantities relative to the suppliers' available capacities, the lower the bound on the final winning bids; this result provides quantitative support for the intuition that the buyer's market power decreases as supply resources become more scarce.

The next step in our analysis is to propose and study the fully specified behavioral model, whereby every supplier uses the MBR calculation device at each round and follows exactly its bidding suggestions. Except for a technical exception, the MBR bidding model satisfies the weak behavioral assumption, so that the previous convergence bounds remain valid. Our third proposition establishes the validity of the MBR model as a solution-concept candidate by showing that the MBR bidding-profile sequence always converges in finite time, and its limit, entirely characterized by the initial bid and the problem data, is comparable to the notion of a *local Nash equilibrium* introduced by Bykowsky et al. (1995). While the MBR model bears some resemblance to the simple, best-response Cournot tatonnement process (Fudenberg and Levine 1998), we feel that it is much more realistic in this setting, mainly because of the lock-in effect induced by the nonreneging rule. More specifically, a supplier incorrectly assuming that the competition will be more aggressive at the next round would then most likely bid lower than necessary, thus wasting some margin and restricting the space of future possible actions.

We use the MBR solution concept to study the impact of initial bids; although our first two propositions suggest that the dependence of the final auction outcome on these bids should be relatively limited when the level of competition among suppliers is relatively high, this dependence is yet unclear in market environments characterized by supply resource scarcity. To address this issue, we turn to the analysis of the MBR bidding dynamics in the symmetric auction with two suppliers and two item types, which constitutes the simplest nontrivial market environment where our smart market mechanism can be applied. We also consider a scarce supply environment, where the buyer has no choice but to use some capacity from both suppliers. In this setting, our fourth proposition is that when the initial bids are sufficiently close to each other and are greater than the component-wise maximum of the production costs, they have relatively little impact on the final outcome; it also shows that the particular information structure used in our mechanism may still enable some amount

Downloaded from informs.org by [216.73.216.192] on 13 July 2026, at 17:43 . For personal use only, all rights reserved.

of bidding competition in this market. When the initial bids are further apart, we demonstrate in a particular example how premature equilibria may then arise and infer from these observations that whenever supply capacity is suspected to be scarce, it seems in the buyer's interest to have initial bids as close to each other as possible, for example, by enforcing a common maximum entry bid.

Finally, we briefly discuss some experiments (reported in Gallien 2000) linked to the incentive-compatibility issues involved in this type of market interaction. Although the impact of incentives linked to the suppliers' production costs input when using the MBR bid-suggestion device is a legitimate question in theory, numerical simulations tend to show that a truthful cost revelation is incentive compatible. Turning to capacity revelation, a first approach is to ignore the potential impact of a supplier's capacity input on the other suppliers' bidding behavior and to perform some sensitivity analysis on this input from the final equilibrium point. Interestingly, by complementarity slackness, a necessary condition for this marginal change to have any impact on this supplier's allocation is that his capacity constraint be binding. Therefore, this preliminary analysis suggests that the only potential suppliers' benefits, induced by overstating their capacity, should be associated with overloading and, thus, entail an observable supplier default. It turns out that our numerical experiments, taking into account the strategic effects, confirm this hypothesis. From this perspective, it seems that a procurement contract specifying an appropriately highly fixed penalty in case of misdelivery and/or quality problems should ensure a relatively truthful capacity input by the suppliers. We note that similar, albeit more informal, penalty costs already exist in practice in the form of reputation effects because a principal motivation for the design of our smart-market mechanism is the observed impact of overload fear on suppliers' bidding behavior in existing procurement auctions.

References

Bykowsky, M. M., R. J. Cull, J. Ledyard. 1995. Mutually destructive bidding: The FCC auction design problem. Working paper 916. California Institute of Technology.

Demange, G., D. Gale, M. Sotomayor. 1986. Mult-item auctions. *J. Polit. Econom.* **94** 863–872.

Fudenberg, D., D. K. Levine. 1998. *The Theory of Learning in Games.* The MIT Press, Cambridge, MA.

Gallien, J. 2000. Optimization-based auctions and stochastic assembly replenishment policies for industrial procurement. Ph.D. thesis. Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA.

Jahnke, A. August 1998. How bazaar-purchasing strategies. *CIO Web Business Magazine.*

Kinney, S. (FreeMarkets cofounder). 1999. Personal communication.

Milgrom, P. 1999. Putting auction theory to work: The simultaneous ascending auction. Dept. of Economics, Stanford University, Palo Alto, CA.

Parkes, D. C., L. H. Ungar. 2000. Iterative combinatorial auctions: Theory and practice. *Proc. 17th National Conf. Artificial Intelligence* (AAAI-00) 74–81.

Rangan, V. K. 1999. FreeMarkets Online. Harvard Business School Case 9-598-109.

Samuelson, L. 1997. *Evolutionary Games and Equilibrium Selection.* The MIT Press, Cambridge, MA.

Wellman, M. P., W. E. Walsh, P. R. Wurman, J. K. MacKie-Mason. 1999. Auction protocols for decentralized scheduling. EECS Dept., University of Michigan, Ann Arbor, MI.

Flexible Service Capacity: Optimal Investment and the Impact of Demand Correlation

Serguei Netessine
W. E. Simon Graduate School of Business, University of Rochester, Rochester, New York 14627

Faculty Mentors: Gregory Dobson • Robert A. Shumsky
W. E. Simon Graduate School of Business, University of Rochester, Rochester, New York 14627

Downloaded from informs.org by [216.73.216.192] on 13 July 2026, at 17:43 . For personal use only, all rights reserved.

Consider a telecommunications company that must construct daily service assignments for its hardware technicians and must also, in the long run, determine the number of technicians to hire and train. There are three major functions that the technicians perform. They are, in order of increasing complexity, equipment installation, testing, and repair. Technicians from the repair department are able to handle the other two jobs, those in the testing department can also accomplish installation, and members of the installation department are not yet cross-trained in any other job. When constructing the daily schedule, if the company runs out of technicians in one department, it may try to borrow a specialist with expertise at the next highest level. Due to union rules and employee preferences, borrowing from a department with skills *two* levels higher may not be feasible. In the long run, this limited flexibility must be taken into account when specifying the number of technicians needed to staff each department.

This paper considers such environments, in which customers may be upgraded to a higher level of service at no cost to the customer. We consider the short-term assignment problem when actual demand is known, as well as the long-term capacity decision in the face of uncertain demand. This model applies to problems in which customers may not be aware of the upgrade (as in the hardware technician example above) as well as service environments in which the customer experiences the upgrade directly. Examples include commercial aviation (economy to business-class upgrades), time-shared executive jets (the use of larger, faster jets to substitute for smaller, slower aircraft), and the car-rental industry (customers who request a compact car may be offered a mid-sized car).

We formulate the problem as a one-period model. The firm must commit to resource capacities at the beginning of the period, when demand is uncertain but a distribution is known. During the period the firm must assign available resources to realized demand. We find that significant benefits can be gained if the possibility of service substitution is accounted for at the time of capacity planning, rather than only at the time-of-service delivery. We also find that the correlation structure of demand has a significant impact on the value of resource flexibility.

Similar problems have been examined by many authors in the literature on both inventory management and capacity investment. The papers closest to this one are van Mieghem (1998) and Bassok et al. (1999). The most important difference between our model and this related literature is in the description of capacity flexibility. Van Mieghem's formulation can be adapted to an arbitrary substitution structure, while Bassok et al. allow ''full downward substitution,'' i.e., a customer can be served by any capacity of a higher, more expensive class. Here we assume that substitution by only one level is profitable. In many service environments this is a reasonable approximation of reality, for substitution by more than one class may

Downloaded from informs.org by [216.73.216.192] on 13 July 2026, at 17:43 . For personal use only, all rights reserved.

require a prohibitively expensive resource, such as a luxury car, to a service generating little revenue, as from a customer only willing to pay for a subcompact. Our assumption enables us to find simple expressions for the first-order optimality conditions, to make explicit comparisons among solutions with and without considering flexibility, and to derive efficient methods for calculating the optimal solutions for problems with large numbers of products and services. With a more general structure it is difficult to find solutions to problems with more than a few products.

Our single-upgrade structure, as well as the assumption that demand follows any elliptically contoured distribution (such as the multivariate Normal and t-distribution), allows us to characterize the effect of changes in demand correlation on the optimal solution. To our knowledge, there are no previous analytical results describing these effects on flexible capacity decisions. Eppen (1979) discusses the impact of correlation on the centralization decision for a single-product inventory system. Fine and Freund (1990) and Van Mieghem (1998) discuss the value of flexible capacity in the presence of perfect correlation. Van Mieghem (1997) describes numerical evidence that the optimal level of dedicated (flexible) capacity increases (decreases) as correlation increases.

Our characterization of the impact of demand correlation on the optimal solution invokes a result known in probability theory as Slepian's inequality (Tong 1980). Our primary analytical results demonstrate that when there are just two types of service, an increase in correlation will indeed lead to a shift from flexible to dedicated capacity. Specifically, we find that as correlation rises, then the rate of increase for specialized capacity is larger than the rate of increase for flexible capacity. Therefore, the optimal capacities for both types of capacities may rise as correlation increases, but the specialized capacity will rise faster than the flexible capacity. Likewise, both types of capacity may fall with increasing correlation, but the specialized capacity will fall more slowly than the flexible capacity. This phenomenon has not been described in the literature. By specifying appropriate ranges of demand parameters, we are also able to characterize the situations where both types of capacity will not either rise or fall simultaneously.

When there are three or more classes, there are also adjustments in the resources not directly affected by the correlation change. The adjustments follow an alternating pattern, with every other resource rising and the remaining resources falling, and this pattern holds for any demand distribution. Moreover, the absolute value of adjustments is diminishing for capacities farther away from the resources directly affected by the correlation change. Therefore, our work analytically characterizes the complex impact of demand correlation on the optimal level of capacity investment.

References

Bassok, Y., R. Anupindi, R. Akella. 1999. Single-period multiproduct inventory models with substitution. *Oper. Res.* **47** 632–642.

Eppen, G. D. 1979. Effects of centralization on expected costs in a multi-location newsboy problem. *Management Sci.* **25** 498–501.

Fine, C. H., R. M. Freund. 1990. Optimal investment in product-flexible manufacturing capacity. *Management Sci.* **36** 449–466.

Tong, Y. L. 1980. *Probability Inequalities in Multivariate Distributions.* Academic Press.

Van Mieghem, J. A. 1997. Investment strategies for flexible resources. Technical report. Center for Mathematical Studies in Economics and Management Science, Northwestern University.

Van Mieghem, J. A. 1998. Investment strategies for flexible resources. *Management Sci.* **44** 1071–1078.

Resequencing and Feature Assignment on a Moving Assembly Line

Maher Lahmar • Hakan Ergan
Department of Mechanical Engineering, University of Minnesota, Minneapolis, Minnesota 55455

Faculty Mentor: Saifallah Benjaafar
Department of Mechanical Engineering, University of Minnesota, Minneapolis, Minnesota 55455

Downloaded from informs.org by [216.73.216.192] on 13 July 2026, at 17:43 . For personal use only, all rights reserved.

We consider the problem of resequencing a prearranged set of jobs on a moving assembly line with the objective of minimizing changeover costs. A changeover cost is incurred whenever two consecutive jobs do not share the same feature. Features are assigned from a job-specific set of feasible features. Resequencing is limited by the availability of offline buffers, where a job removed from the line can be temporarily stored before it is reinserted. Because jobs are continuously moving, resequencing is limited to jobs that are upstream from the offline buffer. Because jobs may visit multiple departments before leaving the system, a sequence that is optimal for one department is usually suboptimal for most other departments. Consequently, there is a need to resequence jobs for each department.

Problems where jobs must be resequenced using offline buffers are frequently encountered in industries where a continuous material-transfer system is used and different products, or the same product with varying features, are produced on the same line. In our case, the problem is motivated by a vehicle-sequencing problem at a North American Ford truck assembly plant. In the plant, vehicles must be resequenced upon leaving the body shop and before entering the paint area. In the paint area, vehicles undergo a series of coating, painting, and drying operations. If two consecutive vehicles are painted different colors, a significant changeover cost is incurred because the current paint must be flushed-out—and disposed of—and paint nozzles must be thoroughly washed and cleaned with solvents. The changeover cost is primarily determined by the volume of paint that must be purged from an intermediate pipe that connects the paint nozzle to the paint tanks. Although the volume of paint that is purged is always the same (approximately one gallon), the cost of different paints can vary widely. This makes it desirable to form large blocks of the expensive colors, while possibly incurring the cost of more frequent color changeovers of the cheaper ones. Note that the cost of changing colors is not sequence-dependent (e.g., the cost of going from darker to lighter colors is not more expensive than going from lighter to darker colors). New paint technology has recently removed this dependency.

Flexibility in assigning colors to vehicles is limited by current demand requirement and available vehicle body type. Therefore, the set of feasible colors (features) for two consecutive vehicles can vary significantly. Because of space constraints and strict requirements on assembly line speed, there is currently little opportunity to minimize paint changeover costs by altering the sequence of vehicles. Therefore, changeover costs are being effected solely through color assignment. However, the plant is contemplating using a limited number of offline buffers, called *pull-off* tables, to allow for improved vehicle sequencing. Because pull-off tables are expensive to implement and operate, only a few can be introduced.

A pull-off table allows a vehicle to be removed from the line with little interruption to line flow and reinserted later. The line is designed to allow pulling and reinsertion at any point in the sequence without

Downloaded from informs.org by [216.73.216.192] on 13 July 2026, at 17:43 . For personal use only, all rights reserved.

significantly affecting line speed (pull-off tables have been originally designed to pull and reinsert vehicles that require rework). However, because pull-off tables are stationary, a pulled vehicle can only be inserted upstream from its initial position. Also, because a pull-off table serves as a temporary storage buffer for pulled vehicles, a pull-off table is unavailable to remove other vehicles until the current one has been reinserted.

The introduction of resequencing buffers to a moving assembly line poses a number of challenges. Decisions must be made regarding which vehicles must be pulled and where they should be reinserted. Simultaneously, decisions must be made regarding which color should be assigned to each vehicle, given the limited assignment flexibility available to each vehicle. A decision must be made at a strategic level regarding how many pull-off tables are desirable or affordable, and how the value of these tables is affected by various system parameters.

In this paper, we address several of these challenges. In particular, we present an integrated model for the joint problem of color assignment and vehicle resequencing, where the objective is to minimize color changeover costs. Because resequencing and color assignment must be carried out for successive blocks of vehicles, we present an efficient solution approach that is amenable to real-time implementation. We also examine the impact of multiple offline buffers on performance and study the relationship between resequencing ability and system parameters, such as block size and density of the color assignment matrix. Using numerical examples, we show that the effect of offline buffers is of the diminishing kind with most of the benefits of total resequencing flexibility achieved with only a few buffers. For the case of Ford, we show that significant savings (in excess of $1.2 million per paint line, per year) can be realized with the introduction of only a single buffer. We also show that the value of resequencing is dependent on the density of the color-assignment matrix, with resequencing having the greatest impact when density is in the middle range.

Despite the fact that there is a large body of literature on sequencing of paced assembly lines (see Baybars 1986 and Yano and Bolat 1989 for a general review), most of it assumes that the original sequence of jobs remains fixed. Therefore, the primary concern is determining a single job sequence for the entire line for an available (or recurring) set of jobs, where the objective typically is to balance workload among the different processing departments (Lee and Vairaktarakis 1997 and Yano and Rachamadugu 1991). Few papers, especially those that deal with mixed assembly lines, consider sequence-dependent setups. However, most of this literature assumes that there is full flexibility in how jobs are sequenced. It also assumes that setups are family or lot-specific, with family or lot membership being known. The problem we treat in this paper is different from this literature in two respects. First, we assume that there is only limited flexibility in how jobs can be resequenced. Second, we allow some flexibility in assigning attributes that determine family membership among jobs. To our knowledge, this is the first paper to consider the problem of where attribute assignment and scheduling are considered simultaneously. More significantly, this is the first paper, as far as we know, to consider the issue of sequencing with limited flexibility and job-dependent setups.

In general, the resequencing and color (feature) assignment problem is NP hard. For example, in the case where there is a sufficient number of pull-off tables to allow for full resequencing flexibility, the problem reduces to a set-covering problem. This means that solving the problem even with a small number of pull-off tables can be computationally demanding. Computational effort in our case is particularly a concern because decisions must be made in real time. The fact that vehicles are continuously moving, with approximately one-minute spacing, means that solutions must be provided in a few seconds.

Because resequencing and color assignment must be carried out for successive blocks of vehicles, we introduce a decomposition-based approach, where we successively solve a series of problems with a single pull-off table. For problems with a single pull-off table, we provide a polynomial-time dynamic-programming algorithm. The algorithm takes advantage of several structural properties that allow us to significantly re-

Downloaded from informs.org by [216.73.216.192] on 13 July 2026, at 17:43 . For personal use only, all rights reserved.

duce the solution space. We benchmark the solution we obtain against optimal solutions for solvable problems. We also develop upper and lower bounds on the solution and show that our approach yields reasonably good solutions. We show that in our application, the decomposition has a limited effect on the quality of the solution because pull-off tables are expensive to implement, and, at most, one or two can be accommodated on a single line. More importantly, we show that generally there is a diminishing effect to having multiple tables, and that most of the benefits of resequencing occurs with a single pull-off table.

References

Baybars, I. 1986. A survey of exact algorithms for the simple assembly line balancing problem. *Management Sci.* **32** 902–952.

Lee, C. Y., G. Vairaktarakis. 1997. Workforce planning in mixed model assembly systems. *Oper. Res.* **45** 553–567.

Yano, C. A., A. Bolat. 1989. Survey, development, and application of algorithms for sequencing paced assembly line. *J. Manufacturing Oper. Management* **2** 172–198.

——, R. Rachamadugu. 1991. Sequencing to minimize work overload in assembly lines with product options. *Management Sci.* **37** 572–586.

Capacity Management in Rental Businesses with Heterogeneous Customer Bases

Sergei V. Savin
Columbia Business School, Columbia University, New York, New York 10027

Faculty Mentor: Morris A. Cohen
The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104

Rental businesses in many sectors of the economy share some fundamental attributes. The rental company invests in the equipment for which there is a potential demand, and a stream of customers patronizes the company, renting its equipment. After each rental, the equipment is returned to the company, and rental durations are typically significantly shorter than the life of the equipment, so each unit may be used repeatedly.

For those who manage rental businesses, important managerial decisions focus on matching rental demand with the equipment supply. These decisions create a hierarchy of managerial levers at the company's disposal. Longer term, capital-investment decisions set the company's overall level of rental capacity and attempt to capture as much demand for rental services as is (marginally) profitable. While they provide for long-term matching between supply and demand, fleet-sizing decisions may not be used to counterbalance the short-term supply and demand mismatches, however. And, at any moment in time, tactical capacity-allocation decisions may be needed to determine which customers get served when rental capacity becomes scarce.

The two sets of decisions are interrelated, and this significantly complicates their analysis: good capacity-allocation policies depend on the overall level-of-service capacity, while the choice for the capacity level may be affected by the allocation policy employed. Thus, both sets of decisions should be analyzed si-

Downloaded from informs.org by [216.73.216.192] on 13 July 2026, at 17:43 . For personal use only, all rights reserved.

multaneously. Unfortunately, to the best of our knowledge, no existing research literature contains such joint analysis.

On the one hand, papers addressing the issue of capacity sizing (e.g., Heyman and Sobel 1984) assume that capacity is not rationed among customers; i.e., that every request for service is satisfied as long as there is capacity available. Such an approach, known as ''complete sharing'' of the service capacity, is also often observed in practice because rental company managers are typically encouraged to maximize the utilization of the available equipment. The structural properties of the optimal capacity decisions derived in these papers are based on the fact that revenues generated in these systems are concave functions of the service capacity.

On the other hand, a substantial number of papers on capacity allocation exists within the literature on yield management (for a recent comprehensive review see McGill and Van Ryzin 1999). Allocation of fixed capacity was studied in the contexts of the airline industry, the hotel business, and telecommunications. The issue of capacity, however, was not addressed in this line of research.

In this paper, we investigate the tactical allocation problem in which a fixed number of pieces of equipment must be rationed among arriving customers. Our approach follows in the spirit of the seat-allocation formulation of traditional airline yield-management problems: Given two classes of customers, when should arriving customers of each of the classes be allowed to rent equipment, and when would they be ''closed out?''

Two common assumptions made in traditional revenue-management models make them inadequate for our purposes: They assume that there exists a finite horizon over which units of capacity can be sold, and that each unit of capacity can be used only once. For example, in aviation, there are c seats on a flight, and once they are sold or the plane takes off they are not available for sale. While hotel problems could (and perhaps should) be formulated as rental problems, they are typically addressed as problems of allocating the rooms available on a single night.

In rental businesses, however, capacity becomes available when a customer's rental period ends. In turn, the time horizon may be thought of as occurring over multiple periods, and the problem is most naturally treated as a problem in dynamic and stochastic control. An arriving customer rents a unit, which becomes unavailable for the length of the person's rental. When the rental period ends, the unit becomes available again. Over any short period of time, the numbers of arriving and departing customers may be uncertain, and managers must develop effective policies for controlling the rental of system capacity.

Car rentals present the most common example of a business with such dynamics. Unfortunately, existing papers dealing with car rentals mainly focus on the descriptive features rather than on systematic analysis of the underlying system dynamics.

In this paper, we view the allocation of rental capacity as a continuous-time, infinite-horizon problem in which arrivals of customers and the durations of rentals are both uncertain. We formulate this problem as one of admission control to a multiple-server loss system. This control problem is a special case of the *stochastic knapsack* problem introduced in the telecommunications literature (Ross and Tsang 1989), one in which arriving ''objects'' (demands) are all of size one. This special structure allows us to characterize properties of the optimal admission control. Our analysis parallels that of Altman et al. (1998), who use dynamic-programming techniques to study optimal capacity-allocation rules and develop a fluid approximation of the problem, which they solve numerically. Given the particular structure of our problem, we are able to develop a number of additional useful structural results. Specifically, in this paper we make the following contributions:

1. We investigate the issue of the optimality for the complete sharing policy that is often used in practice. In particular, we derive two types of conditions that ensure the optimality of complete sharing: We show that not only it is optimal *not* to ration service capacity in the ''off-season,'' when the overall demand for service is low relative to capacity, but also in the ''peak season'' of high demand, as long as different customer classes are sufficiently similar.

2. In cases in which both the demand and the ser-

Downloaded from informs.org by [216.73.216.192] on 13 July 2026, at 17:43 . For personal use only, all rights reserved.

vice capacity are high, a fluid model may serve as a good approximation for the original stochastic-control problem. We derive closed-form expressions that characterize controls and performance obtained by ''aggregate threshold'' allocation policy of the fluid model. These expressions allow us to efficiently calculate admission thresholds that appear to perform well in the original stochastic model.

3. Closed-form expressions for the fluid model also allow us to demonstrate the concavity of the fluid model's revenues with respect to the fleet size when the trunk reservation policy is used. This concavity is the essential property required for the efficient solution of the related, long-term problem of capacity sizing.

Thus, in addition to characterizing the effect of fleet size on capacity allocation decisions, we lay the groundwork for the investigation of the effect of capacity allocation on fleet sizing decisions.

References

Altman, E., Jimenez, T., G. Koole. 1998. On optimal call admission control. Working paper.

Heyman, D. P., M. J. Sobel. 1984. *Stochastic Models in Operations Research.* McGraw-Hill **2** 391–407.

McGill, J. I., G. J. van Ryzin. 1999. Revenue management: Research overview and prospects. *Trans. Sci.* **33** (2) 233–256.

Ross, K. W., D. H. K. Tsang. 1989. The stochastic knapsack problem. *IEEE Trans. Comm.* **37** 740–747.

Coordinating Production and Delivery Under a (z, Z)-type Vendor-Managed Inventory Contract

Michael J. Fry
Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, Michigan 48109

Faculty Mentors: Roman Kapuscinski • Tava Lennon Olsen
Department of Operations Management, University of Michigan Business School, Ann Arbor, Michigan 48109
John M. Olin School of Business, Washington University in St. Louis, St. Louis, Missouri 63130

New information technologies and a continued increase in competition have led to the development of new interorganizational relationships in the supply chain. One such relationship that appears frequently in industry is vendor-managed inventory, or VMI. Under VMI, the supplier assumes management of the retailer's inventory, making such decisions as when and how much inventory to ship to the retailer. The supplier typically uses advanced information technology such as Electronic Data Interchange (EDI) or Internet technology to monitor the retailer's stock level and demand information. As these technologies continue to become cheaper and more common in industry, it is expected that programs such as VMI will also become more popular. Probably the best known VMI program is the partnership between Wal-Mart and Procter and Gamble, but other companies have implemented VMI as well including Campbell's Soup, Barilla SpA, Shell Chemical, General Electric, Intel, and many others.

Reports from industry provide differing accounts on the success and benefits attributable to VMI (Saccomano 1997), Schenck and McInerney (1998). Due to a lack of analytical models describing VMI and the fact that VMI appears in many different forms within industry, it has been very difficult to determine ex-

Downloaded from informs.org by [216.73.216.192] on 13 July 2026, at 17:43 . For personal use only, all rights reserved.

actly why VMI seems to work in some situations, but not in others. It is also well known that implementing VMI can be cumbersome and difficult as each participant must double check their inventories for accuracy and information systems must be made compatible across the supply chain. Therefore, it would be extremely helpful for firms considering VMI to be able to estimate the expected benefits before actual implementation of a VMI program. Our research attempts to model one particular form of VMI agreement that occurs within industry, which we call a (z, Z) VMI contract. We wish to identify if VMI provides opportunities for savings over traditional retailer-managed inventory (RMI) for the overall supply chain, and if so, how much of a savings can be realized.

We model a specific VMI agreement between a single retailer and a single supplier to determine the benefits that are offered by coordinating production and delivery under VMI. The supplier can ship to the retailer in any period, but can produce only once every T periods. This is appropriate for a setting in which a supplier may manufacture different goods on a set production schedule so that the supplier can only produce material for a given retailer once every T periods. Alternatively, we could imagine a supplier which, due to limited availability of necessary raw materials, can produce for a given retailer only once every T periods.

We further assume that the retailer (under RMI) and the supplier (under VMI) follow a periodic-review inventory policy so that they examine the retailer's inventory level at the beginning of each period and that end-user demand occurs only at the retailer. Demand is independently and identically distributed (i.i.d.) according to the distribution function $\Phi(\cdot)$ and density function $\phi(\cdot)$. Excess demand at the retailer is backlogged. We allow the supplier to outsource to obtain material to ship to the retailer in a period in which the supplier cannot produce. In this context outsourcing could represent a form of expediting such as working overtime, producing using less efficient technology, transshipping from another location, or procuring from an outside source.

Our model is based on several direct experiences we have had with VMI agreements in practice from the semiconductor industry, the electronics industry, and the furniture industry. While each of these VMI agreements is unique in regards to the specifics of the agreement, we found many common features across the agreements that helped us to formulate our model of VMI. In particular, we found that in each case the supplier's control over the retailer's inventory was not total, and that VMI was typically initiated by some sort of contract, or up-front agreement, which defined the limits of the supplier's control and any related performance measures. We found that the supplier's performance was commonly evaluated by the retailer using such measures as minimum and maximum inventory levels or inventory turns and end-user customer-service levels.

Our experiences with these firms led us to define a type of initiating VMI contract that we call a (z, Z) VMI contract. Under a (z, Z) VMI contract, the retailer sets a minimum inventory level, z, and a maximum inventory level, Z, that represent the lowest and highest inventory levels, respectively, that the retailer wishes his stock to experience when measured after customer demand. The supplier agrees to pay a penalty amount of b^- (b^+) per unit to the retailer for every unit of the retailer's inventory that is less than z (more than Z) after customer demand. This effectively serves to transfer some of the demand risk from the retailer to the supplier.

Using this model of VMI, we characterize the optimal policy of both the supplier and the retailer operating under a (z, Z) VMI contract. The supplier's decisions under VMI are composed of how much to ship to the retailer in each period (and associated with this, the decision of how much to outsource in each period), as well as the decision of how much to produce every T periods. The retailer is allowed to set the values of Z and z, as these values are abstractions of the typical requirements of such measures as minimum inventory turns and minimum customer service requirements. We further assume that the VMI contractual parameters are set through an up-front negotiation process, and we examine how robust contractual parameters may be chosen.

An adaptation of Clark and Scarf's work with mul-

Downloaded from informs.org by [216.73.216.192] on 13 July 2026, at 17:43 . For personal use only, all rights reserved.

tiechelon inventory models allows us to prove that, under VMI, it is optimal for the supplier to follow a replenish up-to policy. The supplier brings the retailer's inventory up to an amount that is dependent on both the amount of time remaining until the next production period and whether or not the supplier must outsource in the current period. We also show that it is optimal for the supplier to follow a production up-to policy for his/her production decisions. The retailer's decision of choosing Z and z is shown to be based on a newsvendor type of relation. It is also shown that given a value of $Z - z$, the supplier is indifferent to the specific choices of Z and z.

Once we have determined the optimal behavior of both the supplier and the retailer under VMI, we compare the performance of a supply chain operating under a VMI agreement to that of one operating under a more traditional RMI scenario. We do this by conducting a numerical study utilizing infinitesimal perturbation analysis (see Glasserman and Tayur (1995) for a discussion of infinitesimal perturbation analysis) and non-linear optimization techniques to quickly solve for the optimal replenishment up-to and production up-to amounts. The costs generated under both VMI and RMI are then found using simulation.

Our numerical analysis demonstrates that, in most ''reasonable'' scenarios, VMI results in lower total supply-chain costs than RMI, and often the costs are significantly lower. However, we also show that VMI can perform worse than RMI in certain scenarios if care is not taken in setting the upfront VMI contractual parameters. To help prevent this from happening, we investigate and recommend guidelines for choosing contractual parameters (b^- and b^+, as well as the value of $Z - z$), which encourage low supply-chain costs under VMI. We find that the savings from VMI increase as the cost of outsourcing increases and as the amount of variability present in the demand distribution increases.

Our analysis appears to match many claims of VMI implementation in industry that state that VMI can be quite detrimental to suppliers because we can show that in many situations where VMI performs poorly, it is the supplier who bears the brunt of the additional costs. However, under an appropriate choice of contractual parameters, the supplier benefits from a move to VMI. This, for example, validates claims made by a Harvard Business School videotape (HBS #695504) that states that the supplier may receive great benefit from VMI. Finding parameters where both the supplier and the retailer benefit from VMI is, in most cases, nontrivial. Thus, it may be appropriate to consider a fixed payment or other pricing scheme that more fairly distributes the savings from VMI to make VMI attractive to all parties involved.

References

Clark, A. J., H. Scarf. 1960. Optimal policies for a multi-echelon inventory problem. *Management Sci.* **6** (4) 475–490.

Glasserman, P., S. Tayur. 1995. Sensitivity analysis for base-stock levels inmultiechelon production-inventory systems. *Management Sci.* **41**(2) 263–281.

Harvard Business School Videotape #695504. Bose Corp.: The JIT II Program.

Saccomano, A. 1997. Vendor managed inventory: Making it work. *Traffic World.* **34** 250.

Schenck, J., J. McInerney. 1998. Applying vendor-managed inventory to the apparel industry. *Automatic I.D. News.* **14**(6) 36–38.