

Using Imperfect Advance Demand Information in Production-Inventory Systems with Multiple Customer Classes

Jean-Philippe Gayon

Laboratoire G-SCOP, INP Grenoble, 38031 Grenoble Cedex 1, France, jean-philippe.gayon@inpg.fr

Saif Benjaafar

Department of Mechanical Engineering, University of Minnesota,
Minneapolis, Minnesota 55419, saif@umn.edu

Francis de Véricourt

European School of Management and Technology, 10178 Berlin, Germany, devericourt@esmt.org

We consider a make-to-stock supplier that operates a production facility with limited capacity. The supplier receives orders from customers belonging to several demand classes. Some of the customer classes share advance demand information with the supplier by announcing their orders ahead of their due date. However, this advance demand information is not perfect because the customer may decide to order prior to or later than the expected due date or may decide to cancel the order altogether. Customer classes vary in their demand rates, expected due dates, cancellation probabilities, and shortage costs. The supplier must decide when to produce and, whenever an order becomes due, whether or not to satisfy it from on-hand inventory. Hence, the supplier is faced with a joint production-control and inventory-allocation problem. We formulate the problem as a Markov decision process and characterize the structure of the optimal policy. We show that the optimal production policy is a state-dependent base-stock policy with a base-stock level that is nondecreasing in the number of announced orders. We show that the optimal inventory-allocation policy is a state-dependent multilevel rationing policy, with the rationing level for each class nondecreasing in the number of announced orders (regardless of whether the class provides advance information). From numerical results, we obtain several insights into the value of advance demand information for both supplier and customers.

Key words: advance demand information; production-inventory systems; inventory rationing; make-to-stock queues; Markov decision process

History: Received: July 28, 2005; accepted: August 16, 2007. Published online in *Articles in Advance* February 1, 2008.

1. Introduction

Technologies such as the Internet, electronic data interchange, and radio frequency identification are making it increasingly possible for firms to share demand information with other members of their supply chains. Initiatives such as the inter-industry Consortium on Collaborative Planning, Forecasting and Replenishment are frameworks for participating companies to share demand forecasts and coordinate ordering decisions. Large retailers such as Wal-Mart have put in place sophisticated processes that enable them to share real-time inventory usage and point-of-sale data with thousands of their suppliers. Some manufacturers have begun to offer incentives to

encourage their customers to place orders in advance. For example, Dell has recently announced its “Intelligent Fulfillment” initiative under which it offers four delivery options: Premium (next-day delivery), Precision (delivery on a specific date), Standard (promised delivery within five days), and Value (longer delivery times); see Özer and Wei (2004). By offering lower pricing for longer delivery times, Dell induces some customers to place their orders early (in contrast to those that require next-day delivery, for example).

Large manufacturers such General Motors and Boeing share with their suppliers demand forecasts, production schedules, and even future design plans. These various initiatives point to the growing

realization that acquiring and providing information about future demand is beneficial. For suppliers, having information about future customer demand is believed to mitigate the need for inventory. For customers, providing advance information about future requirements is thought to improve the quality of service that customers receive from their suppliers.

Although advanced demand information (ADI) can assume several forms in practice, it typically reduces to having information, either perfect or imperfect, about the timing and quantity of future customer orders. If the information is perfect, customers place orders ahead of time in specified quantities to be delivered at specified due dates. If the information is imperfect, customers place orders ahead of time but provide only estimates of either the actual due dates or quantities. The realized due dates and quantities may therefore differ significantly from these initial estimates. In this paper, we model a setting with imperfect ADI, where customers always place unit orders but provide only an estimate of the due dates and have the option of canceling the order. Hence, ADI is imperfect in both quantity (it can be zero or one) and due date timing.

We are in part motivated by settings where ADI is provided by allowing a supplier to be informed about the internal operations of its customers and to use this information to deduce something about when customers will eventually place orders. For example, an aircraft manufacturer such as Boeing may inform one of its component suppliers each time it starts assembling a new plane (or each time it enters a particular stage of the assembly process). The component is not immediately needed and is required only at a later stage of the assembly process. The manufacturer does not accept early deliveries but wishes to have the component available as soon as it is needed in a *just-in-time* fashion. The supplier uses the information about when the initial assembly has started to estimate when it will need to make a delivery to the manufacturer. In making this estimate, the supplier uses its knowledge of the manufacturer's operations and available data from past interactions. However, this estimate is clearly imperfect, and the manufacturer (because of inherent variability in its own assembly process) may request the component sooner or later than the estimated due date.

Similar ADI scenarios arise elsewhere. van Doneslaar et al. (2001) describe a case study of how ADI

is shared between building constructors and building material suppliers, with building constructors informing suppliers about either initiation or progress of building projects. The suppliers use this information to estimate when their material (e.g., power cables) would be needed. This estimation is imperfect because progress on a building project can be highly variable and because building constructors may decide not to place an order after all (e.g., when a feature of the building is removed or modified).

Although the benefits of sharing ADI are perhaps intuitively clear, it is less clear how this ADI should be used to make production or inventory-allocation decisions. This is particularly true when the demand information is not perfect, is only partially available, or varies in quality. In this paper, we shed some light on this problem in the context of a single firm with a finite production capacity (producing a single item at a time) that serves as a supplier of a common product to multiple customer classes. Customers place orders continuously over time with rates that vary from class to class. Some customer classes provide ADI by announcing their orders before they are actually due (the announcement of orders can be implicit, as in the cases described in the previous paragraph). However, this information is not perfect, and customers may decide to cancel their orders. Customers may also request that their orders be fulfilled prior to or later than the announced expected due dates. Hence, the *demand lead time* (the time between when an order is announced and when it is requested) is a random variable. Both the cancellation probabilities and the distributions of demand lead times may depend on the customer class.

In response to customer orders, the supplier must decide on how much inventory to stock and when to replenish this stock by producing. An order that becomes due but that is not immediately satisfied from inventory is considered lost and incurs a lost sales cost. Because lost sales costs vary by customer class, it may be optimal to reject an order from certain classes to reserve inventory for future orders from more important classes (classes with higher lost sales cost). The supplier must decide on how best to ration available inventory among the different classes. Hence, the supplier is faced with a joint production and inventory-allocation problem that must be resolved each time the state of the system changes.

Using imperfect demand information in a production-inventory system with multiple classes raises several questions. How should the availability of imperfect demand information affect production decisions and the allocation of available inventory among the different classes? How valuable is imperfect demand information, and how is its value affected by system parameters, especially capacity? Is ADI equally valuable to the supplier and to the customers that provide it? Are there customers whose ADI is more valuable to the supplier than others?

In this paper, we provide answers to these and other related questions for production-inventory systems with Poisson demands and exponentially distributed production times and demand lead times. We formulate the problem as a Markov decision process (MDP) and characterize the structure of the optimal policy. We show that the optimal production policy is a state-dependent base-stock policy with a base-stock level $s(y_1, \dots, y_n)$, where y_i is the number of orders from class i ($i = 1, \dots, n$) that are announced but not yet due. The base-stock level $s(y_1, \dots, y_n)$ is nondecreasing in each of the state variables y_i and increasing by at most one with unit increases in y_i . We show that the optimal inventory-allocation policy is a multilevel rationing policy, with state-dependent rationing levels $r_i(y_1, \dots, y_n)$, where orders due from class i are fulfilled as long as inventory level is above its corresponding rationing level. We show that the rationing level for each class (regardless of whether the class provides ADI) is nondecreasing in the number of announced orders from all the classes. This means, perhaps unexpectedly, that orders from a class that provides ADI are rationed at a higher level, as more orders from that class are announced.

Based on numerical results, we obtain several insights into the value of imperfect ADI. In particular, we make the following observations.

- The benefit, in terms of cost reduction, to the supplier from using ADI can be significant. However, the relative cost reduction is sensitive (in some cases in a nonmonotonic fashion) to various operating parameters, including demand lead time, production capacity, and lost sales costs.
- The benefit, in terms of higher service levels, to the customers that provide ADI, can be insignificant, with the supplier using ADI in some cases to reduce

inventory costs at the expense of customer service levels.

- Customers could extract some (or all) of the value of ADI from the supplier by imposing higher lost sales penalties in exchange for ADI.
- It can be more beneficial to the supplier to have ADI on a class with a lower lost sales cost than on one with a higher cost.
- The relative benefit of ADI to the supplier does not exhibit diminishing returns as the fraction of customers providing ADI increases.
- The benefit the supplier derives from inventory rationing while ignoring ADI can be more significant than the benefit derived from ADI without rationing. When both ADI and rationing are used, the benefits are at least complementary.

The rest of the paper is organized as follows. In §2, we provide a brief summary of related literature. In §3, we describe our model and formulate the problem as an MDP. In §4, we characterize the structure of the optimal policy and describe several properties of the optimal policy. In §5, we provide numerical results, which we use to derive additional insights. In §6, we offer a summary of main contributions and discuss possible extensions.

2. Literature Review

Our work is related to two streams of literature, one dealing with *inventory control with ADI* and the other with *inventory rationing with multiple customer classes*. The literature dealing with ADI can be broadly classified into two categories based on the underlying supply process: systems with load-independent supply lead times and systems with load-dependent lead times. In the first case, inventory replenishment lead times are assumed to be independent of the number of outstanding orders. In the second case, replenishment lead times are affected by the number of outstanding orders because of limitations in production capacity and congestion at the production facility. For example, in systems where items are produced one unit at a time, the supply lead time (the time it takes to replenish inventory to a particular level) depends on how many orders are already in the queue. In both streams of literature, ADI is usually assumed to be perfect, with announced orders specifying exact due dates and exact quantities.

The literature dealing with ADI in systems with load-independent lead times is extensive (see Gallego and Ozer 2002 for a recent review). The majority of this literature considers inventory systems with periodic review, where ADI is in the form of demand placed in a period t , but not due until a future period $t + L$, and where the demand lead time L is a known constant. Recent examples include papers by Gallego and Ozer (2001), Chen (2001), and Hu et al. (2003).

Hariharan and Zipkin (1995) were the first to consider ADI in a system with continuous review. They show that, in a system where orders are announced L units of time before their due date, the optimal policy has the form of a base-stock policy. Moreover, they show that demand lead time can be used to directly offset supply lead time, so that having a demand lead time of L is equivalent to reducing supply lead time by the same amount. Schwarz et al. (1997) consider a problem where customers place orders ahead of time, but, with some probability, some customers cancel their orders at the time they become due. They show that the optimal policy is a state-dependent base-stock policy where the state is the vector of demand signals over the fixed demand lead time. There is also a body of related literature dealing with future demand information in the form of demand forecast updates. See, for example, papers by Graves et al. (1986), Heath and Jackson (1994), Güllü (1996), and Zhu and Thonemann (2004) and the references therein.

The literature dealing with ADI in systems with load-dependent lead times (to which our paper belongs) is less extensive. Buzacott and Shanthikumar (1994) consider an integrated production-inventory system with Poisson demand and exponential production times, where orders, as in Hariharan and Zipkin (1995), are announced a fixed L units of time in advance of their due date (ADI is perfect). They evaluate a class of policies that uses two parameters: a fixed base-stock level and a fixed release lead time. Karaesmen et al. (2002) consider a discrete time version of this model and examine the structure of the optimal policy. Karaesmen et al. (2004) evaluate the impact of capacity on the benefit of ADI and show that the benefits of ADI tend to diminish when capacity is tight. Benjaafar et al. (2007) consider a setup similar to ours but with a single customer class and backorders.

The literature on inventory rationing can be similarly classified based on the assumption regarding

the supply lead time, load dependent or load independent. Papers dealing with inventory rationing in systems with load-independent lead times include Topkis (1968), Nahmias and Demmy (1981), Cohen et al. (1988), Frank et al. (2003), and Deshpande et al. (2003) and the references therein. Ha (1997a) appears to be the first to consider rationing in the context of a production-inventory system with Poisson demand and exponentially distributed production times. For a system with N customer classes and lost sales, he shows that the optimal policy is of the threshold type, where orders from the lower-priority class are fulfilled as long as inventory is above a certain threshold level. Ha (1997b) shows that the same structure holds under backordering for a system with two customer classes, and de Véricourt et al. (2002) generalize this result to systems with N classes. Benjaafar et al. (2004) consider inventory rationing in a system with multiple products and multiple production facilities, and Benjaafar and Elhafsi (2006) consider inventory rationing in an assemble-to-order system. Finally, Tan et al. (2005) treat a problem with both rationing and imperfect ADI, but in a discrete time setting with two periods and two demand classes. In their case, imperfect ADI corresponds to a signal that can be used to update the distribution of future demand.

To the best of our knowledge, our paper is the first to model imperfect ADI in systems with multiple demand classes and to consider how ADI can be used to affect both production and inventory allocation decisions. By treating ADI in the context of a system with finite capacity where production takes place one unit at a time, our paper captures important interactions absent from models in which lead times are assumed to be independent of the current loading of the supply process.

3. Model Formulation

We consider a supplier that produces a single item at a single facility for n different classes of customers. Customers place orders continuously over time according to a Poisson process with rate λ_i for customer class i , $i = 1, \dots, n$. Some customer classes provide ADI by announcing their orders before they are actually due. However, this information is not perfect, and customers may decide to request that an order be fulfilled prior to or later than the *expected* due

date, or they may cancel the order altogether. We let L_i denote the demand lead time for customers of class i , where L_i describes the time between when an order is first announced and when it becomes either due or is cancelled. The demand lead time L_i is exponentially distributed with mean $E(L_i) = 1/\nu_i$. The probability that an order is cancelled is $1 - p_i$, $0 \leq p_i \leq 1$. Orders from customers who do not provide ADI are considered immediately due on arrival. We let $\mathcal{A} \subset \{1, \dots, n\}$ denote the set of indices of classes with ADI and $\mathcal{W} \subset \{1, \dots, n\}$ denote the set of indices of classes without ADI, where $\mathcal{A} \cup \mathcal{W} = \{1, \dots, n\}$, and $\mathcal{A} \cap \mathcal{W} = \emptyset$.

The supplier has the option of producing ahead of demand and placing inventory in stock. The supplier can produce at most one unit at a time. Production times are exponentially distributed with mean $1/\mu$. At any time, the supplier has the choice of either producing or not. If a unit is not currently being produced, this means deciding whether to initiate production of a new unit. If a unit is currently being produced, this means deciding whether to interrupt its production. If the production of a unit is interrupted, it can be resumed the next time production is initiated (because of the memoryless property of the exponential distribution, resuming production from where it was interrupted is equivalent to initiating it from scratch). We assume that there are no costs associated with interrupting production. This conforms to earlier treatment of production-inventory systems in the literature; see, for example, Ha (1997a, b).

When a customer order becomes due and is not cancelled, the supplier has the option of either satisfying it from on-hand inventory or rejecting it, in which case a lost sales cost c_i is incurred if the rejected order is from class i (this cost may correspond to the cost of expediting the order through other means, including scheduling overtime or satisfying it from a third-party supplier). If the order is cancelled by the customer, no additional cost is incurred. All customer classes are satisfied from the same common stock. If an order becomes due and there is no on-hand inventory, the order is automatically rejected and the lost sales cost is incurred.

In our model, we assume that demand is Poisson and both production times and demand lead times are exponentially distributed. These assumptions are made in part for mathematical tractability, as they

allow us to formulate the control problem as an MDP and enable us to characterize analytically the structure of the optimal policy. They are also useful in approximating the behavior of systems with high variability. The assumptions of Poisson demand and exponential processing times are consistent with previous treatments of production inventory systems; see, for example, Buzacott and Shanthikumar (1993), Ha (1997a), Zipkin (2000), and de Véricourt et al. (2002) among others. The exponential distribution is appropriate for modeling demand lead time as well, because demand lead time may be determined, as we discussed in the introduction, by processing times of operations at the customer level. For example, in the case of aircraft assembly, the time until certain tasks are completed can be highly variable, justifying the high coefficient variation of an exponential distribution. In §6, we discuss how these assumptions may be partially relaxed.

We denote by $X(t)$ the level of finished goods inventory at time t and by $h(X(t))$ the corresponding inventory holding cost per unit of time, where $X(t)$ belongs to \mathbb{N} , the set of nonnegative integers, and $h(X(t)) = h'X(t)$ is a linear function with rate $h' \geq 0$. We assume that holding cost is incurred only for finished goods inventory because we assume that value is added to raw material only on production completion. We also assume that there is an unlimited supply of raw material (or, alternatively, raw material can always be delivered in a just-in-time fashion) so that production can be initiated at any time. Both of these assumptions are consistent with treatments elsewhere in the literature. They also allow us to focus on holding and shortage costs associated with finished goods inventory. Without loss of generality, we assume that the lost sales costs are ordered such that $c_1 \geq \dots \geq c_n$.

For i in $\mathcal{A} \cup \mathcal{W}$, we denote by $Y_i(t)$ the number of orders of class i that have been announced but that are not due yet at time t . For i in \mathcal{W} , $Y_i(t) = 0$. In the following, we initially assume that the number of announced orders for class i stays bounded by a finite number $m_i < \infty$. Orders of type i that are announced when $y_i = m_i$ are rejected and incur the lost sales cost c_i . Note that although m_i is finite, it can be arbitrarily large. We make this assumption for two reasons: (1) It arises naturally in some settings where the supplier cannot accept more than m_i orders in advance (in practice, of course, the number of

announced orders can never be infinite), and (2) it allows us to formulate the problem as a continuous time MDP with finite transition rates. The latter offers technical advantages because we can transform, via *rate uniformization*, the continuous time problem into a discrete time problem, which in turn simplifies the analysis and makes it easier to characterize the structure of an optimal policy. The version of the problem with finite parameters m_i also serves as the basis for treating the case with infinite m_i . This case involves unbounded transition rates and would be difficult to analyze otherwise. We discuss this case later and show that the optimal cost of the bounded problem converges to the cost of the unbounded problem and that properties of the optimal policy are preserved for the unbounded case.

Let $\mathbf{Y}(t)$ be the vector of announced orders defined by $\mathbf{Y}(t) = (Y_1(t), \dots, Y_n(t))$, where $\mathbf{Y}(t)$ belongs to the set $\mathcal{Y} = \{0, 1, \dots, m_1\} \times \dots \times \{0, 1, \dots, m_n\}$, where $m_i = 0$ for i in \mathcal{W} . Then the variable $(X(t), \mathbf{Y}(t))$ exhaustively describes the state of the system and the state space is $\mathcal{S} = \mathbb{N} \times \mathcal{Y}$.

A control policy π specifies at each time instant if the supplier should produce or not and determines at time instants when orders become due if the order should be fulfilled from on-hand inventory—if there is any—or be rejected. We restrict the analysis to Markovian policies—it will be shown in the proof of Theorem 1 that the optimal policy belongs to this class. Let $\mathbf{a}^\pi(x, \mathbf{y}) = (a_0^\pi(x, \mathbf{y}), \dots, a_n^\pi(x, \mathbf{y}))$ be the control associated with a policy π , where $a_0^\pi(x, \mathbf{y})$ corresponds to the production action when the system is in state (x, \mathbf{y}) such that

$$a_0^\pi(x, \mathbf{y}) = \begin{cases} 1 & \text{if the action is to produce, and} \\ 0 & \text{if the action is not to produce,} \end{cases}$$

and $a_i^\pi(x, \mathbf{y})$, $1 \leq i \leq n$, corresponds to the inventory-allocation action when a class i order becomes due and the system is in state (x, \mathbf{y}) such that

$$a_i^\pi(x, \mathbf{y}) = \begin{cases} 1 & \text{if the action is to satisfy an order from class } i \\ & \text{(possible only when } x \geq 1), \text{ and} \\ 0 & \text{if the action is to reject an order from class } i. \end{cases}$$

Note that the decisions about production and inventory allocation are made simultaneously. When the system is in state (x, \mathbf{y}) , we decide both whether to produce and whether to accept a future order from class i , i in $\mathcal{A} \cup \mathcal{W}$, should one become due while the system is still in state (x, \mathbf{y}) . If we decide to produce, the state of the system does not instantaneously change. Instead, it remains the same until either production completes or an order becomes due and is satisfied from inventory, whichever happens first. This also means that an item is not included in inventory and is not available for allocation to a customer until its production is complete. An underlying assumption is that if production is initiated, it can always be preempted. Of course, this may occur only if the system changes while production has not been completed (i.e., one or more orders become due and are fulfilled from inventory).

Given the control policy π , we define $v^\pi(x, \mathbf{y})$ as the expected total discounted cost of the infinite horizon MDP associated with policy π and initial state (x, \mathbf{y}) where $\mathbf{y} = (y_1, \dots, y_n)$. Let α be the discount factor, $0 < \alpha < 1$, and $N_i(t)$ be the number of orders of class i that have not been satisfied up to time t . Then $v^\pi(x, \mathbf{y})$ is given by

$$v^\pi(x, \mathbf{y}) = E_{x, \mathbf{y}}^\pi \left[\int_0^{+\infty} e^{-\alpha t} h(X(t)) dt + \sum_{i \in \mathcal{A} \cup \mathcal{W}} \int_0^{+\infty} e^{-\alpha t} c_i dN_i(t) \right]. \quad (1)$$

We seek to find a policy π^* that minimizes the expected discounted cost. We introduce the uniform rate $\gamma = \mu + \sum_{i \in \mathcal{A} \cup \mathcal{W}} \lambda_i + \sum_{i \in \mathcal{A}} m_i \nu_i$ and, without loss of generality, rescale time by letting $\alpha + \gamma = 1$. This allows us to transform the continuous time decision process into an equivalent discrete time decision process (Lippman 1975). Then $v^*(x, \mathbf{y}) \equiv v^{\pi^*}(x, \mathbf{y})$ can be shown to satisfy for all $(x, \mathbf{y}) \in \mathbb{N} \times \mathcal{Y}$ the following optimality equations

$$v^*(x, \mathbf{y}) = T v^*(x, \mathbf{y}), \quad (2)$$

where the operator T is defined as

$$T v = h + \mu P v + \sum_{k \in \mathcal{W}} \lambda_k W_k v + \sum_{k \in \mathcal{A}} (\lambda_k A_k^1 v + \nu_k y_k [p_k A_k^2 v + (1 - p_k) A_k^3 v] + (m_k - y_k) \nu_k v),$$

and operators P , W_k , A_k^1 , A_k^2 , and A_k^3 are given by

$$\begin{aligned}
 Pv(x, \mathbf{y}) &= \min[v(x, \mathbf{y}), v(x+1, \mathbf{y})], \\
 W_k v(x, \mathbf{y}) &= \begin{cases} \min[v(x-1, \mathbf{y}), v(x, \mathbf{y}) + c_k] & \text{if } x > 0 \\ v(x, \mathbf{y}) + c_k & \text{if } x = 0, \end{cases} \\
 A_k^1 v(x, \mathbf{y}) &= \begin{cases} v(x, \mathbf{y} + \mathbf{e}_k) & \text{if } y_k < m_k \\ v(x, \mathbf{y}) + c_k & \text{if } y_k = m_k, \end{cases} \\
 A_k^2 v(x, \mathbf{y}) &= \begin{cases} \min[v(x-1, \mathbf{y} - \mathbf{e}_k), v(x, \mathbf{y} - \mathbf{e}_k) + c_k] & \text{if } x > 0 \text{ and } y_k > 0 \\ v(x, \mathbf{y} - \mathbf{e}_k) + c_k & \text{if } x = 0 \text{ and } y_k > 0 \\ 0 & \text{if } y_k = 0, \end{cases} \\
 A_k^3 v(x, \mathbf{y}) &= \begin{cases} v(x, \mathbf{y} - \mathbf{e}_k) & \text{if } y_k > 0 \\ 0 & \text{if } y_k = 0, \end{cases}
 \end{aligned}$$

with \mathbf{e}_k denoting the k th unit vector of dimension n (e.g., \mathbf{e}_2 is the vector $(0, 1, 0, \dots, 0)$). Operator P is associated with the optimal control of production, W_k with the optimal allocation of inventory to a demand from class $k \in \mathcal{W}$, A_k^1 with the announcement of an order from class k where $k \in \mathcal{A}$, A_k^2 with the optimal allocation of inventory to a demand from class $k \in \mathcal{A}$, and A_k^3 with the cancellation of an order from class $k \in \mathcal{A}$. Moreover, a deterministic stationary policy that specifies for each (x, \mathbf{y}) an action that attains the minimum on the right-hand side of (2) is optimal, including among history-dependent (see Theorem 5.5.3b in Puterman 1994) and randomized Markov policies (see Proposition 6.2.1 in Puterman 1994).

4. Characterizing the Optimal Policy

We investigate the structure of the optimal policy by identifying a set of structured value functions that is preserved under the optimal operator T . The following definition introduces this set. First, we define the operators Δ_0 , Δ_i , and Δ_{0+i} , where $\Delta_0 v(x, \mathbf{y}) = v(x+1, \mathbf{y}) - v(x, \mathbf{y})$, $\Delta_i v(x, \mathbf{y}) = v(x, \mathbf{y} + \mathbf{e}_i) - v(x, \mathbf{y})$, $\Delta_{0+i} v(x, \mathbf{y}) = v(x+1, \mathbf{y} + \mathbf{e}_i) - v(x, \mathbf{y})$ and combinations of these operators (for example, $\Delta_i \Delta_0 v(x, \mathbf{y}) = \Delta_0 v(x, \mathbf{y} + \mathbf{e}_i) - \Delta_0 v(x, \mathbf{y})$). Next, we define \mathcal{U} , a set

of real-valued functions in $\mathbb{N} \times \mathcal{Y}$, with the following properties.

DEFINITION 1. If $v \in \mathcal{U}$, then for all $(x, \mathbf{y}) \in \mathbb{N} \times \mathcal{Y}$ and for all $i \in \mathcal{A}$:

Condition C.1 $\Delta_i \Delta_0 v(x, \mathbf{y}) \leq 0$ if $y_i < m_i$

Condition C.2 $\Delta_{0+i} \Delta_0 v(x, \mathbf{y}) \geq 0$ if $y_i < m_i$

Condition C.3 $\Delta_0 \Delta_0 v(x, \mathbf{y}) \geq 0$

Condition C.4 $\Delta_0 v(x, \mathbf{y}) \geq -c_1$.

Condition C.1 states that $\Delta_0 v$ is nonincreasing in y_i or equivalently v is submodular in the direction of \mathbf{e}_0 and \mathbf{e}_i . Condition C.2 states that v is supermodular in the direction of \mathbf{e}_0 and $\mathbf{e}_0 + \mathbf{e}_i$. Condition C.3 states that $\Delta_0 v$ is nondecreasing in x or equivalently v is convex in x . For more on sub/supermodularity, see Veatch and Wein (1992) for an example.

PROPERTY 1. Define $s(\mathbf{y}) = \min[x \geq 0 \mid \Delta_0 v(x, \mathbf{y}) > 0]$, $r_i(\mathbf{y}) = \min[x \geq 1 \mid \Delta_0 v(x-1, \mathbf{y} - \mathbf{e}_i) + c_i > 0]$ if $i \in \mathcal{A}$, and $r_i(\mathbf{y}) = \min[x \geq 1 \mid \Delta_0 v(x-1, \mathbf{y}) + c_i > 0]$ if $i \in \mathcal{W}$, where $v(x, \mathbf{y} - \mathbf{e}_i) = 0$ when $y_i = 0$. Then C.3 in Definition 1 implies the following:

$$\Delta_0 v(x, \mathbf{y}) > 0 \quad \text{if and only if } x \geq s(\mathbf{y}), \quad (3)$$

$$\Delta_0 v(x-1, \mathbf{y} - \mathbf{e}_i) + c_i > 0$$

$$\text{if and only if } x \geq r_i(\mathbf{y}) \quad \text{for } i \in \mathcal{A}, \quad \text{and} \quad (4)$$

$$\Delta_0 v(x-1, \mathbf{y}) + c_i > 0$$

$$\text{if and only if } x \geq r_i(\mathbf{y}) \quad \text{for } i \in \mathcal{W}. \quad (5)$$

In the proof of Theorem 1, we will show that the optimal value function v^* satisfies Conditions C.1–C.4; i.e., $v^* \in \mathcal{U}$. We will also show that if we construct a policy π^* such that the actions specified by π^* are given by

$$a_0^{\pi^*}(x, \mathbf{y}) = \begin{cases} 0 & \text{if } x \geq s^*(\mathbf{y}) \\ 1 & \text{otherwise, and} \end{cases}$$

$$a_i^{\pi^*}(x, \mathbf{y}) = \begin{cases} 1 & \text{if } x \geq r_i^*(\mathbf{y}) \\ 0 & \text{otherwise,} \end{cases}$$

then π^* is optimal, with $s^*(\mathbf{y}) = \min[x \geq 0 \mid \Delta_0 v^*(x, \mathbf{y}) > 0]$, $r_i^*(\mathbf{y}) = \min[x \geq 1 \mid \Delta_0 v^*(x-1, \mathbf{y} - \mathbf{e}_i) + c_i > 0]$ if $i \in \mathcal{A}$, and $r_i^*(\mathbf{y}) = \min[x \geq 1 \mid \Delta_0 v^*(x-1, \mathbf{y}) + c_i > 0]$ if $i \in \mathcal{W}$. In other words, under policy π^* , there is a base-stock level $s^*(\mathbf{y})$ and rationing levels $r_i^*(\mathbf{y})$

such that we produce if $x < s^*(\mathbf{y})$ (and idle otherwise) and we fulfill due orders from class j if $x \geq r_j^*(\mathbf{y})$ (and reject otherwise).

The shapes of the n -dimensional *switching surfaces*, defined by the base-stock and rationing levels, are a priori very general and complex. The following result states, however, that $s(\mathbf{y})$ and $r_j(\mathbf{y})$, defined for any $v \in \mathcal{U}$, are nondecreasing in each of the variables y_i , with unit increases in y_i leading to at most unit increases in $s(\mathbf{y})$ and $r_j(\mathbf{y})$. This means for the optimal policy that the increase in the base-stock and rationing levels (as a function of any of the variables y_i) is bounded by a linear function with unit slope. The proof of this result and of all subsequent ones (unless stated otherwise) are included in the online appendix.

PROPERTY 2. *Let $v \in \mathcal{U}$; then*

$$s(\mathbf{y} + \mathbf{e}_i) = s(\mathbf{y}) \quad \text{or} \quad s(\mathbf{y} + \mathbf{e}_i) = s(\mathbf{y}) + 1$$

$$\text{if } y_i < m_i, \quad i \in \mathcal{A}, \quad \text{and} \quad (6)$$

$$r_j(\mathbf{y} + \mathbf{e}_i) = r_j(\mathbf{y}) \quad \text{or} \quad r_j(\mathbf{y} + \mathbf{e}_i) = r_j(\mathbf{y}) + 1$$

$$\text{if } y_i < m_i, \quad j \in \mathcal{A} \cup \mathcal{W}, \quad i \in \mathcal{A}. \quad (7)$$

Furthermore, the rationing levels are ordered as stated in the following property.

PROPERTY 3. *Let $v \in \mathcal{U}$ and consider two classes of customers i and j with lost sales costs $c_i \geq c_j$. Let $\mathbf{y} \in \mathcal{Y}$; then:*

- *If $i \in \mathcal{A} \cup \mathcal{W}$ and $j \in \mathcal{W}$, then $r_i(\mathbf{y}) \leq r_j(\mathbf{y})$.*
- *If $i \in \mathcal{A} \cup \mathcal{W}$ and $j \in \mathcal{A}$, then $r_i(\mathbf{y}) \leq r_j(\mathbf{y}) + 1$.*

Although the second bullet in Property 3 suggests that in the case of $i \in \mathcal{A} \cup \mathcal{W}$ and $j \in \mathcal{A}$, we either have $r_i(\mathbf{y}) \leq r_j(\mathbf{y})$ or $r_i(\mathbf{y}) = r_j(\mathbf{y}) + 1$, we suspect that the first inequality is always true. This has proven difficult to show analytically, but it is supported by numerical results for the optimal policy. The nested structure of the rationing levels has intuitive appeal. When on-hand inventory is sufficiently high, orders from all classes are fulfilled. When inventory drops below a certain threshold, orders from the least important class (the class with the lowest sales cost) are rejected and on-hand inventory is reserved for future orders from more important classes. When inventory drops below another threshold, orders from the two least important classes are rejected while

orders from other classes continue to be fulfilled. This process continues with more classes being rejected as inventory drops below successively lower thresholds. Eventually, only orders from the most important class are fulfilled from on-hand inventory. The exact value of the rationing threshold for each class is determined by comparing the optimal expected cost associated with rejecting an order from a class (and reserving the existing inventory for future orders from more important classes) to the optimal expected cost of using existing inventory to fulfill an order from this class (and face a potential shortage for more important classes in the future).

We should note that inventory rationing of this kind is not uncommon in practice, for example, in assembly systems where a component is shared among multiple products. Inventory of this component is often rationed among the different products based on their profitability. Inventory rationing is also common in settings where inventory cannot be replenished (e.g., the seats on an airplane for a particular flight). In that case, the rationing levels are referred to as booking limits, with each booking limit corresponding to a threshold on the number of seats below which customers from a particular fare class are rejected in favor of reserving seats for potential future customers from higher fare classes. The specificity of our model is to consider thresholds modulated by an imperfect demand information.

The following lemma shows that the application of the operator T to a function $v \in \mathcal{U}$ preserves the properties of v . This result is critical to showing that $v^* \in \mathcal{U}$ and that the properties of the optimal policy described in Theorem 1 hold.

LEMMA 1. *If $v \in \mathcal{U}$, then $Tv \in \mathcal{U}$.*

The proof of Lemma 1 makes use of Property 1 to simplify the operator expressions. We then proceed by showing that Tv for $v \in \mathcal{U}$ satisfies Conditions C.1–C.4 in Definition 1 and, therefore, is in \mathcal{U} .

We are now ready to state our main result.

THEOREM 1. *The optimal value function v^* belongs to \mathcal{U} . Further, there exists a stationary optimal policy that consists of a base-stock production policy with a state-dependent base-stock level and a multilevel inventory-rationing policy with state-dependent rationing levels for*

each class. Specifically, the optimal policy has the following properties:

Property P.1. For each vector $\mathbf{y} = (y_1, \dots, y_n)$, there is a corresponding base-stock level $s^*(\mathbf{y})$ such that it is optimal to produce if $x < s^*(\mathbf{y})$ and not to produce otherwise.

Property P.2. The base-stock level $s^*(\mathbf{y})$ is nondecreasing in each of the variables y_i ($i \in \mathcal{A}$) with $s^*(\mathbf{y}) \leq s^*(\mathbf{y} + \mathbf{e}_i) \leq s^*(\mathbf{y}) + 1$.

Property P.3. For each vector $\mathbf{y} = (y_1, \dots, y_n)$, there is a corresponding rationing level for each product j , $r_j^*(\mathbf{y})$, such that it is optimal to fulfill an order from class j if $x \geq r_j^*(\mathbf{y})$ and not to fulfill it otherwise.

Property P.4. The rationing level $r_j^*(\mathbf{y})$ is nondecreasing in each of the variables y_i ($i \in \mathcal{A}$) with $r_j^*(\mathbf{y}) \leq r_j^*(\mathbf{y} + \mathbf{e}_i) \leq r_j^*(\mathbf{y}) + 1$.

Property P.5. If $c_i \geq c_j$ with $i \in \mathcal{A} \cup \mathcal{W}$, then:

- $r_i^*(\mathbf{y}) \leq r_j^*(\mathbf{y})$ when $j \in \mathcal{W}$,
- $r_i^*(\mathbf{y}) \leq r_j^*(\mathbf{y}) + 1$ when $j \in \mathcal{A}$.

Property P.6. It is always optimal to fulfill orders from class 1 customers whenever there is inventory; i.e., $r_1(\mathbf{y}) = 1$.

In Figure 1, we illustrate the optimal policy for a system with two classes, where class 2 provides ADI but not class 1. It is perhaps surprising to observe that the rationing level for class 2 can increase as the number of its announced orders increases. In other words, knowing that there are more orders announced from class 2 induces the supplier to reserve more inventory for class 1. However, this does not necessarily mean, in the long run, that fewer orders from class 2 would be fulfilled, because the overall base-stock level also increases with the number of announced orders. The fact that the base-stock level increases reduces

the probability that class 2 orders would be rejected. In turn, this allows for a more aggressive protection against shortages for class 1 by increasing the rationing level for class 2 (without necessarily negatively affecting class 2 shortages). In general, the rationing level for any class can be affected by the number of announced orders from any class, regardless of whether these classes have higher or lower lost sales costs. For example, we observed numerically, for a system with three classes all with ADI and $c_1 > c_2 > c_3$, that the rationing level for class 2, $r_2(y_1, y_2, y_3)$, can strictly increase not only in y_1 and y_2 but also in y_3 .

Theorem 1 characterizes the optimal policy for a general n -dimensional problem. Known results from the literature for simpler problems can be retrieved as special cases. For instance, without ADI, the state of the system can be described by the inventory level x only and the switching curves reduce to fixed thresholds as stated in the following corollary, which corresponds to the main result in Ha (1997a).

COROLLARY 1. When none of the demand classes provides ADI, the optimal policy consists of a base-stock production policy with a fixed-based stock level s^* and an inventory-rationing policy with n fixed rationing levels r_j^* , $1 \leq j \leq n$ such that

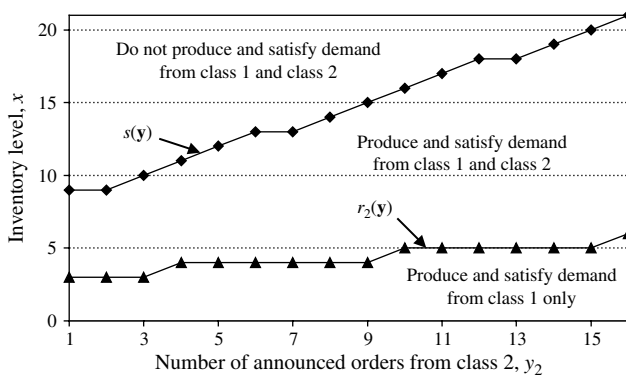
1. It is optimal to produce if $x < s^*$ and not to produce otherwise.
2. It is optimal to fulfill an order from class j if $x \geq r_j^*$ and not to fulfill it otherwise.
3. The rationing thresholds are ordered such that $r_1^* \leq r_2^* \leq \dots \leq r_n^*$.
4. It is always optimal to fulfill orders from class 1; i.e., $r_1^* = 1$.

PROOF. Immediate from Theorem 1 with $\mathcal{A} = \emptyset$. \square

When a single demand class is considered, it has been shown for other models with ADI that the optimal policy is a state-dependent base-stock policy, which increases in the number of announced orders (see, for instance, Karaesmen et al. 2002). The following corollary provides a similar result for our continuous time production-inventory setting.

COROLLARY 2. In a system with a single demand class, the optimal policy consists of a base-stock policy with a

Figure 1 The Structure of the Optimal Policy



Note. $\lambda_1 = \lambda_2 = 0.4$, $c_1 = 1,000$, $c_2 = 100$, $p_1 = p_2 = 1$, $v_2 = 0.1$.

state-dependent base-stock level $s^*(\mathbf{y})$, where y is the number of announced orders, such that:

1. It is optimal to produce if $x < s^*(\mathbf{y})$ and not to produce otherwise.
2. The base-stock level $s(\mathbf{y})$ is nondecreasing in y with $s(\mathbf{y} + 1) \leq s(\mathbf{y}) + 1$.
3. It is always optimal to fulfill orders whenever inventory is available.

PROOF. Immediate from Theorem 1 with $n = 1$. \square

Note that when the demand classes have identical lost sales costs and demand lead time distributions (i.e., $c_i = c_j = c$ and $v_i = v_j = v$, for any pair i and j), the system is equivalent to a single demand class system with arrival rate $\lambda = \lambda_1 + \dots + \lambda_n$. This follows from the fact that the superposition of n Poisson processes is a Poisson process. Corollary 2 thus implies that the optimal policy is a base-stock policy with a base-stock level that only depends on the sum of announced orders, such that $s^*(\mathbf{y}) = s^*(y_1 + \dots + y_n)$ and $s^*(\mathbf{y})$ is nondecreasing in $y_1 + \dots + y_n$.

In a system where the lost sales costs are identical but the mean demand lead times are different, the optimal production policy is a base-stock policy but with a base-stock level $s^*(\mathbf{y})$ that depends on the individual values of y_i for $i = 1, \dots, n$. There is no inventory rationing in this case, because the lost sales costs are identical, and it is optimal to allocate inventory on a first-come, first-served (FCFS) basis. This case can be used to model settings where there is a single demand class but information about the distribution of demand lead times is updated when an order is announced. That is, with probability $\lambda_i / \sum_{i=1}^n \lambda_i$, the lead time demand of an announced order is exponential with parameter v_i , where this information becomes available when the order is announced.

In some settings, inventory rationing is not used or is not an option (for reasons exogenous to our model). In this case, orders from all demand classes must be fulfilled whenever there is available inventory. The system manager decides only when to produce. Our original MDP formulation can still be used to treat this problem if we restrict ourselves to policies π with $a_i^\pi(x, \mathbf{y}) = 1$, for $i = 1, \dots, n$. In particular, we can show that the optimal policy is a base-stock policy with a state-dependent base-stock level $s(\mathbf{y})$, with the base-stock level satisfying Properties P.1 and P.2 of Theorem 1. We will refer to the optimal policy in this

case as the FCFS policy. This policy can be used to benchmark the optimal rationing policy and to study the benefit of rationing, with or without ADI. We do this in the numerical experiments described in §5.

So far, we have assumed that the number of announced orders for class $i \in \mathcal{A}$ is bounded by a finite number $m_i < \infty$. We have done so because the case of finite m_i is of interest by itself and because it serves as a basis for treating the case of infinite m_i . The latter is difficult to analyze directly because it involves a problem with unbounded transition rates, making rate uniformization impossible. Nevertheless, in the following theorem, we show that our results extend to the infinite m_i case, with the structure of the optimal policy remaining unchanged.

THEOREM 2. For systems with $m_i = \infty$ for $i \in \mathcal{A}$, there exists an optimal stationary policy. Furthermore, the optimal policy consists of a state-dependent base-stock $s(\mathbf{y})$ and rationing levels $r_i(\mathbf{y})$ for $i = 1, \dots, n$, with Properties P.1–P.6 described in Theorem 1.

5. Numerical Study

In this section, we describe results from a numerical study that we carried out to examine the benefit of using ADI with both suppliers and customers and to compare the value of ADI to that of inventory rationing.

5.1. Computational Procedure

Numerical results are obtained by solving the dynamic programs corresponding to each problem instance using the value-iteration method. The value-iteration algorithm is terminated only when a five-digit accuracy is achieved. The state space is truncated at $[0, m_0] \times \dots \times [0, m_n]$, where m_i is a positive integer for $i = 1, \dots, n$. The size of the state space is increased until the average cost is no longer sensitive to the truncation level. For all problem instances, we assume that the holding costs are linear, and in this context we set, without loss of generality, $h(x) = x$. Also without loss of generality, we set $\mu = 1$.

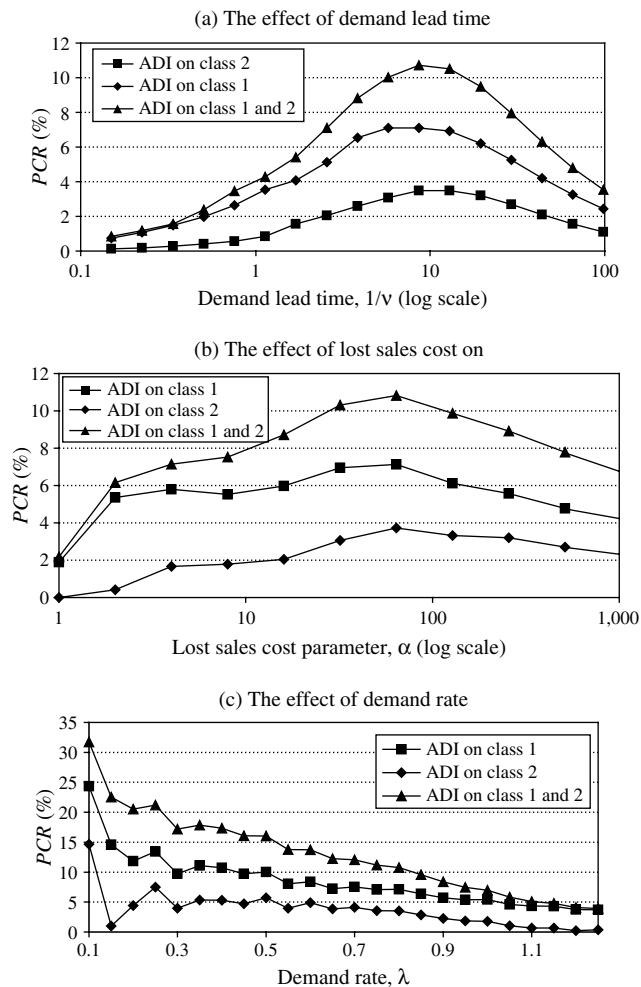
5.2. The Benefit of ADI to the Supplier

We consider a system with two customer classes. We obtain the optimal average costs $g_{A(1)}^*$, $g_{A(2)}^*$, and $g_{A(12)}^*$ that correspond respectively to systems with ADI on only class 1, ADI on only class 2, and ADI on both classes 1 and 2. We compare these costs to the optimal

average cost g_W^* obtained for a system without ADI on both classes. The systems are the same, except that in the absence of ADI, announced orders are due immediately. This means that if an announced order is not immediately fulfilled from stock, it is considered lost and incurs a lost sales penalty.

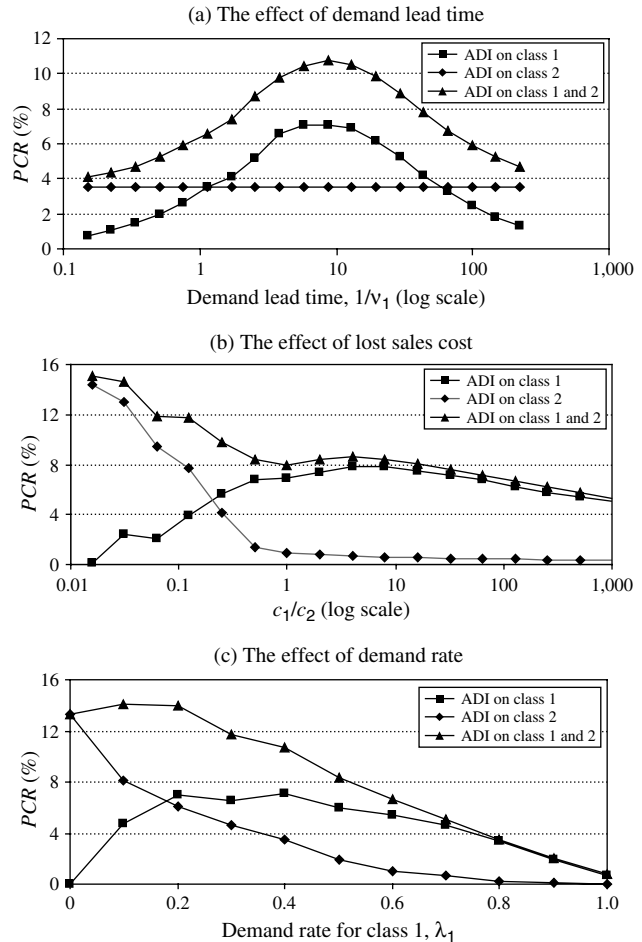
Representative numerical results are shown in Figures 2 and 3 where the percentage cost reduction $PCR(i) = (g_W^* - g_{A(i)}^*)/g_W^*$, $i = 1, 2$, and 12 has been obtained for the three ADI scenarios and for a wide range of values of the three main parameters; demand lead time, lost sales costs, and demand rates. In each

Figure 2 The Effect of Varying System Parameters for Both Classes on the Benefit of ADI to the Supplier Under Different ADI Scenarios



Notes. (a) $\lambda_1 = \lambda_2 = 0.4$, $c_1 = 200$, $c_2 = 50$, $p_1 = p_2 = 1$, $v_1 = v_2 = v$; (b) $\lambda_1 = \lambda_2 = 0.4$, $c_1 = 4\alpha$, $c_2 = \alpha$, $p_1 = p_2 = 1$, $v_1 = v_2 = 0.1$; (c) $\lambda_1 = \lambda_2 = 0.5$, $c_1 = 200$, $c_2 = 50$, $p_1 = p_2 = 1$, $v_1 = v_2 = 0.1$.

Figure 3 The Effect of Varying System Parameters for Class 1 Only on the Benefit of ADI to the Supplier Under Different ADI Scenarios



Notes. (a) $\lambda_1 = \lambda_2 = 0.4$, $c_1 = 200$, $c_2 = 50$, $p_1 = p_2 = 1$, $v_2 = 0.1$; (b) $\lambda_1 = 0.7$, $\lambda_2 = 0.1$, $c_2 = 100$, $p_1 = p_2 = 1$, $v_1 = v_2 = 0.1$; (c) $\lambda_2 = 0.4$, $c_1 = 200$, $c_2 = 50$, $p_1 = p_2 = 1$, $v_1 = v_2 = 0.1$.

figure, we vary the value of one parameter, over the entire range of plausible values, while keeping the other parameters fixed. Based on Figure 2, the following observations can be made:

- The benefit of ADI to the supplier can be relatively significant, with cost savings in excess of 30% in some cases (the average cost saving is 9.8% for the cases shown).
- The benefit of ADI is higher when all suppliers provide information.
- When customer classes, except for their lost sales costs, have similar parameters, it is more valuable to have ADI from the class with the higher cost.

- The benefit of ADI is sensitive (in a nonmonotonic fashion in some cases) to system operating parameters, namely, demand lead times, lost sales costs, and capacity utilization.

- The relative benefit of ADI tends to be insignificant when expected lead times and lost sales costs are either very small or very large. There are values in the middle range for expected demand lead times and lost sales costs for which the relative benefit of ADI is maximum.

- The relative benefit of ADI tends to be decreasing in the demand rates (or, more generally, capacity utilization), and it becomes insignificant when the demand rates (capacity utilization) are very high.

The fact that the relative benefit of ADI is not monotonic in demand lead time and lost sales cost can be explained as follows. Consider first the effect of demand lead time. When expected demand lead time is small, orders are due shortly after they are announced. Hence, there is little opportunity to take advantage of this information to affect either production or inventory allocation. When expected demand lead time is large, orders are announced far in advance of their expected due date, leading on average to a large number of announced orders. This makes ADI less useful, because as expected value of demand lead time increases, so does the variance. In the limit, our estimate of the due date of the next order reduces to our a priori estimate (i.e., the due date of the next order of type i is treated as exponentially distributed with rate λ_i). We would expect this effect to be absent if the expected value of demand lead time were to increase but the variance were to stay the same or decrease. For example, in systems when demand lead times are deterministic, we expect longer lead times to be always more valuable.

The fact that the relative benefit of ADI is not monotonic in the lost sales costs is somewhat easier to explain. When the lost sales costs are small, the penalty from ignoring ADI (e.g., not producing when we should) is small relative to the inventory-holding cost. When the lost sales costs are very large, the base-stock level tends to be high regardless of the number of announced orders. This also means that the fraction of unfulfilled orders tends to be relatively small. Hence, the impact of ADI on how inventory is allocated among the demand classes is insignificant.

The effect of the demand rate (shown in Figure 2(c)) can be explained as follows. When the total demand rate is very small, it is optimal, without ADI, to hold no inventory and to incur the lost sales costs instead (the optimal decision is to never produce). In contrast, with ADI, it is optimal to produce whenever the number of announced orders becomes sufficiently large. Hence with ADI, there is an opportunity to satisfy at least a fraction of the demand, even when it is very small. Although the absolute difference between systems with and without ADI is small when the total demand rate is small, the relative difference can be significant. As the total demand rate increases, the advantage of systems with ADI tends to diminish. When demand is very high, the optimal production decision, with or without ADI, is to always produce. Because demand from both classes is large, only demand from the class with the higher lost sales cost can be satisfied. Therefore, the optimal inventory allocation, regardless of ADI, is to always reject orders from the class with the lower cost.

In Figure 3, we illustrate the impact of varying the parameters of class 1 only, instead of varying the parameters of both classes as we do in Figure 2. The results reveal that having ADI on class 1 (the class with the higher lost sales cost) is not always more desirable than having ADI on class 2. If the demand lead time for class 1 is either very short or very long, it may be more desirable to have ADI on class 2 if its demand lead time is in the middle range. Similarly, if the demand rate of class 1 is sufficiently small relative to the demand rate of class 2, then having ADI on class 2 is more beneficial. As shown in Figure 3(b), the point at which having ADI on a particular class becomes preferable depends, of course, on the ratio of the lost sales costs c_1/c_2 . In general, whether or not having ADI on class 1 or class 2 is more preferable does not appear to follow a simple rule. There is a complex relationship among the demand lead times, demand rates, and lost sales cost.

Finally, we should note that the lack of “smoothness” in the curves displaying the impact of various parameters is due to the discreteness of the base-stock and rationing levels. This effect is most pronounced when the base-stock levels are small (e.g., when demand rates or lost sales costs are small).

5.3. The Benefit of ADI to Customers

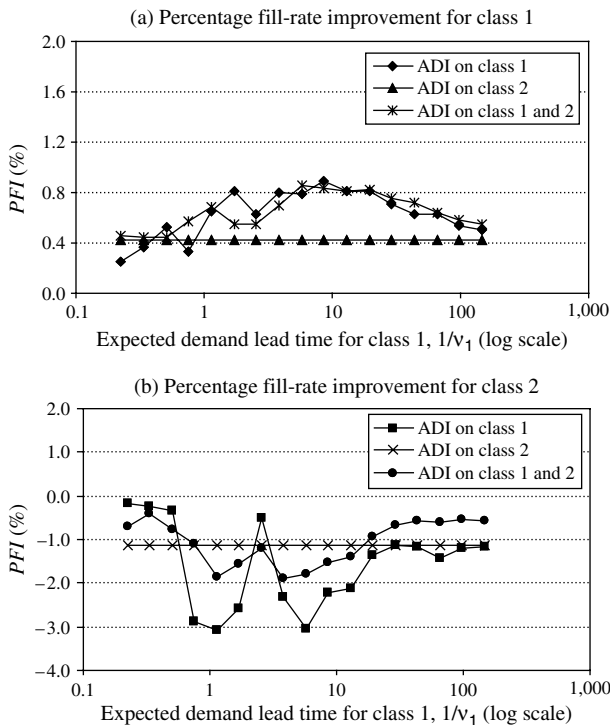
The results of the previous section show that, depending on system parameters, suppliers can realize significant benefits by having customers provide ADI. However, it is not clear if customers benefit equally. To study the impact of ADI on customers, we examine the quality of service received by customers with and without ADI. We measure a customer’s service quality by fill rate, which corresponds to the long-run fraction of the customer’s orders filled from on-hand inventory. We present results for systems with two customer classes. We let $f_A(i)^*$ and $f_W(i)^*$ denote, respectively, the fill rate with and without ADI for customer class $i = 1, 2$. In Figure 4, we present sample results for a system with two customer classes showing the percentage fill-rate improvement $PFI(i) = (f_A(i)^* - f_W(i)^*)/f_W(i)^*$ for customer class i due to ADI for varying values of demand lead time for class 1 and for scenarios with and without ADI on the other class.

Perhaps surprisingly, ADI has for the most part little effect on the quality of service customers receive

(these results are consistent with those obtained from a larger data when other parameters are varied). In fact, in some cases, the quality of service diminishes with ADI, regardless of whether it is on class 1 or class 2. The supplier appears to use ADI in some cases to reduce inventory costs at the expense of customer service. Moreover, a class that offers ADI can in some cases negatively affect the service level another class receives. It may also negatively impact its own service level while improving the service level of another class. Note that the impact of demand lead time on the percentage fill-rate improvement is somewhat erratic. This is again due to the discreteness in the base-stock and rationing levels and the complex relationship among various parameters.

The fact that the supplier appears to benefit more raises the obvious question of why customers would be willing to provide ADI. In practice, the answer may be that customers would agree to provide ADI only if there is a contractual agreement that service levels would be improved or, alternatively, that the penalties for poor service would be increased. For example, this could be achieved by the customer negotiating an increase in the penalty for not fulfilling demand immediately—i.e., a higher lost sales cost. Based on numerical results (not shown), we observed that customers can indeed negotiate a significantly higher lost sales penalty, by up to 50% in some observed cases, in exchange for providing ADI.

Figure 4 The Effect of Varying Demand Lead Time on Fill-Rate Improvement




Note. $\lambda_1 = \lambda_2 = 0.4$, $c_1 = 200$, $c_2 = 50$, $v_2 = 0.1$, $p_1 = p_2 = 1$.

5.4. The Benefit of Partial ADI

In settings where ADI is available only from a subset of the customer classes, an important question that arises is what the marginal benefit is to the supplier of increasing the fraction of customers that provide ADI. To shed some light on this question, we consider a system with two classes. Class 1, with demand rate λ_1 , offers ADI, and class 2, with demand rate λ_2 , does not. We examine the effect of increasing the fraction of customers with ADI by varying the ratio $\beta = \lambda_1 / (\lambda_1 + \lambda_2)$ while maintaining $\Lambda = \lambda_1 + \lambda_2$. Higher values of β indicate higher availability of ADI.

Representative results from numerical experiments are shown in Figure 5 (note that to isolate the effect of β , we let the lost sales penalty be the same for both classes; i.e. $c_1 = c_2 = c$ and $v_1 = v_2 = v$). As we can see, the relative benefit of ADI does not exhibit

Figure 5 The Effect of Partial Advanced Demand Information on Supplier's Cost




Note. $c = 100$, $p = 1$, $v = 0.1$.

diminishing returns with increases in the fraction β of customers with ADI. In addition, the benefit is growing almost linearly in β (this appears due to the fact that parameter values other than demand rates for both classes are the same). This suggests that in this case ADI on any particular order yields a benefit that is independent of whether or not there is ADI on other orders. In practice, this means that additional investments in ADI can remain equally beneficial, regardless of previous investments.

5.5. ADI vs. Inventory Rationing

In some settings, inventory rationing is not possible (e.g., withholding available inventory from certain customers is not an acceptable practice). In those settings, it is useful to evaluate the extent to which ADI continues to be useful and to compare the benefits gained from ADI to those gained from rationing. To carry out such comparisons, we consider a system with two classes with penalty costs c_1 and c_2 , where $c_1 > c_2$. We obtained the optimal average cost for a system with (1) neither ADI nor inventory rationing, (2) ADI only (on both classes), (3) inventory rationing only, and (4) both ADI and inventory rationing. For system (1), orders are fulfilled on a FCFS basis, regardless of their class, and production decisions are made without the benefit of ADI. For system (2), orders are fulfilled on a FCFS basis regardless of their class, but production decisions take into account announced orders. For system (3), inventory is optimally rationed among the classes, but production decisions are made without the benefit of ADI. Finally for system (4), inventory is optimally rationed, and production decisions are optimally made taking into account announced orders.

Figure 6 The Effects of Inventory Rationing vs. Advanced Demand Information on Supplier's Cost



Note. $\lambda_1 = \lambda_2 = 0.4$, $c_1 + c_2 = 100$, $p_1 = p_2 = 1$, $v_1 = v_2 = 0.1$.

In Figure 6, we show results depicting the impact of ADI alone, rationing alone, and joint ADI and rationing. The results indicate that the benefits of inventory rationing and ADI are complementary. The benefit of inventory rationing is more significant than that of ADI if the lost sales cost ratio is relatively high. The fact that the benefit of jointly using ADI and inventory rationing is at least the sum of the individual benefits of ADI (alone) and rationing (alone) suggests that one cannot be used as a substitute for the other. ADI and rationing appear to bring two different types of benefits to the supplier. This is supported by the fact that ADI tends to affect primarily decisions about production (although rationing levels are also affected) and rationing primarily affects decisions about inventory allocation. The value of rationing increases with the cost ratio c_1/c_2 (it becomes beneficial to reserve inventory for the more important class). In general, the value of ADI is sensitive to the ratio c_1/c_2 . It is not in the example shown because of our choice of parameter values. We vary c_1/c_2 while keeping $c_1 + c_2$ constant and $\lambda_1 = \lambda_2$, which in some cases makes the results insensitive to changes in c_1/c_2 .

6. Conclusion and Future Research

We have considered production control and inventory allocation in an integrated production-inventory system with multiple customer classes and imperfect ADI. In our model, ADI is not perfect because (1) order due dates are not known exactly, (2) orders can be cancelled by the customers, and (3) ADI is available only from a subset of the customers. We showed that the optimal production policy consists of a base-stock policy with state-dependent base-stock

levels, where the state is determined by the inventory level and the number of announced orders from each class. We showed that the optimal inventory-allocation policy consists of a rationing policy with state-dependent rationing levels such that it is optimal to fulfill orders from a particular class only if the inventory level is above the rationing level corresponding to that class.

Using numerical results, we showed that taking into account ADI can be beneficial to the supplier. However, we found that these benefits can be sensitive (sometimes in a nonmonotonic fashion) to various system parameters. Somewhat surprisingly, we found that customers benefit less from ADI than suppliers, with suppliers using ADI in some cases to reduce inventory costs at the expense of customer service levels. We showed how customers could extract some of this value from the supplier by imposing higher lost sales penalties in exchange for ADI. For the supplier, we showed that more benefits can be realized by appropriately rationing inventory among the different customer classes, when their lost sales costs are sufficiently different, than by collecting ADI. However, when both rationing and ADI are used, we found their benefits to be cumulative. Furthermore, we found that the benefit of ADI to the supplier does not exhibit diminishing returns with increases in the fraction of customers that provides ADI.

There are several possible avenues for future research. Our model could be generalized by substituting the exponential distribution for demand lead time, production time, or order interarrival time by phase-type distributions that are useful in approximating other distributions. The use of phase-type distributions retains the Markovian property of the system and continues to allow the formulation of the problem as an MDP. For demand lead time, the phase-type distribution would also allow us to model settings where due dates are progressively updated. For example, the distribution of demand lead time could be modeled as an Erlang distribution with k stages. As announced orders move from stage to stage, the expected due date of the order is updated. The order becomes due when it leaves the k th stage.

Our model could be extended to the case where backorders are allowed. Although we do not expect

the structure of the optimal policy to change significantly, the analysis does become less tractable, because the state space must include the number of backorders for each customer class. Our model could also be embedded within models that explicitly encompass decisions by both customers and suppliers. For example, using a game theory framework, our model could serve as a building block for exploring how customers should negotiate service levels with or without ADI and/or how to set price discounts in exchange for ADI.

Another worthwhile area for future research is the development of simple yet effective heuristics. In particular, it may be possible to construct heuristics that mimic the optimal policy by specifying the base-stock and rationing levels in terms of simple functions (e.g., linear functions) of the state variables. Such a heuristic could be designed to preserve the properties of the optimal policy yet be simpler to communicate and perhaps implement.

Electronic Companion

An electronic companion to this paper is available on the *Manufacturing & Service Operations Management* website (<http://msom.pubs.informs.org/ecompanion.html>).

Acknowledgments

The authors would like to thank the associate editor and two anonymous reviewers for valuable comments. They are also thankful to Yves Dallery for the many useful discussions and for generous feedback throughout the course of this work.

References

- Benjaafar, S., M. Elhafi. 2006. Production and inventory control of a single product assemble-to-order system with multiple customer classes. *Management Sci.* **52** 1896–1912.
- Benjaafar, S., W. L. Cooper, S. Mardan. 2007. Production-inventory systems with imperfect advance demand information and updating. Working Paper, University of Minnesota, Minneapolis.
- Benjaafar, S., M. Elhafi, F. de Véricourt. 2004. Demand allocation in multi-product, multi-facility make-to-stock systems. *Management Sci.* **50** 1431–1448.
- Bertsekas, D. 2001. *Dynamic Programming and Optimal Control*, Vol. 2, 2nd ed. Optimization and Computation Series, Athena Scientific, Nashua, NH.
- Buzacott, J. A., J. G. Shanthikumar. 1993. *Stochastic Models of Manufacturing Systems*. Prentice-Hall, Upper Saddle River, NJ.
- Buzacott, J. A., J. G. Shanthikumar. 1994. Safety stock versus safety time in MRP controlled production systems. *Management Sci.* **40** 1678–1689.

- Chen, F. 2001. Market segmentation, advance demand information, and supply chain performance. *Manufacturing Service Oper. Management* **3** 53–67.
- Cohen, M. A., P. R. Kleindorfer, H. L. Lee. 1988. Service constrained (s, S) inventory systems with priority demand classes and lost sales. *Management Sci.* **34** 482–499.
- de Véricourt, F., F. Karaesmen, Y. Dallery. 2002. Stock allocation for a capacitated supply system. *Management Sci.* **48** 1486–1501.
- Deshpande, V., M. A. Cohen, K. Donohue. 2003. A threshold inventory rationing policy for service-differentiated demand classes. *Management Sci.* **49** 683–703.
- Frank, K. C., R. Q. Zhang, I. Duenyas. 2003. Optimal policies for inventory systems with priority demand classes. *Oper. Res.* **51** 993–1002.
- Gallego, G., O. Ozer. 2001. Integrating replenishment decisions with advance order information. *Management Sci.* **47** 1344–1360.
- Gallego, G., O. Ozer. 2002. Optimal use of demand information in supply chain management. J. S. Song, D. D. Yao, eds. *Supply Chain Structures: Coordination, Information and Optimization*. Kluwer Academic, Norwell, MA, 119–160.
- Graves, S. C., H. C. Meal, S. Dasu, Y. Qiu. 1986. Two-stage production planning in a dynamic environment. S. Axsäter, C. Schneeweiss, E. Silver, eds. *Multi-Stage Production Planning and Control—Lecture Notes in Economics and Mathematical Systems*. Springer-Verlag, Berlin, 9–43.
- Güllü, R. 1996. On the value of information in dynamic production/inventory problems under forecast evolution. *Naval Res. Logist.* **43** 289–303.
- Guo, X., O. Hernández-Lerma. 2003. Continuous-time controlled Markov chains with discounted rewards. *Acta Applicandae Mathematicae* **79** 195–216.
- Ha, A. Y. 1997a. Inventory rationing in a make-to-stock production system with several demand classes and lost sales. *Management Sci.* **43** 1093–1103.
- Ha, A. Y. 1997b. Stock-rationing policy for a make-to-stock production system with two priority classes and backordering. *Naval Res. Logist.* **44** 457–472.
- Hariharan, R., P. Zipkin. 1995. Customer-order information, lead-times and inventories. *Management Sci.* **41** 1599–1607.
- Heath, D. C., P. L. Jackson. 1994. Modeling the evolution of demand forecasts with application to safety-stock analysis in production/distribution systems. *IIE Trans.* **26** 17–30.
- Hu, X., I. Duenyas, R. Kapuscinski. 2003. Advance demand information and safety capacity as a hedge against demand and capacity uncertainty. Working paper, University of Michigan, Ann Arbor.
- Karaesmen, F., J. A. Buzacott, Y. Dallery. 2002. Integrating advance order information in production control. *IIE Trans.* **34** 649–662.
- Karaesmen, F., G. Liberopoulos, Y. Dallery. 2004. The value of advance demand information in production/inventory systems. *Ann. Oper. Res.* **126** 135–158.
- Lippman, S. 1975. Applying a new device in the optimization of exponential queueing systems. *Oper. Res.* **23** 687–710.
- Nahmias, S., S. Demmy. 1981. Operating characteristics of an inventory system with rationing. *Management Sci.* **27** 1236–1245.
- Özer, O., W. Wei. 2004. Inventory control with limited capacity and advance demand information. *Oper. Res.* **52** 988–1000.
- Puterman, M. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley and Sons, Inc., New York.
- Schwarz, L. B., N. C. Petrucci, K. Wee. 1997. The value of advance-order information and the implications for managing the supply chain: An information/control/buffer portfolio perspective. Working paper, Purdue University, West Lafayette, IN.
- Tan, T., R. Güllü, N. Erkip. 2005. Using imperfect advance demand information in ordering and rationing decisions. Working paper, Eindhoven University of Technology, The Netherlands.
- Topkis, D. M. 1968. Optimal ordering and rationing policies in a nonstationary dynamic inventory model with n demand classes. *Management Sci.* **15** 160–176.
- van Donselaar, K., L. R. Kopczak, M. Wouters. 2001. The use of advance demand information in a project-based supply chain. *Eur. J. Oper. Res.* **130** 519–538.
- Veatch, M. H., L. M. Wein. 1992. Monotone control of queueing networks. *Queueing Systems* **12** 391–408.
- Zhu, K., U. Thonemann. 2004. Modeling the benefits of sharing future demand information. *Oper. Res.* **52** 136–147.
- Zipkin, P. H. 2000. *Foundations of Inventory Management*. McGraw-Hill, New York.