

Production-Inventory Systems with Imperfect Advance Demand Information and Updating

Saif Benjaafar,¹ William L. Cooper,¹ Setareh Mardan²

¹ Program in Industrial and Systems Engineering, University of Minnesota, 111 Church Street S.E., Minneapolis, Minnesota 55455

² PROS, 3100 Main Street, Suite #900, Houston, Texas 77002

Received 19 September 2008; revised 27 October 2010; accepted 16 November 2010

DOI 10.1002/nav.20443

Published online 3 February 2011 in Wiley Online Library (wileyonlinelibrary.com).

Abstract: We consider a supplier with finite production capacity and stochastic production times. Customers provide advance demand information (ADI) to the supplier by announcing orders ahead of their due dates. However, this information is not perfect, and customers may request an order be fulfilled prior to or later than the expected due date. Customers update the status of their orders, but the time between consecutive updates is random. We formulate the production-control problem as a continuous-time Markov decision process and prove there is an optimal state-dependent base-stock policy, where the base-stock levels depend upon the numbers of orders at various stages of update. In addition, we derive results on the sensitivity of the state-dependent base-stock levels to the number of orders in each stage of update. In a numerical study, we examine the benefit of ADI, and find that it is most valuable to the supplier when the time between updates is moderate. We also consider the impact of holding and backorder costs, numbers of updates, and the fraction of customers that provide ADI. In addition, we find that while ADI is always beneficial to the supplier, this may not be the case for the customers who provide the ADI. © 2011 Wiley Periodicals, Inc. *Naval Research Logistics* 58: 88–106, 2011

Keywords: advance demand information; production-inventory systems; make-to-stock queues; continuous-time Markov decision processes

1. INTRODUCTION

It is increasingly common for members of the same supply chain to share advance demand information (ADI). This practice has been facilitated by information technologies such as the Internet, electronic data interchange (EDI), and radio frequency identification (RFID). It has also been supported by initiatives such as the inter-industry consortium on Collaborative Planning, Forecasting and Replenishment (CPFR), which provides a framework for participating companies to share future demand projections and coordinate ordering decisions. Large manufacturers, such as Toyota and Boeing, are tightly integrated with their first tier suppliers with whom they share production status, inventory usage, and even future design plans. Large retailers, such as Wal-Mart and Best Buy, have invested in sophisticated information collecting and processing infrastructure that enables them to share real-time

inventory usage and point-of-sale (POS) data with thousands of their suppliers. Several manufacturers that sell directly to the consumer, such as Dell, and online retailers, such as Amazon, encourage their customers to place their orders early by offering discounts to those that accept later delivery dates. In some industries, suppliers allow their long-term customers to place soft orders far ahead of due dates, which may then later be firmed up, modified, or canceled.

Although ADI can take on different forms and may be enabled by a variety of technologies, it typically reduces to customers providing advance notice to their suppliers about the timing and size of future orders. This information can be perfect (exact information about future orders) or imperfect (estimates of timing or quantity of future orders). The information can also be explicit, with customers directly stating their intent about future orders, or implicit, with customers allowing suppliers to observe their internal operations and to determine estimates of future orders (the systems we consider in this article are in part motivated by settings with such implicit information; we provide examples and further discussion later in this section). It is generally believed that ADI,

Additional Supporting Information may be found in the online version of this article.

Correspondence to: W.L. Cooper (billcoop@me.umn.edu)

even if imperfect, can improve supply chain performance. In particular, with information about future demand, a supplier may be able to reduce the need for inventory or excess capacity. Customers may also benefit through improved service quality or lower costs.

However, the availability of ADI raises questions. How should a supplier use ADI to make decisions? How valuable is ADI to suppliers and customers and how is this value affected by operating characteristics of the supplier and the quality of information provided by customers? How significant are benefits from receiving information further in advance or increasing the portion of customers that provide ADI? Is ADI equally beneficial to all parties in the supply chain and could it be harmful, particularly to the customer who provides it?

We address these and other related questions for a supplier that produces a single product. Customers furnish the supplier with ADI by announcing orders ahead of their due dates. However, this information is not perfect, and orders may become due prior to or later than the announced expected due date or they can be canceled altogether. Hence, the demand leadtime (the time between when an order is announced and when it is requested or canceled) is random. Customers provide status updates as their orders progress toward becoming due but the time between consecutive updates is also random and independent of updates for other orders. In this article, we are primarily motivated by settings where customers implicitly provide ADI to the supplier by allowing the supplier to observe their internal operations (e.g., order fulfillment, manufacturing, inventory usage), thereby enabling the supplier to estimate when customers will eventually place orders. We refer to such internal operations as the demand leadtime system.

For most of the article, we assume that the actual due dates of different orders are independent of each other and orders that are announced later can become due before (or after) those that are announced earlier. Updates are also independent and do not follow a first-announced, first-updated rule. We refer to this as a system with independent due dates (IDD). In Section 5, we show how our treatment can be extended to systems where announced orders are updated and the orders become due in the sequence in which they are announced. We refer to this as a system with sequential due dates (SDD).

The following examples illustrate the types of settings we model in this article. Consider a supplier that provides a component to a manufacturer, such as Boeing, of a large and complex product (e.g., an aircraft). The manufacturer informs the supplier each time it initiates the production of a new product and each time it completes a stage of the production process. The component provided by the supplier is not immediately needed and is required only at a later stage of the production process. The manufacturer does not accept early deliveries, but wishes to receive the component as soon as it is needed in a just-in-time fashion. The supplier uses the

information about the progression of the product through the manufacturer's production process to estimate when it will need to make a delivery to the manufacturer. To make such estimates, the supplier uses its knowledge of the manufacturer's operations and available data from past interactions. However, the estimates are imperfect and the manufacturer (due to inherent variability) may complete a production stage sooner or later than expected. The manufacturer may initiate, in response to its own demand, the production of multiple products simultaneously (e.g., an aircraft manufacturer may assemble multiple airplanes in parallel). The evolution of these products through the production process is largely independent, so that a product that enters a particular stage of production later than another product may complete it sooner.

This type of ADI may also arise in settings other than manufacturing. For example, van Donselaar et al. [30] present a case study of how builders provide material suppliers with ADI about the start and progress of construction projects. The suppliers use the information to estimate when a builder will need materials. This estimation is not perfect because progress on a construction project can be variable and because design specifications may change over the course of a project, sometimes leading the builders ultimately not to place orders.

In this article, we provide a general framework for modeling systems with this type of ADI. The framework is broad enough to model a wide range of demand leadtime systems with various assumptions regarding due date updating. The demand leadtime system can be viewed in general as a queueing system composed of parallel servers, with service times consisting of multiple stages of random duration and the completion of a stage of corresponding to an update. Arrivals to the demand leadtime system correspond to orders being announced; similar to service times, interarrival times are random and may have multiple stages, with the completion of a stage indicating an update. Departures from the demand leadtime system correspond to orders becoming due.

In the systems we study, the supplier has finite capacity, producing items one at a time with stochastic production times. Hence, the supplier itself can be viewed as a single server queue, with arrivals corresponding to orders becoming due (i.e., an arrival to the supplier is a departure from the demand leadtime system). The supplier has the ability to produce items ahead of their due dates in a make-to-stock fashion. However, items in inventory incur a holding cost. When an order becomes due and it cannot be immediately satisfied from inventory, it is backordered but it incurs a backorder cost. The supplier's objective is to find a production control policy to minimize the expected total discounted cost or the expected average cost per unit time.

We formulate the problem as a continuous-time Markov decision process (MDP). We show that there is an optimal production policy that is a state-dependent base-stock policy, wherein the supplier produces if and only if the net inventory

is below a base-stock level, which depends only upon the numbers of orders in various stages of update. We also derive results on the sensitivity of the state-dependent base-stock levels to the numbers of orders in each stage of update. For SDD systems, we obtain similar results. In our analysis, we develop a method for proving structural properties of optimal policies of continuous-time Markov decision processes (CTMDPs) with unbounded jump rates. The derived structure is useful as it allows one to compute and store an optimal policy in terms of just the base-stock levels, simplifying the policy's implementation. The structure of the optimal policy can also guide construction of simpler heuristics if needed, or in assessing the effectiveness of heuristics that may already be in use.

We also conduct a numerical study to examine the benefits of ADI to both suppliers and customers by comparing systems with ADI, without ADI, and with partial ADI. The study yields several managerial insights, a few of which we now summarize. Increasing the average demand leadtime by increasing the number of updates always reduces the supplier's cost. However, given a fixed number of updates, increasing the average time between updates may increase or decrease cost. ADI is most valuable when the average time between updates is moderate. ADI is less valuable when the average time between updates is short, because there is little time to react to information. It is also less valuable when the average time between updates is long, because the earlier notice comes with an increase in variability of the demand leadtime. This points out that obtaining earlier notice (on average) of orders is not necessarily desirable, and that when evaluating the benefit of ADI, it is important to account for the mechanism by which this ADI might be obtained. The incremental cost reduction from updating is often small compared with that from announcing orders ahead of their due dates. Typically, much of the benefit of ADI can be realized if customers provide just initial advance order announcements and few or no updates. Although ADI leads to an overall reduction in cost, in some cases it may be used by the supplier to reduce inventory at the expense of more backorders. Therefore, customers that provide ADI may witness a decline in service levels. However, in exchange for ADI, customers are in position to negotiate an increase in the backorder penalty they apply to the supplier. Higher backorder penalties can serve as a mechanism for customers to deter suppliers from reducing service levels, or as a mechanism to share, indirectly, the cost savings from ADI.

The remainder of the article is organized as follows. In Section 2, we give a brief literature review and summarize our contribution. In Section 3, we formulate the problem and describe the structure of an optimal policy. We also discuss extensions including systems with variable numbers of updates, order cancelations, multiple customer classes, and lost sales. In Section 4, we present numerical results. In

Section 5, we extend our analysis to systems with sequential updating. In Section 6, we offer a summary and concluding comments. Proofs are in the Appendix and Supporting Information.

2. LITERATURE REVIEW AND SUMMARY OF CONTRIBUTIONS

There is a growing literature on inventory systems with ADI. A review of much of this work can be found in Gallego and Özer [9]. Models can be broadly classified into two categories based on whether inventory is reviewed periodically or continuously.

For systems with periodic review, ADI is typically modeled as information available about demand in future periods. Under varying assumptions, Gallego and Özer [8], Özer and Wei [25], and Schwarz et al. [27] have shown the existence of optimal state-dependent base-stock policies for periodic-review problems with ADI. In these articles, the base-stock levels depend upon a vector of advance orders for future periods. Özer [24] extends the analysis to distribution systems with multiple retailers, Gallego and Özer [10] to serial systems, and Wang and Toktay [33] to systems with flexible delivery. Other papers that consider periodic review systems with ADI include Thonemann [29], Gavirneni et al. [11], and DeCroix and Mookerjee [6].

For continuous-review inventory systems with ADI, Buzacott and Shanthikumar [3] consider production-inventory systems with ADI and evaluate policies that use two parameters: a base-stock level and a release leadtime. Hariharan and Zipkin [17] introduced the notion of demand leadtime in a system where orders are announced a fixed amount of time before they are due. For constant supply leadtimes and Poisson order arrivals, they show that there is an optimal base-stock policy with a fixed base-stock level. Karaesmen et al. [20] analyze a discrete-time model with constant demand leadtimes that is similar to our SDD model with no due-date updating (see Section 5). They prove the optimality of state-dependent base-stock policies. Gallego and Özer [9; Section 2.4] consider a system similar to a special case of our SDD setting but with exogenous load-independent supply leadtimes. Gayon et al. [12] study a system similar to our IDD scheme but with multiple demand classes, lost sales, and no due-date updates. Other papers that deal with continuous-review systems include Liberopoulos et al. [22] and Karaesmen et al. [21].

It is possible to view the demand leadtime system in our model as a Markovian demand-modulating process with transition probabilities between states determined by the dynamics of order announcements and due dates. Previous literature (see, e.g., [4]) has established the optimality of state-dependent base-stock policies for periodic review

inventory problems with Markov-modulated demand and exogenous leadtimes. The results of Chen and Song do not directly apply to our setting of endogenous leadtimes and continuous review. Nevertheless, it might be possible to develop an alternate analysis of the systems considered herein using techniques from the study of inventory models with Markov-demand demand.

ADI can be viewed as a form of forecast updating. Examples of papers that deal with inventory systems with periodic forecast updates include Graves et al. [13], Heath and Jackson [18], Güllü [14], Sethi et al. [28], Zhu and Thonemann [34], and references therein. The models we present in this article can be viewed as dealing with forecast updates. However, in our case the updates are with respect to the timing of future demand.

Finally, there is a literature that deals with how a supplier should quote delivery leadtimes to its customers; see, for example, Duenyas and Hopp [7], Hopp and Sturgis [19], and references therein. The setting studied in this literature is quite different from ours and typically concerns make-to-order systems where no finished goods inventory is held in advance of customer orders.

Relative to the above literature, we make the following contributions. Our article is the first to consider imperfect ADI with updates for continuous-review production-inventory systems. It also appears to be the first to directly model stochastic demand leadtimes and distinguish between systems with independent and sequential due date updates, and to derive the structure of an optimal policy for each. The article offers one of the most general models of ADI in the literature (e.g., systems with no updates or with a single update can be treated as special cases). The modeling framework is flexible and can accommodate additional features such as random numbers of updates, order cancelations, multiple demand classes, and lost sales. Moreover, the numerical study yields new insights on the benefit of ADI to suppliers, highlighting important effects due to capacity, demand leadtime, and cost parameters. It also contrasts the impact of (i) increasing the number of updates, (ii) increasing the fraction of customers who give ADI, and (iii) increasing the length of individual update stages. The numerical results also shed light on effects of ADI on customers. We show that customers may see their service quality deteriorate if they provide ADI to their suppliers.

Beyond the context of ADI, our article also contains an approach for proving structural properties of optimal policies of CTMDPs with unbounded jump rates. (The IDD model has unbounded jump rates.) The usual approach for proving structural properties for CTMDPs with bounded jump rates is to first uniformize (see, e.g., [23]) the CTMDP to get an equivalent discrete-time Markov decision process (DTMDP), and then to show that certain properties of functions are preserved by the DTMDP transition operator. Results then follow

using induction and the convergence of value iteration. With unbounded jump rates, uniformization cannot be applied, and hence the “usual approach” does not work. Our method for CTMDPs with unbounded jump rates involves proving the desired structural properties for each of a sequence of problems with bounded jump rates, and then extending to the problem with unbounded jump rates by passing to a limit via a suitably chosen subsequence and appealing to results of Guo and Hernández-Lerma [15]. Although the method is somewhat intuitive, it involves resolving a number of non-trivial technical issues, some of which are problem specific. A variant of the approach was used in the paper by Gayon et al. [12] cited above. The approach may prove useful in other problems with unbounded jump rates.

3. PROBLEM FORMULATION AND STRUCTURE OF AN OPTIMAL POLICY

Consider a supplier of a single product, who can produce at most one unit of the product at a time. The supplier may hold completed units of the product in inventory. Any such unit of inventory incurs a holding cost of h per unit time.

We model ADI through the notion of a demand leadtime system. (As mentioned in the introduction, such a system may represent the internal operations of customers; ADI is provided implicitly by allowing the supplier to view these internal operations.) Orders for the product are announced before their due dates. Such announcements may be viewed as arrivals to the demand leadtime system. We assume that the announcements arrive continuously over time according to a Poisson process with rate λ . The amount of time between when an order is announced and when it becomes due is random. We refer to this random variable as the demand leadtime. The demand leadtime of an order is the amount of time it spends in the demand leadtime system. We assume orders are homogeneous in the sense that demand leadtimes have the same distribution for all orders, and hence the expected demand leadtime is the same for all orders.

After an order is announced, it progresses through the demand leadtime system before becoming due. Specifically, it undergoes a series of $k - 1$ updates ($k \geq 1$). For $i = 1, \dots, k - 1$, the time between the $(i - 1)$ th and i th update is exponentially distributed with mean v_i^{-1} . (The 0th update is the order’s initial announcement.) The time between the $(k - 1)$ th update and when the order becomes due is exponentially distributed with mean v_k^{-1} . Hence, each demand leadtime consists of k exponentially distributed stages with the expected demand leadtime of each order equal to $v_1^{-1} + \dots + v_k^{-1}$. (The case with $k = 1$ represents a situation with no updates and exponential demand leadtimes.) When an order has undergone exactly $i - 1$ updates we say that it is in stage i . Viewed in this fashion, the i th update corresponds to an order moving from stage i to stage $(i + 1)$. When an

order undergoes its i th update, the supplier learns that the order's expected remaining demand leadtime has decreased from $v_i^{-1} + \dots + v_k^{-1}$ to $v_{i+1}^{-1} + \dots + v_k^{-1}$. Equivalently, we may think of demand leadtime as having a phase-type distribution with k phases in series. Information is provided each time the demand leadtime completes a phase. In the case where $v_i = v$ for $i = 1, \dots, k$, demand leadtimes have an Erlang distribution. Note that the process by which orders are announced, updated, and become due can be viewed as an M/G/ ∞ queue.

After an order becomes due (i.e., after it leaves the demand leadtime system), the supplier fills the order if it has inventory on hand. If the supplier does not have inventory on hand, then the order is backordered and incurs a backorder cost of b per unit time. Orders do not incur backorder costs before they are due; i.e., orders in the demand leadtime system do not incur backorder costs. As mentioned above, the supplier can produce one item at a time. We assume that production times are exponentially distributed with mean μ^{-1} . Hence, the production process can itself be viewed as a queue, whose input is provided by the output of the demand leadtime system.

The assumptions of Poisson arrivals of announcements and exponential production and update times are made in part for mathematical tractability as they allow the problem to be cast as a CTMDP. They are also appropriate for approximating systems with high variability. Such Markovian assumptions are consistent with previous studies of production-inventory systems; see e.g., Buzacott and Shanthikumar [2], Ha [16], Zipkin [35], de Véricourt et al. [5], and others. Later, we partially relax these assumptions.

In the remainder of this section, we develop the formulation and describe the structure of the optimal policy. This is done by first analyzing in Section 3.1 a simplified version (with a truncated state space) of the problem. We then extend the analysis in Section 3.2 to systems without the truncation and state our main result for this section in Theorem 2.

3.1. Bounded Jump Rates

In this section, we assume that the total number of announced orders (i.e., the number of orders in the demand leadtime system) at any instant remains bounded by a finite integer $m < \infty$, so that $\sum_{i=1}^k y_i \leq m$, where y_i is the number of orders in stage i . Order announcements are rejected and leave without entering the leadtime demand system (and hence never become due) if $\sum_{i=1}^k y_i = m$. This assumption allows us to formulate the problem as a Markov decision process with bounded jump rates. From a queueing perspective, the introduction of the finite m means that we approximate the M/G/ ∞ queue mentioned above by an M/G/ m/m queue (an Erlang loss system). When m is chosen to be large, very few order announcements are rejected and hence the arrival rate of due orders to the production facility will be roughly the same as the arrival rate of order announcements to the demand leadtime system. In fact, the precise arrival rate

of due orders to the production facility will be $\lambda(1 - B(m))$ where $B(m)$ is the probability that an M/G/ m/m queue is full. The probability $B(m)$ approaches 0 as $m \rightarrow \infty$. The exact value of $B(m)$ is given by the well-known Erlang loss formula. Using the results of this section as a building block, in Section 3.2 we extend our results to the case with no bound on the total announced orders.

Let \mathbb{Z} and \mathbb{Z}_+ be respectively the sets of integers and non-negative integers, and let \mathbb{Z}^k and \mathbb{Z}_+^k be their k -dimensional cross products. Let \mathbb{R} be the real numbers. Throughout, $\mathbf{y} = (y_1, \dots, y_k)$. The MDP has state space $\mathcal{S}_m := \mathbb{Z} \times \mathbb{Z}_+^k(m)$, where $\mathbb{Z}_+^k(m) := \{\mathbf{y} \in \mathbb{Z}_+^k : \sum_{i=1}^k y_i \leq m\}$. To keep notation clean, we will indicate the dependence on m only in notation that is used later for extending to the case without m . It is, however, important to keep in mind that most of the quantities in this section do depend upon m , even if this is not reflected in the notation.

The state of the system is determined by $X(t)$, which represents the net inventory at time t , and $\mathbf{Y}(t) = (Y_1(t), \dots, Y_k(t))$, where $Y_i(t)$ is the number of announced orders in stage i at time t . In each state, two actions are possible: produce or idle (do not produce). The objective is to find a production policy that minimizes the long-run expected discounted cost. Let the set of such production policies be denoted by Π . A deterministic stationary policy $\pi := \{\pi(x, \mathbf{y}) : (x, \mathbf{y}) \in \mathcal{S}_m\}$ specifies the action taken at any time as a function only of the state of the system, where $\pi(x, \mathbf{y}) = 1$ means produce in state (x, \mathbf{y}) , and $\pi(x, \mathbf{y}) = 0$ means idle in state (x, \mathbf{y}) .

We will work with a uniformized version (see, e.g., [23]) of the problem in which the transition rate in each state under any action is $\Lambda := \lambda + \mu + m \sum_{i=1}^k v_i$ so that the transition times $0 = \tau_0 \leq \tau_1 \leq \tau_2 \leq \dots$ are such that $\{\tau_{n+1} - \tau_n : n \geq 0\}$ is a sequence of i.i.d. exponential random variables, each with mean Λ^{-1} . Let $\{(X_n, \mathbf{Y}_n) : n \geq 0\}$ denote the embedded Markov chain of states; that is, $(X_n, \mathbf{Y}_n) := (X(\tau_n), \mathbf{Y}(\tau_n))$ is the state immediately after the n th transition. For $i = 1, \dots, k$, let \mathbf{e}_i be the k -dimensional vector with 1 in position i and zeros elsewhere. Let \mathbf{e}_0 be the k -dimensional vector of zeros. If action $a \in \{0, 1\}$ is selected in state (x, \mathbf{y}) , then the next state of the embedded Markov chain is (x', \mathbf{y}') with probability

$$P_{(x, \mathbf{y}), (x', \mathbf{y}')} (a) := \begin{cases} \Lambda^{-1} \mu \mathbb{I}_{\{a=1\}} & \text{if } (x', \mathbf{y}') = (x + 1, \mathbf{y}) \\ \Lambda^{-1} \lambda \mathbb{I}_{\{\bar{y} < m\}} & \text{if } (x', \mathbf{y}') = (x, \mathbf{y} + \mathbf{e}_1) \\ \Lambda^{-1} v_i y_i \mathbb{I}_{\{y_i \geq 1\}} & \text{if } (x', \mathbf{y}') = (x, \mathbf{y} + \mathbf{e}_{i+1} - \mathbf{e}_i) \\ \Lambda^{-1} v_k y_k \mathbb{I}_{\{y_k \geq 1\}} & \text{if } (x', \mathbf{y}') = (x - 1, \mathbf{y} - \mathbf{e}_k) \\ \Lambda^{-1} [\Lambda - \mu \mathbb{I}_{\{a=1\}} - \lambda \mathbb{I}_{\{\bar{y} < m\}} - \sum_{i=1}^k v_i y_i \mathbb{I}_{\{y_i \geq 1\}}] & \text{if } (x', \mathbf{y}') = (x, \mathbf{y}) \\ 0 & \text{otherwise,} \end{cases}$$

where $\bar{y} := \sum_{i=1}^k y_i$ and $\mathbb{I}_{\{\cdot\}}$ is the indicator function. The cost rate when the state is (x, \mathbf{y}) is $c(x, \mathbf{y}) = c(x) := hx^+ + bx^-$,

where $h > 0$ and $b > 0$ are the per-unit holding and back-order cost rates, and $x^+ = \max\{x, 0\}$ and $x^- = -\min\{x, 0\}$. Here, we again emphasize that backorder costs are incurred only when an order becomes due and is not immediately satisfied. Jobs inside the demand leadtime system (which have been announced, but which are not yet due) do not incur backorder costs.

The value function, which specifies the optimal expected total discounted cost, is given by

$$\begin{aligned} v_m^*(x, \mathbf{y}) &:= \inf_{\pi \in \Pi} E_{(x, \mathbf{y})}^\pi \left[\int_{t=0}^{\infty} e^{-\beta t} c(X(t)) dt \right] \\ &= \inf_{\pi \in \Pi} E_{(x, \mathbf{y})}^\pi \left[\sum_{n=0}^{\infty} \left(\frac{\Lambda}{\gamma} \right)^n \frac{c(X_n)}{\gamma} \right], \end{aligned} \quad (1)$$

where $\beta > 0$ is the discount rate, $\gamma := \beta + \Lambda$, and $E_{(x, \mathbf{y})}^\pi$ denotes expectation with respect to the probability measure determined by policy π and $(X(0), \mathbf{Y}(0)) = (x, \mathbf{y})$.

Let V be the set of real-valued functions on \mathcal{S}_m and let v be an arbitrary element of V . Define $T_\lambda, T'_i, T_\mu : V \rightarrow V$ as follows

$$\begin{aligned} T_\lambda v(x, \mathbf{y}) &:= v(x, \mathbf{y} + \mathbf{e}_1 \mathbb{I}_{\{\bar{y} < m\}}) \\ T'_i v(x, \mathbf{y}) &:= v(x, \mathbf{y} + [\mathbf{e}_{i+1} - \mathbf{e}_i] \mathbb{I}_{\{y_i \geq 1\}}) \quad i = 1, \dots, k-1 \end{aligned} \quad (2)$$

$$T'_k v(x, \mathbf{y}) := v(x - \mathbb{I}_{\{y_k \geq 1\}}, \mathbf{y} - \mathbf{e}_k \mathbb{I}_{\{y_k \geq 1\}}) \quad (3)$$

$$T_\mu v(x, \mathbf{y}) := \min\{v(x, \mathbf{y}), v(x+1, \mathbf{y})\}.$$

Consider also the operator $T : V \rightarrow V$ defined by

$$Tv(x, \mathbf{y}) := \min_{a \in \{0,1\}} \left\{ \frac{c(x)}{\gamma} + \frac{\Lambda}{\gamma} \sum_{(x', \mathbf{y}') \in \mathcal{S}} p_{(x, \mathbf{y}), (x', \mathbf{y}')} (a) v(x', \mathbf{y}') \right\} \quad (4)$$

$$\begin{aligned} &= \gamma^{-1} \left[c(x) + \lambda T_\lambda v(x, \mathbf{y}) + \sum_{i=1}^k v_i y_i T'_i v(x, \mathbf{y}) \right. \\ &\quad \left. + \sum_{i=1}^k v_i (m - y_i) v(x, \mathbf{y}) + \mu T_\mu v(x, \mathbf{y}) \right]. \end{aligned} \quad (5)$$

The function v_m^* defined in (1) is the minimum non-negative solution of the optimality equation

$$v = Tv, \quad (6)$$

and moreover a stationary policy that specifies for each (x, \mathbf{y}) an action that attains the minimum on the right-hand side of (6) is optimal. See, e.g., Section 11.5 of Puterman [26].

In the optimality equation (6), operator T_λ corresponds to the arrival of a customer. More precisely, if $v(x, \mathbf{y})$ represents the “value” of being in state (x, \mathbf{y}) , then $T_\lambda v(x, \mathbf{y})$ is the value

just after an arrival occurs when the state is (x, \mathbf{y}) . Similarly, operator T'_i ; $i = 1, \dots, k-1$ corresponds to an update of an order from stage i to stage $i+1$ and operator T'_k corresponds to an order becoming due. Operator T_μ corresponds to the production decision. When $v(x+1, \mathbf{y}) < v(x, \mathbf{y})$, it is better to produce a unit of inventory than it is to idle. In this case, $T_\mu v(x, \mathbf{y}) = v(x+1, \mathbf{y})$ represents the value just after the completion of the unit of inventory when the state is (x, \mathbf{y}) . When $v(x+1, \mathbf{y}) \geq v(x, \mathbf{y})$, it is instead better to idle, in which case $T_\mu v(x, \mathbf{y}) = v(x, \mathbf{y})$. The other term — $\sum_{i=1}^k v_i (m - y_i) v(x, \mathbf{y})$ —in the optimality equation corresponds to null transitions introduced through uniformization of the jump rate. To understand the term λ that multiplies $T_\lambda v(x, \mathbf{y})$, note that in state (x, \mathbf{y}) , the next event will be an arrival with probability λ/Λ . Similar interpretations are possible for the other multipliers. The term λ also represents the rate of order announcements. Likewise, μ is the rate of potential production completions, $v_i y_i$ is the rate of updates at stage i when y_i orders are in stage i , and $\sum_{i=1}^k v_i (m - y_i)$ is the rate of null transitions when there are \mathbf{y} jobs in the demand leadtime system. Hence Λ is the overall rate of (real and null) transitions.

In preparation for Theorem 1, let $\Delta v(x, \mathbf{y}) := v(x+1, \mathbf{y}) - v(x, \mathbf{y})$ for $v \in V$ and let $U := \{v \in V : v \text{ satisfies conditions (C1)–(C4)}\}$, where conditions (C1)–(C4) are defined as follows:

- (C1) $\Delta v(x, \mathbf{y}) \leq \Delta v(x+1, \mathbf{y})$ for all $x \in \mathbb{Z}$, $\mathbf{y} \in \mathbb{Z}_+^k(m)$.
- (C2) $\Delta v(x, \mathbf{y} + \mathbf{e}_j) \leq \Delta v(x+1, \mathbf{y} + \mathbf{e}_j)$ for all $x \in \mathbb{Z}$, $\mathbf{y} \in \mathbb{Z}_+^k(m-1)$, $j = 0, \dots, k-1$, and $l = j+1, \dots, k$.
- (C3) $\Delta v(x, \mathbf{y} + \mathbf{e}_{j+1}) \leq \Delta v(x, \mathbf{y} + \mathbf{e}_j)$ for all $x \in \mathbb{Z}$, $\mathbf{y} \in \mathbb{Z}_+^k(m-1)$, and $j = 0, \dots, k-1$.
- (C4) $\Delta v(x, \mathbf{y}) \leq 0$ for all $x \in \mathbb{Z}$ with $x < 0$, $\mathbf{y} \in \mathbb{Z}_+^k(m)$.

As we can see from Proposition 1 below, the value function satisfies these conditions. The fact that the value function satisfies these conditions implies certain structural properties for the optimal policy. In particular, Condition (C1) is a convexity property that can be used to show the existence of a state-dependent base-stock optimal policy. Conditions (C2) and (C3) can be used to show that the announcement of a new order or the update of an existing order will cause the base-stock level either to increase by one or to remain unchanged. Condition (C4) can be used to show that it is optimal to produce whenever there are backorders.

PROPOSITION 1: The value function is an element of U ; that is, $v_m^* \in U$.

The proof of Proposition 1 is in Section S-1 of the Supporting Information. We are now ready for the main result of the section. Theorem 1 describes the structure of an optimal policy; a proof is in the Appendix.

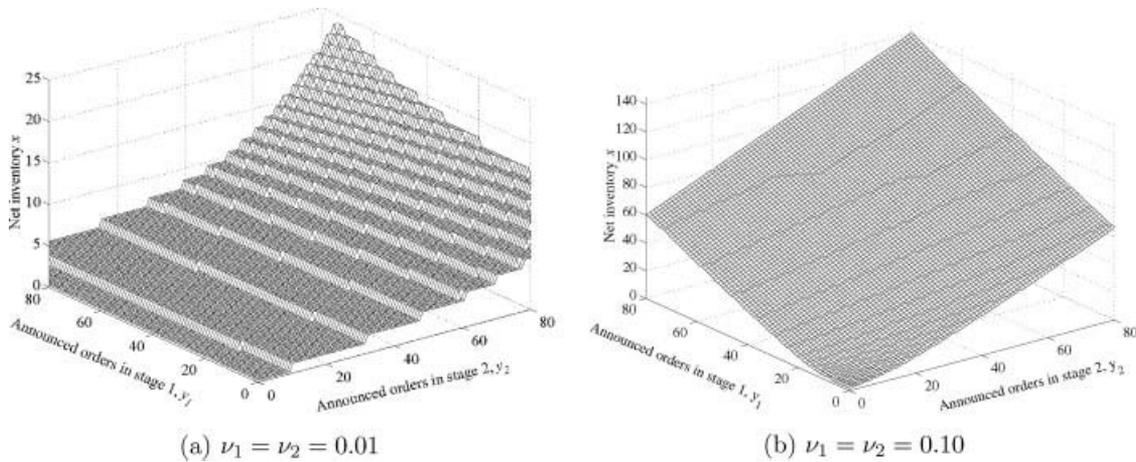


Figure 1. Optimal policies for two different systems with IDD and $k = 2$: The surfaces depict the state-dependent base-stock levels. For a given $\mathbf{y} = (y_1, y_2)$, if the net inventory on hand x is below the surface, it is optimal to produce; if the net inventory on hand x is on or above the surface, it is optimal to idle ($m = 200, \mu = 1, \lambda = 0.6, h = 10, b = 100$).

THEOREM 1: The stationary state-dependent base-stock policy $\pi^* = \{\pi^*(x, \mathbf{y})\}$ given by

$$\pi^*(x, \mathbf{y}) := \begin{cases} 0 & \text{if } x \geq s_{\mathbf{y}} \\ 1 & \text{if } x < s_{\mathbf{y}} \end{cases} \quad (7)$$

where $s_{\mathbf{y}} := \min\{x : v_m^*(x + 1, \mathbf{y}) - v_m^*(x, \mathbf{y}) \geq 0\}$ is optimal. In addition, (a) the base-stock levels satisfy $s_{\mathbf{y}+\mathbf{e}_l} \in \{s_{\mathbf{y}+\mathbf{e}_j}, s_{\mathbf{y}+\mathbf{e}_j} + 1\}$ for $j = 0, \dots, k - 1; l = j + 1, \dots, k$ and (b) $\pi^*(x, \mathbf{y}) = 1$ if $x < 0$.

Theorem 1 states that for each vector \mathbf{y} of announced orders there exists a threshold $s_{\mathbf{y}}$ such that it is optimal to produce if net inventory is less than $s_{\mathbf{y}}$, and it is optimal to idle if net inventory is at least $s_{\mathbf{y}}$. We refer to the parameters $\{s_{\mathbf{y}}\}$ as the \mathbf{y} -dependent base-stock levels. Part (a) with $l = j + 1$ indicates that the \mathbf{y} -dependent base-stock level increases by at most one if an order is updated or if a new order is announced. It also follows from part (a) that $s_{\mathbf{y}}$ is increasing in each component of \mathbf{y} ; i.e., $s_{\mathbf{y}} \leq s_{\mathbf{y}+\mathbf{e}_i}$ for $i \in \mathbb{Z}_+$ and $j = 1, \dots, k$. Part (b) states that it is optimal to produce if there are any backorders. These results are consistent with those obtained by Özer and Wei [25], who show a similar structure to the optimal policy in periodic-review systems where ADI consists of confirmed demand for future periods (e.g., see Theorem 2 in [25]).

Figure 1 illustrates the structure described in Theorem 1 for two examples, each with $k = 2$ stages. In part (a) of the figure, the mean time $1/\nu_i$ spent in each of the stages is relatively long, and hence the production policy is much less sensitive to orders in Stage 1 than it is to orders in Stage 2. In part (b) the mean time $1/\nu_i$ is shorter, and hence the production policy treats orders in Stage 1 almost the same as orders in Stage 2.

Above, we focused on the discounted-cost optimality criteria. A treatment of average cost can be found in Section S-2 of the Supporting Information, where Theorem S-1 shows that a direct analog of Theorem 1 holds for the average-cost optimality criteria under the additional assumption that $\lambda < \mu$ (which ensures that production can keep up with demand and prevent backorders from growing “infinitely large”). In the discounted-cost case, we do not need this assumption.

We close this section by illustrating the flexibility of our modeling framework in accommodating additional features. In particular, we consider four extensions to our basic model: (1) systems with random numbers of updates, (2) systems with order cancelations, (3) systems with two demand classes, one providing ADI and the other one not, and (4) systems with lost sales.

3.1.1. Systems with Random Numbers of Updates

Suppose that customers update their orders a random number of times. In particular, suppose that given an announced order is at stage i , it will, independent of everything else, become due after the end of stage i with probability q_i and progress to next stage with probability $1 - q_i$, for $i = 1, \dots, k - 1$. To extend the model to such a setting with a random number of updates, we need to replace T'_i in (5) by \tilde{T}'_i defined by

$$\begin{aligned} \tilde{T}'_i v(x, \mathbf{y}) := & (1 - q_i) T'_i v(x, \mathbf{y}) \\ & + q_i v(x - \mathbb{1}_{\{y_i \geq 1\}}, \mathbf{y} - \mathbf{e}_i \mathbb{1}_{\{y_i \geq 1\}}) \quad i = 1, \dots, k. \end{aligned} \quad (8)$$

In this case, there is again an optimal state-dependent base stock policy and it is optimal to produce whenever backorders are present. However, to prove properties (a) and (b) of the

base-stock levels in Theorem 1 we impose the condition that $v_1q_1 \leq v_2q_2 \leq \dots \leq v_kq_k$; that is, v_iq_i is nondecreasing in i . A proof is in Section S-3 of the Supporting Information. To understand the importance of this condition, suppose temporarily that the condition does not hold. More specifically, suppose, e.g., that v_1q_1 is much larger than v_2q_2 . Then an order in Stage 1 is, in a sense, “closer” to becoming due than is an order in Stage 2. To see this, observe that an order in Stage 1 tends to quickly become due after just one stage of update whereas, in the event that the order progresses to Stage 2, it tends to remain there a (relatively) long time. Hence, there may be (x, \mathbf{y}) such that it is best to produce when in state $(x, \mathbf{y} + \mathbf{e}_1)$, but best to idle when in state $(x, \mathbf{y} + \mathbf{e}_2)$.

3.1.2. Systems with Order Cancelations

In some settings, customers may cancel their orders after they have been announced. For example, consider a situation where, with each update, an order is either canceled or its due date is updated. In this case, ADI is imperfect with regard to both timing and realization of future orders. For example, in the context of a building construction project, changes to building specifications at some stage of the project may lead the builder to cancel orders for certain material. To incorporate this into the model, let p_i now denote the probability that an order is canceled at the end of its i th stage. The case where $p_i = 0$ corresponds to a system with no cancelations. The state space, action space, and cost rates are as in a system without order cancelations. To handle cancelations, we need only replace T'_i in (5) by

$$\begin{aligned} \tilde{T}'_i v(x, \mathbf{y}) := & (1 - p_i)T'_i v(x, \mathbf{y}) \\ & + p_i v(x, \mathbf{y} - \mathbf{e}_i \mathbb{I}_{\{y_i \geq 1\}}) \quad i = 1, \dots, k. \end{aligned}$$

See Section S-4 of the Supporting Information for a proof that Theorem 1 holds for systems with order cancelations under the additional assumption that $v_i p_i$ is nonincreasing in i .

3.1.3. Systems with two Customer Classes

In some situations, it may be the case that not all customers provide ADI. In other words, there may be a fraction of customers that does not announce orders ahead of demand. This is plausible in settings where the supplier has a mix of long-term and short-term (or nonrecurring) customers. Long-term customers are more likely to share information and to invest in the necessary infrastructure. Suppose that a fraction η of orders provides ADI, and that a fraction $1 - \eta$ does not. Equivalently, we may view customers as belonging to two separate classes. Class 1, with arrival rate $\lambda_1 = \eta\lambda$, provides ADI, and Class 2, with arrival rate $\lambda_2 = (1 - \eta)\lambda$, does not. To incorporate this into the model, we replace the operator T_λ by

$$\tilde{T}_\lambda v(x, \mathbf{y}) := \eta T_\lambda v(x, \mathbf{y}) + (1 - \eta)v(x - 1, \mathbf{y}).$$

The operator \tilde{T}_λ preserves conditions (C1)–(C4) because T_λ does; see the proof of Proposition 1. Hence, Theorem 1 holds in this setting.

3.1.4. Systems with Lost Sales

So far we have assumed that when orders become due and there is no on-hand inventory, orders can wait. In many applications, orders cannot wait and instead are lost if they become due and cannot be filled immediately. With each lost order, a lost sales cost is incurred. This cost may be a negotiated penalty with the customer or may reflect the cost of expediting the order or fulfilling it from an outside supplier (it may also reflect the loss of good will). To incorporate lost sales, we need only take the state space to be $\mathbb{Z}_+ \times \mathbb{Z}_+^k(m)$, re-define the cost function to be $c(x) = hx$, and replace the operator T'_k by \tilde{T}'_k defined as follows:

$$\begin{aligned} \tilde{T}'_k v(x, \mathbf{y}) := & v([x - \mathbb{I}_{\{y_k \geq 1\}}]^+, \mathbf{y} - \mathbf{e}_k \mathbb{I}_{\{y_k \geq 1\}}) \\ & + c_{\text{LS}} \mathbb{I}_{\{y_k \geq 1 \text{ and } x=0\}}, \quad (9) \end{aligned}$$

where c_{LS} corresponds to the lost sale cost per unit. Section S-5 of the Supporting Information contains a proof that Theorem 1 [except property (b), which pertains to backorders] holds in this setting.

3.2. Unbounded Jump Rates

In this section, we again consider the basic IDD system under the discounted-cost criterion. The model is identical to that considered in the previous section, except that here we do not place the bound m on the number of announced orders in the system. There are no rejected orders, and all arrivals enter the demand leadtime system. The state space is now $\mathcal{S} := \mathbb{Z} \times \mathbb{Z}_+^k$. Without the bound m , we have a continuous-time Markov decision process with unbounded transition rates. In particular, the conditional rate of transitions out of state $(x, \mathbf{y}) \in \mathcal{S}$ under action $a \in \{0, 1\}$ is $\lambda + \sum_{i=1}^k v_i y_i + \mu \mathbb{I}_{\{a=1\}}$. With no bound m on $\sum_{i=1}^k y_i$, this conditional rate is not bounded.

Theorem 2 below shows that the results in Theorem 1 also hold in the setting with unbounded jump rates (and discounted costs). We conjecture that similar results for unbounded jump rates hold under the average-cost criterion, but we do not have a proof at this time. Although this may not be surprising because Theorem 1 holds for any finite m , it is important to highlight that the presence of unbounded transition rates poses a technical challenge. In particular, it is not possible to apply uniformization to a problem with unbounded jump rates, and hence the problem cannot be transformed into an

“equivalent” discrete-time problem as in Section 3.1. Such a transformation is typically a crucial step for proving structural properties of optimal policies using inductive approaches (as in the proof of Proposition 1). An additional difficulty is that only recently has there developed a theory for problems with both unbounded jump rates and unbounded cost rates that characterizes the value function as a particular solution of the optimality equation, and ensures the existence of stationary optimal policies. See Guo and Hernández-Lerma [15]—hereafter called GH—for results and references.

Our proof of Theorem 2 establishes the structure of an optimal policy and of the value function for the problem with unbounded jump rates by letting m grow to infinity through a particular sequence of problems such as those considered in Section 3.1. In doing so, there are a number of technical points, such as the existence of various limits, that must be treated with care. The approach may provide a template that could be used for analyzing other CTMDPs with unbounded jump rates and cost rates.

Let v^* denote the value function of the problem with unbounded jump rates. The optimality equation for the problem with unbounded jump rates is $v = Lv$, where L is given by

$$Lv(x, \mathbf{y}) := \frac{1}{Q(\mathbf{y})} \left[c(x) + \lambda v(x, \mathbf{y} + \mathbf{e}_1) + \sum_{i=2}^k v_{i-1} y_{i-1} v(x, \mathbf{y} + \mathbf{e}_i - \mathbf{e}_{i-1}) + v_k y_k v(x - 1, \mathbf{y} - \mathbf{e}_k) + \mu \min\{v(x, \mathbf{y}), v(x + 1, \mathbf{y})\} \right]$$

and $Q(\mathbf{y}) := \beta + \lambda + \sum_{i=1}^k v_i y_i + \mu$.

In preparation for the main theorem of the section, define function $R(\cdot)$ by

$$R(x, \mathbf{y}) := |x| + \sum_{i=1}^k y_i \quad (10)$$

and consider the set of functions $B_R(\mathcal{S}) := \{v : \text{there exist constants } c_1, c_2 \geq 0 \text{ so that } |v(x, \mathbf{y})| \leq c_1 + c_2 R(x, \mathbf{y}) \text{ for all } (x, \mathbf{y}) \in \mathcal{S}\}$ where $R(x, \mathbf{y})$ is given in (10). To employ the theory of GH, we must identify a non-negative function R suitable for defining $B_R(\mathcal{S})$. GH do not specify an R for the use of their theory. “Suitable” means that R must satisfy some conditions that relate to the cost and transition rates of the CTMDP. In Lemma S-2 in Section S-6 of the Supporting Information we verify that our choice of R in (10) is indeed suitable. For convenience, Section S-6 also summarizes results we use from GH.

The following is the main result of the section. A proof is in the appendix.

THEOREM 2: For the system with IDD and unbounded jump rates, the value function v^* is the unique function in $B_R(\mathcal{S})$ that solves the optimality equation $v = Lv$. Moreover, v^* satisfies conditions (C1)–(C4), and there exists a stationary state-dependent base-stock policy that is optimal. The base-stock levels satisfy the conditions in (a) and (b) in Theorem 1.

4. NUMERICAL RESULTS

In this section, we present results from a numerical study. The goal is to examine the benefits from using ADI, to assess the value of updating, and to compare the impact of having full versus partial ADI. The insights we obtain for production-inventory systems with continuous review complement results in the literature for systems with periodic review and deterministic leadtimes.

We use average cost instead of discounted cost, because average cost is independent of the initial state and the discount factor. In all cases we set $\rho = \lambda/\mu < 1$, so that Theorem S-1 applies (see Section S-2 of the Supporting Information). The holding-cost rate is $h = 10$ and the production rate is $\mu = 1$, unless stated otherwise. In the numerical study, we used values of m large enough that further increases in m would not alter the average costs at the level of accuracy shown in our tables. For each problem instance, we obtained the long-run average cost by solving the MDP using value iteration.

4.1. Benefits of ADI

To assess the benefit of ADI, we compare the optimal average cost, J_A , for a system with ADI to the optimal average cost, J_N , for a system with no ADI and obtain the percentage cost reduction PCR := $100 \times (J_N - J_A)/J_N$. The two systems are identical in all respects, except that in the system with no ADI, orders are not announced ahead of their due dates; rather, information about when orders enter the demand leadtime system and when they move from one stage to the next is withheld. Only departures from the last stage of the demand leadtime system are observed. In general, the distribution of the departure process from the demand leadtime system is different from the distribution of its arrival process. However, the arrival and departure processes in steady state have identical (Poisson with rate λ) distributions for systems with no bound m . This follows from the fact that the departure process from an M/G/ ∞ queue in steady state is a homogeneous Poisson process with the same rate as the exogenous input Poisson process. For systems with finite m , the departure process in steady state is closely approximated by a Poisson process when m is large.

Table 1. The percentage cost reduction (PCR) for a system with $k = 1$.

		ν									
		λ	0.01	0.02	0.05	0.1	0.2	0.5	1	2	5
$b = 10$	0.1	0.00	0.00	0.00	0.02	0.37	3.94	8.07	17.84	12.60	7.53
	0.2	0.00	0.00	0.06	0.64	2.36	8.39	14.11	18.41	11.74	6.84
	0.4	0.51	1.66	4.84	8.52	12.73	16.21	20.97	17.74	9.60	5.35
	0.6	0.84	2.45	6.13	9.46	12.03	12.82	9.81	1.38	1.08	0.81
	0.8	6.04	8.47	10.54	10.61	9.06	5.87	3.26	2.02	0.54	0.04
	0.9	9.02	9.25	7.58	5.37	3.03	1.21	1.17	1.03	0.99	0.92
$b = 50$	0.1	0.36	1.96	8.01	16.33	25.47	44.92	40.43	28.71	14.76	8.12
	0.2	3.06	6.18	12.88	20.44	28.61	38.63	28.61	13.15	0.38	0.29
	0.4	5.15	7.57	12.41	17.58	22.52	17.64	15.29	10.67	5.17	2.74
	0.6	4.24	7.58	12.88	15.46	14.77	7.77	4.75	1.51	0.86	0.47
	0.8	11.19	13.77	13.29	9.87	5.95	2.59	1.31	0.89	0.38	0.16
	0.9	11.73	12.78	5.30	2.74	1.14	0.44	0.23	0.20	0.13	0.02
$b = 100$	0.1	8.93	14.28	23.84	32.73	49.49	51.82	39.08	23.15	6.63	0.05
	0.2	0.13	1.08	6.14	13.96	25.43	18.72	6.63	5.90	3.11	1.55
	0.4	2.08	4.78	10.96	16.58	19.37	15.63	5.06	4.07	2.06	1.04
	0.6	5.85	9.72	15.17	16.82	13.96	6.95	3.14	1.89	0.86	0.52
	0.8	12.93	14.98	12.63	8.04	4.42	1.64	0.83	0.46	0.24	0.13
	0.9	9.85	10.09	6.98	4.01	2.03	0.41	0.23	0.20	0.12	0.03

Long-run average cost is unaffected by the transient behavior of the demand leadtime system. As the demand leadtime system approaches steady state, its departure process converges in distribution to a Poisson process with rate λ . Hence, we may, for the purpose of computing long-run average cost for a system without ADI, assume arrivals to the system without ADI form a Poisson process. Finally, we note that a (state-independent) base-stock policy is optimal for a system with no ADI with Poisson arrivals and exponential production times; see, e.g., Veatch and Wein [32].

Representative numerical results comparing systems with and without ADI can be found in Table 1, where PCR is shown for varying values of parameters ν , λ , and b . The results are shown for a system with a single stage ($k = 1$). The effect of multiple update stages is discussed in Section 4.2.

The effect of ν on PCR, when all other parameter values are fixed, is not monotonic, with PCR initially increasing in the mean demand leadtime $1/\nu$ and then decreasing. ADI offers the greatest benefit in terms of PCR when the size of $1/\nu$ is moderate. The percentage cost reduction is relatively small when either $1/\nu$ is very large or very small. This can be explained as follows. When $1/\nu$ is small, the mean time between when an order is announced and when it becomes due is small. Hence, the information is of little use. When $1/\nu$ is large, the mean of the time between an order's announcement and due date is large, but so is the variance. This makes the information about future demand relatively less useful. The meanings of "large," "small," and "moderate" $1/\nu$ depend upon the value of λ . Although the joint effect of λ , μ , and ν on PCR is complicated, it appears that the value of $1/\nu$ that maximizes PCR for a given λ is increasing in λ . The

largest value of $1/\nu$ shown in Table 1 is $1/\nu = 100$; however, computations for larger values support the claim that the relative benefit of ADI is small for large $1/\nu$. For instance, with $b = 100$, $\lambda = 0.8$, and $1/\nu = 500$, we find that PCR is 3.22. These results highlight an important insight: having earlier notice of future orders may not always be desirable since the quality of this information tends also to deteriorate (i.e., the variance in the demand leadtime increases when the average demand leadtime increases). In our model, this is due to the fact that demand leadtime is assumed to have the exponential distribution. However, this also captures the fact that in practice the earlier an order is announced, the less reliable will be the estimate of its due date (see Section 4.2 for further results and discussion for systems with multiple stages of updating).

For each fixed ν , the effect of λ on PCR is also not monotonic. For fixed ν , the relative benefit of ADI is small when λ is large (close to $\mu = 1$). When λ is large, the optimal policy with or without ADI is for the production facility to produce most of the time. Hence, the availability of ADI makes little difference for the decisions taken. When λ is small, the absolute cost reduction from ADI is small, because costs in the systems with and without ADI both approach zero as $\lambda \downarrow 0$, but the value of PCR depends on the value of ν ; see Section S-7 of the Supporting Information for further discussion on this.

The effect of the ratio b/h is also not monotonic, with the value of PCR relatively small when b/h is either small or large. When b/h is small, ignoring ADI and producing to order (i.e., holding little or no inventory in anticipation of future demand) carries a relatively small penalty. When b/h is large, the base-stock levels are high for systems both with

and without ADI, and the probability of backorders is relatively small in both systems. Hence, ADI becomes relatively less useful.

We conclude this section by comparing the preceding observations with those obtained by Gavirneni et al. [11], who also evaluated the benefit of ADI with respect to similar parameters, but in a different context. They study a periodic review system with zero leadtimes and limited replenishment capacity per period, where ADI is obtained by having the supplier observe the demand of a retailer that uses an (s, S) ordering policy. Some of their qualitative insights (for example, regarding the effect of the ratio b/h) are similar to those above. However, there are some notable differences. For example, they found that the percentage cost reduction due to ADI is increasing in capacity. Interestingly, Özer and Wei [25], who consider a different model of capacitated periodic-review inventory systems with ADI, concluded the opposite. That is, in their modeling framework, they found ADI to be most beneficial when capacity is tight. Both of these findings can be contrasted to the effect of varying λ (which varies production system loading) that we describe above.

Some effects observed in both Gavirneni et al. [11] and in our study appear to have different causes in the different settings. For instance, they find that long demand leadtimes (measured in their case by the difference $S - s$) diminish the value of ADI, and they attribute this to the fact that long leadtimes result in large orders from the retailer, forcing the supplier to build inventory over time because of capacity limits. In our case, long demand leadtimes also diminish the value of ADI, but for a different reason: the information regarding due dates becomes less reliable because both the mean and the variance of demand leadtime increase simultaneously.

4.2. Benefits of Updating

In settings where the demand leadtime system consists of multiple stages, the supplier and customer may have a choice of how much information is shared. For example, should the customer inform the supplier as soon as an order enters the first stage or wait until an order has progressed further before forwarding the information to the supplier? Similarly, should the customer update the supplier each time an order enters a new stage or should it wait until the order has passed a specified number of stages? These questions are relevant when there is a cost associated with collecting the information, transmitting it from one party to another, and then making decisions based on it. To explore the benefit of full versus partial information sharing, we consider a system where the demand leadtime has two stages ($k = 2$) and compare the performance of this system when there is full ADI (information is shared as soon as orders enter the first stage and as they leave one stage and enter the next) to its performance when

there is no ADI and when there is only partial ADI (information is shared only when orders enter the second stage). The systems can be viewed as identical in all respects except for the number of update stages, with full ADI corresponding to $k = 2$, partial ADI to $k = 1$, and no ADI to $k = 0$. In the system with full ADI, the order is announced and then progresses through two stages of update, each exponentially distributed with mean $1/\nu$, before becoming due. In the system with partial ADI, the order is announced and then progresses through a single stage of update, exponentially distributed with mean $1/\nu$, before becoming due. In the system with no ADI, the order is announced and becomes due immediately.

Representative numerical results are displayed in Table 2, which not surprisingly shows that full ADI is superior. (A proof of this observation follows by noting that any policy for the system with partial ADI can be reproduced for the system with full ADI by basing decisions in the latter only on the net inventory and the number of orders in the second stage.) Additional numerical results for $b = 10$ and $b = 50$ can be found in Section S-9 of the Supporting Information. The value of full ADI is most significant when both ν and λ are in the mid-range, and least significant when both ν and λ are either very small or very large, as in the upper left and lower right corners of Table 2. This is consistent with results from Section 4.1. In most of the examples considered, the incremental benefit from full ADI ($k = 2$) over partial ADI ($k = 1$) is small. This suggests that partial ADI may be sufficient if updating is expensive to implement.

We close this section by noting a subtle difference between the effect of increasing ADI by increasing the number of stages observable to the supplier and increasing ADI by increasing the length of a particular stage. Compare a system with k stages in which each stage has mean $1/\nu$ to a system with a single stage with mean k/ν . Both systems have the same overall mean, k/ν , but the variance of the system with k stages is k/ν^2 while the one for the system with a single stage is k^2/ν^2 (i.e., k times larger). This helps explain why observing more stages of the demand process is always beneficial, but increasing the average length of a particular stage may not be.

4.3. Benefits of Full versus Partial ADI

As discussed at the end of Section 3.1, there are settings where ADI is not available from all customers. An important question that arises in these settings is how beneficial is it to increase the fraction of customers that provide ADI. In particular, is there a diminishing value to increasing this fraction or is the marginal benefit from ADI insensitive to how many customers already provide ADI? To address this question, we consider the version of our problem where there are two customer classes described at the end of Section 3.1. Class 1, with demand rate $\lambda_1 = \eta\lambda$, provides ADI and class 2, with

Table 2. Average cost and percentage cost reduction (PCR) for systems with $b = 100$. The columns labeled “ $k = 0$ ” show the average cost for systems without ADI.

$v_1 = v_2$	$\lambda = 0.4$			$\lambda = 0.6$			$\lambda = 0.8$		
	$k = 0$	$k = 1$	$k = 2$	$k = 0$	$k = 1$	$k = 2$	$k = 0$	$k = 1$	$k = 2$
0.01	25.07	24.55 2.08%	24.53 2.14%	46.38	43.67 5.85%	43.65 5.90%	107.24	93.38 12.93%	92.10 14.12%
0.02	25.07	23.87 4.78%	23.85 4.86%	46.38	41.87 9.72%	41.74 10.00%	107.24	91.18 14.98%	87.81 18.12%
0.05	25.07	22.32 10.96%	22.26 11.21%	46.38	39.35 15.17%	38.47 17.06%	107.24	93.70 12.63%	86.93 18.94%
0.10	25.07	20.91 16.58%	20.40 18.62%	46.38	38.58 16.82%	36.60 21.10%	107.24	98.62 8.04%	91.50 14.69%
0.20	25.07	20.21 19.37%	19.00 24.20%	46.38	39.91 13.96%	36.33 21.68%	107.24	102.51 4.42%	98.04 8.58%
0.50	25.07	21.15 15.63%	18.60 25.80%	46.38	43.16 6.95%	39.82 14.14%	107.24	105.49 1.64%	103.47 3.52%
1.00	25.07	23.80 5.06%	20.62 17.75%	46.38	44.93 3.14%	43.05 7.18%	107.24	106.35 0.83%	105.48 1.65%
1.50	25.07	23.90 4.66%	22.65 9.64%	46.38	45.38 2.16%	43.93 5.30%	107.24	106.68 0.53%	105.90 1.26%
2.00	25.07	24.05 4.07%	23.63 5.74%	46.38	45.51 1.89%	44.89 3.22%	107.24	106.75 0.46%	106.33 0.85%

demand rate $\lambda_2 = (1 - \eta)\lambda$, does not. We examine the effect of increasing the fraction of customers with ADI by varying the fraction η while maintaining $\lambda = \lambda_1 + \lambda_2$ constant, so that higher values of η correspond to more customers providing of ADI.

Representative results from numerical experiments are shown in Fig. 2 (with $k = 1$, $\lambda = 0.8$, and $b = 100$). As we can see in this example, the relative benefit of ADI does not exhibit diminishing returns with increases in the fraction η of customers with ADI. (The “nearly linear” pattern in the figure is present for other parameter settings as well. In some cases it is less pronounced.) This is in contrast to the typical effect of updating. These differences might, in part, be due to the fact that with additional updating we provide more information for the same customers while with expanding ADI to additional customers we provide new information for different customers. This is significant since we assume that customers announce orders independently of each other, so having information on some customers does not provide information on when other customers might announce their own orders. A managerial implication from these observations is that, all else being equal, and given our independence assumptions, a supplier may be better off expanding ADI (with limited or no updating) to more of its customers than obtaining more updates from those customers that already provide ADI.

These results appear to be different from those reported in the literature for systems with periodic review. For example, Özer and Wei [25] carried out a set of experiments where they varied the number of periods ahead of due dates that demand is announced (they refer to this as the information horizon),

as well as the effective fraction of customers that announce their demand ahead of due dates. Their results show a diminishing marginal cost reduction from increasing the fraction of customers whose demand is announced ahead of its due date. These differences might be due to the fact that for a continuous review system, production decisions are made for one unit at a time. Consequently, the system manager is able to use information about the anticipated due date of each order to make a decision about whether or not to initiate the production of a replenishment unit.

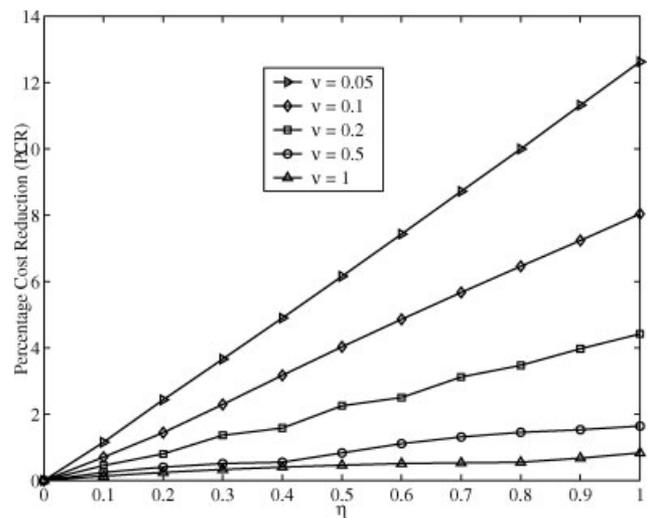


Figure 2. PCR versus fraction of customers that provide ADI.

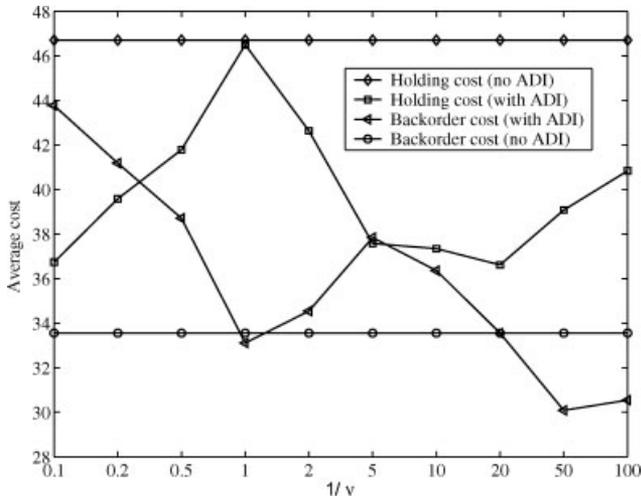


Figure 3. Average holding and backorder costs, with and without ADI.

4.4. Benefits of ADI to the Customer

In evaluating the benefit of ADI, we have so far taken the perspective of the supplier who manages the production process. In this section, we consider the impact of ADI on the customer who provides it. In particular, we address the question of whether or not both supplier and customer benefit from sharing demand information. It is often argued that ADI can reduce costs of the supplier (which we observed to be true) and improve the quality of service received by the customer. In our setting, the latter assertion would mean that customers experience fewer backorders and shorter fulfillment delays. In Fig. 3, we show the breakdown of supplier cost in terms of inventory holding and backorder cost for examples with $b = 50$. (By Little’s Law, the average waiting time of customers is proportional to the average backorder cost, so average waiting times can be deduced from Fig. 3. We emphasize that “waiting time” here refers to the span of time between when an order becomes due and when the order is satisfied. With ADI, an order becomes due when it exits the demand leadtime system.) As we can see, ADI does not always reduce the backorder costs. In some cases, the supplier uses ADI to reduce inventory holding cost at the expense of backorder cost. Generally, whether backorder cost, holding cost, or both decrease (both cannot increase) depends upon problem parameters. Hence, there is no guarantee that sharing ADI will lead to improved service levels to the customers.

This, of course, raises the question of why a customer would be willing to provide ADI only to see service levels suffer. One possible answer is that in practice customers who provide ADI also require a contractual agreement that service levels be improved or, alternatively, that the penalties for poor service be increased. For example, the customer could offer

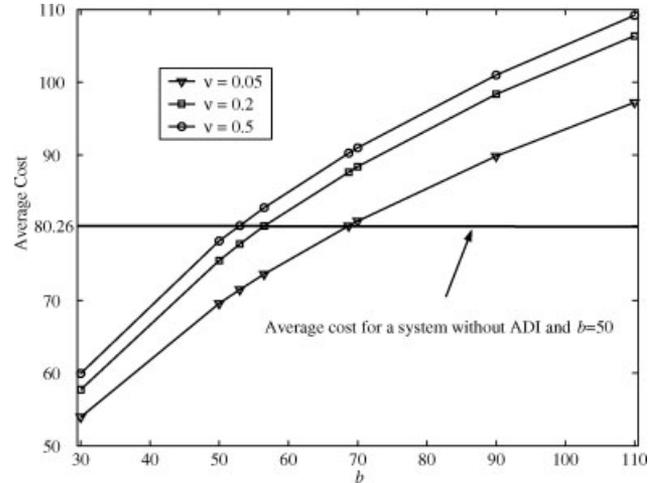


Figure 4. Average cost with ADI versus unit backorder cost rate.

ADI, but simultaneously increase the penalty for backorders. In Fig. 4, for a system with a single stage and $\lambda = 0.8$, we illustrate the impact on the cost of the supplier of having customers simultaneously provide ADI and increase unit backorder costs. The figure shows the average cost for the system with no ADI and $b = 50$, as well as the average cost for the system with ADI for different values of b . Figure 5 shows the average number of backorders as a function of b for systems with ADI. Of course, average costs increase and backorder levels decrease as the backorder penalty increases.

For given b and v there is a value $b^{\max} = b^{\max}(b, v)$ such that the supplier is indifferent to not receiving ADI while using backorder cost b and receiving ADI with mean demand leadtime v^{-1} while using backorder cost b^{\max} . From Fig. 4,

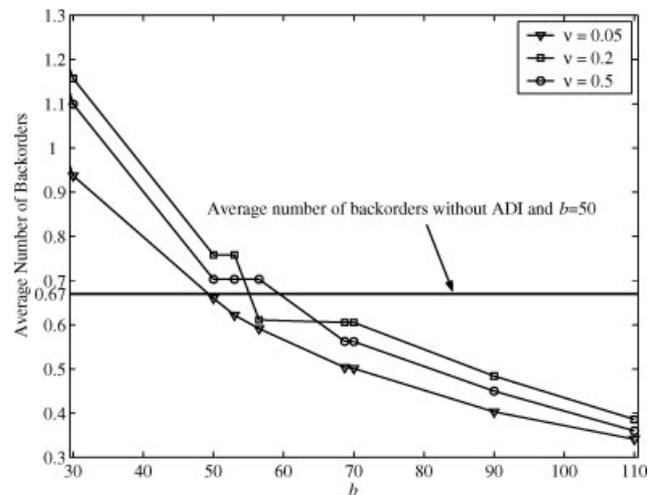


Figure 5. Average number of backorders versus unit backorder cost rate.

it can be seen that, for instance, if $\nu = 0.05$ then the supplier is indifferent between operating without ADI at $b = 50$ and operating with ADI at $b^{\max}(50, 0.05) \approx 69$. [The figure also shows that $b^{\max}(50, 0.2) \approx 57$ and $b^{\max}(50, 0.5) \approx 53$.] In other words, in exchange for receiving ADI with $\nu = 0.05$, the supplier is willing to accept up to a 40% increase in the backorder penalty rate. Figure 5 shows that providing ADI with $\nu = 0.05$ in combination with an increase in the backorder cost rate to $b^{\max} \approx 69$ causes the average number of backorders to decrease to roughly 0.50 from the value of 0.67 obtained in the absence of ADI.

However, given a value of ν , even the maximum increase (from b to b^{\max}) in the backorder cost rate that the supplier will accept may not be sufficient to reduce backorders to the level found without ADI. For example, if $b = 50$ and $\nu = 0.5$, then the supplier is willing to increase the backorder cost to at most $b^{\max} \approx 53$ in exchange for ADI; see Fig. 4. However, Fig. 5 shows that even with a backorder cost of 53, the average number of backorders with ADI and $\nu = 0.5$ is about 0.70, which exceeds 0.67—the average number of backorders for the system without ADI and $b = 50$. Obviously, with backorder penalties there are additional financial transfers from the supplier to the customer, which may compensate for the lower service levels.

In practice, there may be other strategies available to customers to mitigate the negative impact of ADI. For example, customers could charge their suppliers a fee for the demand information they provide. Customers could also use the possibility of providing demand information to strengthen their bargaining position during price negotiation with their suppliers.

5. EXTENSION TO SYSTEMS WITH SEQUENTIAL DUE DATES

In this article, we have focused on a particular form of ADI updating, in which orders that have been announced are updated independently of each other. In this section we briefly discuss how our analysis can be extended to systems where the independent updating assumption does not hold. In particular, we consider a system in which ADI is revealed through a process where orders are updated and become due in the same order they are announced. We call this a system with sequential due dates (SDD). Systems with SDD are similar in all aspects to systems we have considered so far, except for how orders progress through the leadtime system. For example consider a supplier who produces a component for a manufacturer. The manufacturer's production process is a serial production line comprised of a series of workstations that process jobs on a first-come, first-served (FCFS) basis. If a workstation is busy, incoming jobs wait in its queue. The manufacturer informs the supplier each time it releases a job

to the line (this may correspond to the manufacturer receiving an order from its own customers) and updates the supplier each time a job completes processing at one of the workstations. The supplier uses this information to estimate when it will receive a delivery request from the manufacturer. Such a request coincides with a job arriving at the workstation where the component provided by the supplier is needed.

SDD may arise in settings other than manufacturing. Consider, for example, a supplier that produces a product sold through a single retailer, which continuously reviews its inventory and follows a (Q, r) ordering policy. This means that the retailer places an order for Q units each time its inventory position drops to r . The supplier has real-time access to retailer's POS data and is aware of the retailer's ordering policy. Each order placed with the retailer can be used by the supplier to update its estimate of the time until the next replenishment order. This updating process progresses through Q stages, culminating in the placement of an order for a single batch of Q units. From the perspective of the supplier, there is always exactly one announced order, whose due date is updated each time the retailer's inventory position changes. Once an order becomes due, another order is simultaneously announced and starts the same process.

In a system with SDD, the leadtime system can be viewed as a serial queueing system consisting of n servers (we refer to these as nodes). As we shall illustrate with several examples shortly, the demand leadtime system could describe the internal processes of the customers that are observable to the supplier. External arrivals to the demand leadtime system occur at node 1 and progress sequentially through nodes $1, \dots, n$. Completion of service at the n th node corresponds to an order becoming due. Service time at node i consists of k_i stages, with the duration of each stage being exponentially distributed with mean $1/v_{ij}$ for the j th stage at node i . That is, service times at node i have a phase-type distribution with k_i phases in series. Special cases include the Erlang distribution where all the phases have the same mean and the exponential distribution where the number of phases is equal to one. Orders at each server are processed one at a time on a FCFS basis.

Inter-arrival times to the demand leadtime system have a phase type distribution, with one or more phases in series. To describe this arrival process, we introduce an additional node (node 0) and let k_0 denote the number of phases (or stages) associated with this node. The j th stage in node 0 has the exponential distribution with mean $1/v_{0j}$ for $j = 1, \dots, k_0$. Special cases are a Poisson arrival process or an arrival process with Erlang inter-arrival times.

The state of the system is described by the pair (x, \mathbf{y}) where the scalar x represents the net inventory level and $\mathbf{y} = (y_{ij} : i = 0, \dots, n, j = 1, \dots, k_i)$ represents the state of the demand leadtime system. For $i \neq 0$, y_{ij} represents the number of orders in node i that are in stage j . For stage 1,

the state variable y_{i1} indicates the number of orders that are either waiting for service at node i or have initiated the first stage of service at node i . Therefore, y_{i1} is a non-negative integer that can be arbitrarily large. For $j \neq 1$, y_{ij} indicates whether or not there is an order that has initiated the j th stage of service. Since servers process units one at a time, there can be at most one order at a time in stage $j \neq 1$. Hence, y_{ij} is either 0 or 1 for $j = 2, \dots, k_i$ and $\sum_{j=2}^{k_i} y_{ij} \in \{0, 1\}$. For node $i = 0$, y_{0j} is either 0 or 1 and $\sum_{j=1}^{k_0} y_{0j} = 1$, since there is exactly one order at node 0 at all times. In summary, the state space is $\hat{S} := \mathbb{Z} \times \mathbb{Y}$, where $\mathbb{Y} := \mathbb{Y}_0 \cap \mathbb{Y}_1$, $\mathbb{Y}_0 := \{\mathbf{y} : y_{0j} \in \{0, 1\} \text{ for } j = 1, \dots, k_0 \text{ and } \sum_{j=1}^{k_0} y_{0j} = 1\}$, and $\mathbb{Y}_1 := \{\mathbf{y} : y_{ij} \in \mathbb{Z}_+ \text{ for } j = 1, \dots, k_i, i = 1, \dots, n \text{ and } \sum_{j=2}^{k_i} y_{ij} \in \{0, 1\} \text{ for } i = 1, \dots, n\}$.

It is possible to model a wide variety of settings through different combinations of the parameters k_i and n . We describe three examples below.

EXAMPLE 1: Consider the example described earlier where a supplier provides a component to a manufacture whose production system consists of a series of workstations. Let production times at the supplier be exponentially distributed. Let also the manufacturer produce on a make-to-order basis while facing a Poisson external demand process. The component provided by the supplier is used in the $(n + 1)$ th workstation and is expected to be delivered by the supplier as soon as an order goes through the first n workstations. The manufacturer shares information about when orders are released into the production system and when they complete an operation at any of the workstations. In our general framework, this system corresponds to $k_i = 1$ for $i = 0, \dots, n$.

EXAMPLE 2: Consider a system similar to the one described in Example 1, except that items at the manufacturer are now processed one unit at time with all the operations carried out on a single workstation (instead of a series of workstations). At any given time, there may be multiple orders waiting to be processed in the queue of the workstation, but at most one order undergoing processing. Each time the workstation completes an operation, the manufacturer informs the supplier. The manufacturer also informs the supplier each time a new order arrives to the workstation. This system corresponds to $k_0 = 1$ and $n = 1$.

EXAMPLE 3: Consider the example mentioned earlier of a supplier who has a single customer in the form of a retailer. Let the retailer face a Poisson demand process with rate ν . The retailer uses a (Q, r) ordering policy so that it places an order of size Q whenever its own inventory position (sum of inventory on order and inventory on hand less backorders) reaches r . The retailer's inventory position takes on values $Q + r, Q + r - 1, \dots, r + 1$ with transition times between

consecutive values being exponentially distributed with rate ν . If the supplier has access to the retailer's inventory position, then the supplier can use the information to update the expected time at which the retailer will place an order. From the perspective of the supplier there is exactly one announced order at a time and the time between updates (there is a total of Q updates) is exponentially distributed with rate ν . In this setting, one "unit" for the supplier is an order of size Q from the retailer and the production time for this unit is exponential with rate μ . This case corresponds to $k_0 = Q$ and $n = 0$ with $\nu_{0j} = \nu$ for $j = 1, \dots, k_0$.

Define \mathbf{e}_{ij} to be the vector with 1 in the (i, j) -position and zeros elsewhere, and let \mathbf{e}_{00} be the vector of zeros. Let \hat{V} be the set of real-valued functions on \hat{S} . We denote by $\hat{v}^* \in \hat{V}$ the value function of the MDP. That is, $\hat{v}^*(x, \mathbf{y})$ is the minimum expected total discounted cost, given the system starts in state (x, \mathbf{y}) . Let $\hat{\gamma} := \mu + \sum_{i=0}^n \sum_{j=1}^{k_i} \nu_{ij}$. The optimality equation is $v = \hat{T}v$ where $\hat{T} : \hat{V} \rightarrow \hat{V}$ is defined by

$$\hat{T}v(x, \mathbf{y}) := \hat{\gamma}^{-1} \left[c(x) + \sum_{i=0}^n \sum_{j=1}^{k_i} \nu_{ij} T_{ij}v(x, \mathbf{y}) + \mu T_\mu v(x, \mathbf{y}) \right]. \tag{11}$$

The operator T_μ corresponds to the production decision and is defined in Section 3.1. The operators $\{T_{ij}\}$ in (11) are defined in (12)–(17) below. The operators $\{T_{0j} : j < k_0\}$ correspond to transitions in the phase of the external arrival process, T_{0k_0} corresponds to an external arrival to the demand leadtime system or to an order coming due in case $n = 0$, $\{T_{ij} : i > 0, j < k_i\}$ correspond to a transition in the phase of a service time, and $\{T_{ik_i} : i > 0\}$ correspond to a transition of an order between nodes (for $i < n$) and to an order coming due (for $i = n$). We have

$$T_{0j}v(x, \mathbf{y}) := v(x, \mathbf{y} + [\mathbf{e}_{0,j+1} - \mathbf{e}_{0j}]\mathbb{I}_{\{y_{0j}=1\}}) \quad j = 1, \dots, k_0 - 1 \tag{12}$$

$$T_{0k_0}v(x, \mathbf{y}) := \begin{cases} v(x, \mathbf{y} + [\mathbf{e}_{01} - \mathbf{e}_{0k_0} + \mathbf{e}_{11}]\mathbb{I}_{\{y_{0k_0}=1\}}) & \text{when } n > 0 \\ v(x - \mathbb{I}_{\{y_{0k_0}=1\}}, \mathbf{y} + [\mathbf{e}_{01} - \mathbf{e}_{0k_0}]\mathbb{I}_{\{y_{0k_0}=1\}}) & \text{when } n = 0. \end{cases} \tag{13}$$

$$T_{i1}v(x, \mathbf{y}) := v(x, \mathbf{y} + [\mathbf{e}_{i2} - \mathbf{e}_{i1}]\mathbb{I}_{\{y_{i1} \geq 1 \text{ and } y_{i\ell}=0 \text{ for } \ell=2, \dots, k_i\}}) \quad i = 1, \dots, n \text{ when } k_i \geq 2 \tag{14}$$

$$T_{ij}v(x, \mathbf{y}) := v(x, \mathbf{y} + [\mathbf{e}_{i,j+1} - \mathbf{e}_{ij}]\mathbb{I}_{\{y_{ij} \geq 1\}}) \quad i = 1, \dots, n; j = 2, \dots, k_i - 1 \tag{15}$$

$$T_{ik_i}v(x, \mathbf{y}) := v(x, \mathbf{y} + [\mathbf{e}_{i+1,1} - \mathbf{e}_{ik_i}]\mathbb{I}_{\{y_{ik_i} \geq 1\}}) \quad i = 1, \dots, n - 1 \tag{16}$$

$$T_{nk_n}v(x, \mathbf{y}) := v(x - \mathbb{I}_{\{y_{nk_n} \geq 1\}}, \mathbf{y} - \mathbf{e}_{nk_n}\mathbb{I}_{\{y_{nk_n} \geq 1\}}). \tag{17}$$

Note that T_{i1} for $i = 1, \dots, n$ is defined by (16)–(17) when $k_i = 1$.

In preparation for our main result of this section, we introduce an ordering on the space of (node, phase)-indices. For $(p, q), (r, w) \in \{(i, j) : i = 0, \dots, n; j = 1, \dots, k_i\} \cup \{(0, 0)\}$, we define $(p, q) < (r, w)$ to mean that one of the following two conditions holds: (i) $p < r$, or (ii) $p = r$ and $q < w$. Intuitively, $(p, q) < (r, w)$ means that an order at stage w of node r is closer to being due than an order at stage q of node p . We will use the following analogs of Conditions (C1)–(C4).

- ($\hat{C}1$) $\Delta v(x, \mathbf{y}) \leq \Delta v(x + 1, \mathbf{y})$ for all $(x, \mathbf{y}) \in \hat{\mathcal{S}}$.
- ($\hat{C}2$) $\Delta v(x, \mathbf{y} + \mathbf{e}_{pq}) \leq \Delta v(x + 1, \mathbf{y} + \mathbf{e}_{rw})$ for all (x, \mathbf{y}) and $(p, q) < (r, w)$ such that $\mathbf{y} + \mathbf{e}_{pq}, \mathbf{y} + \mathbf{e}_{rw} \in \mathbb{Y}$.
- ($\hat{C}3$) $\Delta v(x, \mathbf{y} + \mathbf{e}_{rw}) \leq \Delta v(x, \mathbf{y} + \mathbf{e}_{pq})$ for all (x, \mathbf{y}) and $(p, q) < (r, w)$ such that $\mathbf{y} + \mathbf{e}_{pq}, \mathbf{y} + \mathbf{e}_{rw} \in \mathbb{Y}$.
- ($\hat{C}4$) $\Delta v(x, \mathbf{y}) \leq 0$ for all $(x, \mathbf{y}) \in \hat{\mathcal{S}}$ with $x < 0$.

When $n \geq 1$ and $k_0 \geq 2$ we also use Condition ($\hat{C}5$) below, which is related to the “arrival” node, 0. Condition ($\hat{C}5$) is needed to ensure that Condition ($\hat{C}3$) is preserved by \hat{T} ; see the proof of Proposition S-1 in the Supporting Information. If $n = 0$ or $k_0 = 1$, then Condition ($\hat{C}5$) is vacuous.

- ($\hat{C}5$) $\Delta v(x, \mathbf{y} + \mathbf{e}_{01} + \mathbf{e}_{11}) \leq \Delta v(x, \mathbf{y} + \mathbf{e}_{0q})$ for all (x, \mathbf{y}) and $q \geq 2$ such that $\mathbf{y} + \mathbf{e}_{0q}, \mathbf{y} + \mathbf{e}_{01} + \mathbf{e}_{11} \in \mathbb{Y}$.

The next theorem, which describes the structure of optimal policies for SDD systems, is the main result of this section. We give a proof in Section S-8 of the Supporting Information. The argument parallels the proof of Theorem 1 that is detailed in Section 3.1 of the text and in Section S-1 of the Supporting Information. Specifically, we show that \hat{T} preserves Conditions ($\hat{C}1$)–($\hat{C}5$), which allows us to conclude that \hat{v}^* satisfies Conditions ($\hat{C}1$)–($\hat{C}5$), from which the theorem then follows.

THEOREM 3: The state-dependent base-stock policy $\hat{\pi}^* = \{\hat{\pi}^*(x, \mathbf{y})\}$ given by

$$\hat{\pi}^*(x, \mathbf{y}) := \begin{cases} 0 & \text{if } x \geq \hat{s}_{\mathbf{y}} \\ 1 & \text{if } x < \hat{s}_{\mathbf{y}} \end{cases} \quad (18)$$

where $\hat{s}_{\mathbf{y}} := \min\{x : \hat{v}^*(x + 1, \mathbf{y}) - \hat{v}^*(x, \mathbf{y}) \geq 0\}$ is optimal. In addition, (a) the base-stock levels satisfy $\hat{s}_{\mathbf{y} + \mathbf{e}_{rw}} \in \{\hat{s}_{\mathbf{y} + \mathbf{e}_{pq}}, \hat{s}_{\mathbf{y} + \mathbf{e}_{pq}} + 1\}$ for $(p, q) < (r, w)$ such that $\mathbf{y} + \mathbf{e}_{pq}, \mathbf{y} + \mathbf{e}_{rw} \in \mathbb{Y}$; and (b) $\hat{\pi}^*(x, \mathbf{y}) = 1$ if $x < 0$.

As we did for systems with independent updating, it is possible to evaluate the benefit of ADI with SDD features. Table S-3 in the Supporting Information contains

numerical results for systems with SDD, and shows the effects of parameters ν , k , and ρ .

6. CONCLUDING COMMENTS

In this article, we considered production-inventory systems where the production facility has access to ADI in the form of advance order announcements and subsequent updates. The ADI is not perfect because orders may become due before or after announced expected due dates, the time between updates is random, and announced orders may be canceled. In addition, only a fraction of the customers may provide ADI. We considered two schemes through which demand information is revealed, one in which due dates are independent and the other in which they are sequential. For each scheme, we formulated the production control problem as a continuous-time Markov decision process and showed that there is an optimal state-dependent base-stock policy, with base-stock levels that are non-decreasing in the number of announced orders at each stage of update. We also showed that the base-stock level increases by at most one unit with a unit increase in the number of orders at any stage.

In numerical experiments, we observed that the cost reduction to the supplier from the introduction of ADI is sensitive to the number of update stages and the length of each stage. Although adding more update stages is always beneficial, increasing the average length of stages may increase or decrease cost, with ADI being most valuable when the average stage length is moderate. We also observed that in many cases, much of the benefit of updating can be achieved with one update. Although ADI is always beneficial to the supplier, we observed that this may not be the case for the customers who provide the ADI. In some cases, the supplier uses ADI to reduce inventory at the expense of higher backorders. We showed that a possible remedy is for the customers to negotiate higher backorder penalties in exchange for ADI.

There are several avenues for future research. It would be of interest to consider systems where order sizes are variable and where the actual size of an order is not known exactly until the order becomes due. This would generalize the model with cancellations by assigning a probability distribution to orders that allows sizes other than zero or one. It would also be of interest to consider multiple customer classes with differing backorder costs. Although the problem would be made difficult by the need for a higher-dimensional state space, we expect there would again be an optimal state-dependent base-stock policy, where the state would include backorder levels of orders from each customer class. In addition to production, the policy would specify whether an order that becomes due should be satisfied from available inventory, if there is any, or backordered. This decision would of course depend on the backorder cost associated with the order’s class.

APPENDIX

PROOF OF THEOREM 1: Any stationary policy that uses for each $(x, \mathbf{y}) \in \mathcal{S}_m$ an action that attains the minimum in $Tv_m^*(x, \mathbf{y})$ is optimal. Hence, the policy that prescribes action $a = 1$ in states $\mathcal{S}_m^1 := \{(x, \mathbf{y}) \in \mathcal{S}_m : v_m^*(x + 1, \mathbf{y}) < v_m^*(x, \mathbf{y})\}$ and action $a = 0$ in states $\mathcal{S}_m^0 := \{(x, \mathbf{y}) \in \mathcal{S}_m : v_m^*(x + 1, \mathbf{y}) \geq v_m^*(x, \mathbf{y})\}$ is optimal. By Proposition 1, v_m^* satisfies Condition (C1), so π^* defined in (7) satisfies $\pi^*(x, \mathbf{y}) = 1$ for $(x, \mathbf{y}) \in \mathcal{S}_m^1$ and $\pi^*(x, \mathbf{y}) = 0$ for $(x, \mathbf{y}) \in \mathcal{S}_m^0$. Hence π^* is optimal.

To prove (a), note that by Proposition 1, the value function v_m^* satisfies conditions (C2) and (C3). Applying condition (C3) $l - j$ times and using the definition of $s_{\mathbf{y}+\mathbf{e}_j}$, we have $\Delta v_m^*(s_{\mathbf{y}+\mathbf{e}_j}, \mathbf{y} + \mathbf{e}_j) \geq \Delta v_m^*(s_{\mathbf{y}+\mathbf{e}_j}, \mathbf{y} + \mathbf{e}_j) \geq 0$, which implies $s_{\mathbf{y}+\mathbf{e}_j} \geq s_{\mathbf{y}+\mathbf{e}_j}$. Also, by condition (C2) and the definition of $s_{\mathbf{y}+\mathbf{e}_j}$, we have $\Delta v_m^*(s_{\mathbf{y}+\mathbf{e}_j} + 1, \mathbf{y} + \mathbf{e}_j) \geq \Delta v_m^*(s_{\mathbf{y}+\mathbf{e}_j}, \mathbf{y} + \mathbf{e}_j) \geq 0$, which implies $s_{\mathbf{y}+\mathbf{e}_j} + 1 \geq s_{\mathbf{y}+\mathbf{e}_j}$. Therefore, $s_{\mathbf{y}+\mathbf{e}_j} \leq s_{\mathbf{y}+\mathbf{e}_j} + 1$ and hence $s_{\mathbf{y}+\mathbf{e}_j}$ is equal to either $s_{\mathbf{y}+\mathbf{e}_j}$ or $s_{\mathbf{y}+\mathbf{e}_j} + 1$. Finally, part (b) is a consequence of the fact that v_m^* satisfies condition (C4).

An alternative proof of the optimality of a state-dependent base-stock policy and for part (a) can be obtained by casting the problem in terms of service rate control and using results on monotone optimal policies for continuous-time MDPs in Veatch and Wein [31]. \square

PROOF OF THEOREM 2: The first statement follows from Theorem 3.2 of GH. Lemma S-2 in the Supporting Information shows that the conditions needed to apply their theorem hold for our problem.

For each m , extend the domain of v_m^* from \mathcal{S}_m to \mathcal{S} by defining $v_m^*(x, \mathbf{y}) := 0$ for $(x, \mathbf{y}) \in \mathcal{S} \setminus \mathcal{S}_m$. To prove the remaining statements, we begin by showing that there exists a pointwise convergent subsequence of $\{v_m^*\}$. Lemma 1 below implies for each $(x, \mathbf{y}) \in \mathcal{S}$ that $\{v_m^*(x, \mathbf{y})\}$ is a bounded sequence of real numbers (note that the bound does *not* depend upon m). Hence, for each (x, \mathbf{y}) , the sequence $\{v_m^*(x, \mathbf{y})\}$ has a convergent subsequence in \mathbb{R} .

Let $\{z_1, z_2, z_3, \dots\}$ be an enumeration of the countable space \mathcal{S} [each z_i is some element (x, \mathbf{y}) of \mathcal{S}]. We now proceed with a diagonalization argument to construct the pointwise convergent subsequence of $\{v_m^*\}$. Let $\{m_{1,j} : j = 1, 2, \dots\}$ be such that $\lim_{j \rightarrow \infty} v_{m_{1,j}}^*(z_1)$ exists. Next, let $\{m_{2,j} : j = 1, 2, \dots\}$ be a subsequence of $\{m_{1,j} : j = 2, 3, \dots\}$ such that $\lim_{j \rightarrow \infty} v_{m_{2,j}}^*(z_2)$ exists. Note also that $\lim_{j \rightarrow \infty} v_{m_{2,j}}^*(z_1)$ exists, because $\{m_{2,j} : j = 1, 2, \dots\} \subseteq \{m_{1,j} : j = 2, 3, \dots\}$. Continuing in this fashion, we proceed sequentially to extract subsequences of subsequences so that for each n we have $\{m_{n,j} : j = 1, 2, \dots\} \subseteq \{m_{n-1,j} : j = n, n + 1, \dots\}$ and $\lim_{j \rightarrow \infty} v_{m_{n,j}}^*(z_i)$ exists for $i = 1, \dots, n$. Let $m_j := m_{j,j}$. It can now be seen that $\lim_{j \rightarrow \infty} v_{m_j}^*$ exists pointwise. (Alternatively, we may appeal to Tychonoff's Theorem to reach this conclusion; see, e.g., [1].) Let v^{**} denote the limit; that is, $v^{**} : \mathcal{S} \rightarrow \mathbb{R}$ is defined to be the function for which $\lim_{j \rightarrow \infty} v_{m_j}^*(x, \mathbf{y}) = v^{**}(x, \mathbf{y})$ for all $(x, \mathbf{y}) \in \mathcal{S}$.

Next, we show that the limit v^{**} is in fact the value function v^* for the problem with unbounded jump rates. To do so, it will be helpful to re-write the optimality equation (6) from Section 3.1 as $v = L_m v$ where operator L_m is given by

$$L_m v(x, \mathbf{y}) := \frac{1}{Q_m(\mathbf{y})} \left[c(x) + \lambda \mathbb{1}_{\{\bar{y} < m\}} v(x, \mathbf{y} + \mathbf{e}_1) + \sum_{i=2}^k v_{i-1} y_{i-1} v(x, \mathbf{y} + \mathbf{e}_i - \mathbf{e}_{i-1}) + v_k y_k v(x - 1, \mathbf{y} - \mathbf{e}_k) + \mu \min\{v(x, \mathbf{y}), v(x + 1, \mathbf{y})\} \right]$$

and $Q_m(\mathbf{y}) := \beta + \lambda \mathbb{1}_{\{\bar{y} < m\}} + \sum_{i=1}^k v_i y_i + \mu$. By rearranging terms, it can be checked that the equation $v = L_m v$ is equivalent to (6). Keep in mind that T in (6) depends upon m .

Naval Research Logistics DOI 10.1002/nav

For any function v on \mathcal{S} and any $(x, \mathbf{y}) \in \mathcal{S}$, observe that $L_m v(x, \mathbf{y}) = Lv(x, \mathbf{y})$ if $m > \bar{y}$. Hence, for any $(x, \mathbf{y}) \in \mathcal{S}$ it follows that

$$v^{**}(x, \mathbf{y}) = \lim_{j \rightarrow \infty} v_{m_j}^*(x, \mathbf{y}) = \lim_{j \rightarrow \infty} L_{m_j} v_{m_j}^*(x, \mathbf{y}) = \lim_{j \rightarrow \infty} L v_{m_j}^*(x, \mathbf{y}). \quad (\text{A1})$$

For any function v on \mathcal{S} and any $(x, \mathbf{y}) \in \mathcal{S}$ we next re-express $Lv(x, \mathbf{y})$. To this end, for given (x, \mathbf{y}) consider the continuous function $\mathcal{L}_{(x, \mathbf{y})} : \mathbb{R}^{k+3} \rightarrow \mathbb{R}$ defined by

$$\mathcal{L}_{(x, \mathbf{y})}(\varphi_1, \dots, \varphi_{k+3}) := \frac{1}{Q(\mathbf{y})} \left[c(x) + \lambda \varphi_1 + \sum_{i=2}^{k+1} v_{i-1} y_{i-1} \varphi_i + \mu \min\{\varphi_{k+2}, \varphi_{k+3}\} \right].$$

It can now be seen that

$$Lv(x, \mathbf{y}) = \mathcal{L}_{(x, \mathbf{y})}(v(x, \mathbf{y} + \mathbf{e}_1), v(x, \mathbf{y} + \mathbf{e}_2 - \mathbf{e}_1), \dots, v(x, \mathbf{y} + \mathbf{e}_k - \mathbf{e}_{k-1}), v(x - 1, \mathbf{y} - \mathbf{e}_k), v(x, \mathbf{y}), v(x + 1, \mathbf{y})).$$

For any sequence of functions $\{u_j\}$ with $u_j \rightarrow u$ pointwise, we have

$$\begin{aligned} \lim_{j \rightarrow \infty} L u_j(x, \mathbf{y}) &= \lim_{j \rightarrow \infty} \mathcal{L}_{(x, \mathbf{y})}(u_j(x, \mathbf{y} + \mathbf{e}_1), \\ &\quad u_j(x, \mathbf{y} + \mathbf{e}_2 - \mathbf{e}_1), \dots, u_j(x, \mathbf{y} + \mathbf{e}_k - \mathbf{e}_{k-1}), \\ &\quad u_j(x - 1, \mathbf{y} - \mathbf{e}_k), u_j(x, \mathbf{y}), u_j(x + 1, \mathbf{y})) \\ &= \mathcal{L}_{(x, \mathbf{y})}(u(x, \mathbf{y} + \mathbf{e}_1), \\ &\quad u(x, \mathbf{y} + \mathbf{e}_2 - \mathbf{e}_1), \dots, u(x, \mathbf{y} + \mathbf{e}_k - \mathbf{e}_{k-1}), \\ &\quad u(x - 1, \mathbf{y} - \mathbf{e}_k), u(x, \mathbf{y}), u(x + 1, \mathbf{y})) \\ &= Lu(x, \mathbf{y}). \end{aligned} \quad (\text{A2})$$

Note that in (A2), we may pass the limit inside $\mathcal{L}_{(x, \mathbf{y})}$ because $\mathcal{L}_{(x, \mathbf{y})}$ is continuous. Applying the preceding observation with $\{u_j\} = \{v_{m_j}^*\}$ and $u = v^{**}$ and using (A1), it follows that $v^{**}(x, \mathbf{y}) = Lv^{**}(x, \mathbf{y})$. Now, because (x, \mathbf{y}) was arbitrary, we see that $v^{**} = Lv^{**}$. That is, v^{**} solves the optimality equation for the problem with unbounded jump rates. Moreover, it can be seen from Lemma 1 that v^{**} is in $B_R(\mathcal{S})$ with $c_1 = \beta^{-2}(h + b)(\lambda + \mu)$ and $c_2 = \beta^{-1}(h + b)$. Therefore, it follows from the first part of the theorem that $v^{**} = v^*$; that is, v^{**} is the value function for the problem with unbounded jump rates. It can now readily be verified that $v^* = \lim_{j \rightarrow \infty} v_{m_j}^*$ satisfies conditions (C1) through (C4) with $\mathbb{Z}_+^k(m - 1)$ and $\mathbb{Z}_+^k(m)$ replaced by \mathbb{Z}_+^k .

Theorem 3.3 of GH implies that the stationary policy that produces in states $\mathcal{S}^1 := \{(x, \mathbf{y}) \in \mathcal{S} : v^*(x + 1, \mathbf{y}) < v^*(x, \mathbf{y})\}$ and that idles in states $\mathcal{S}^0 := \{(x, \mathbf{y}) \in \mathcal{S} : v^*(x + 1, \mathbf{y}) \geq v^*(x, \mathbf{y})\}$ is optimal for the problem with unbounded jump rates. By an argument identical to the proof of Theorem 1, such a policy is a state-dependent base-stock policy that satisfies (a) and (b). \square

LEMMA 1: $0 \leq v_m^*(x, \mathbf{y}) \leq \beta^{-1}(h + b)R(x, \mathbf{y}) + \beta^{-2}(h + b)(\lambda + \mu) < \infty$, where v_m^* is the value function for the problem with bounded jump rates in Section 3.1 and the function R is defined in (10).

PROOF: Fix $m < \infty$ and $(x, \mathbf{y}) \in \mathcal{S}_m$. To bound $v_m^*(x, \mathbf{y})$ from above, it suffices to obtain an upper bound on the expected discounted cost of using the policy π^+ that "always produces" [$\pi^+(x', \mathbf{y}') = 1$ for all $(x', \mathbf{y}') \in \mathcal{S}_m$]. To this end, we begin by developing an explicit construction of a version of the continuous-time Markov chain (CTMC) induced by π^+ . Suppose that $\{A_i : i = 1, 2, \dots\}$ is an i.i.d. sequence of uniform $[0, 1]$ random variables and that $\{E_i : i = 1, 2, \dots\}$ is an i.i.d. sequence of exponential random variables each with mean Λ^{-1} , independent of $\{A_i : i = 1, 2, \dots\}$. Let $\hat{E}_0 := 0$, $\hat{E}_n := \sum_{j=1}^n E_j$ for $n = 1, 2, \dots$, and $N(t) := \max\{n \geq 0 : \hat{E}_n \leq t\}$.

Recall the notation $\bar{\mathbf{y}}' = \sum_{i=1}^k y'_i$. Consider the function $f : \mathcal{S}_m \times [0, 1] \rightarrow \mathcal{S}_m$ given by

$$f((x', \mathbf{y}'), a) := \begin{cases} (x' + 1, \mathbf{y}') & \text{if } a \in [0, \mu/\Lambda] \\ (x', \mathbf{y}' + \mathbf{e}_1 \mathbb{1}_{\{\bar{\mathbf{y}}' < m\}}) & \text{if } a \in (\mu/\Lambda, (\lambda + \mu)/\Lambda] \\ (x', \mathbf{y}' + \mathbf{e}_{i+1} - \mathbf{e}_i) & \text{if } a \in \left(\left(\lambda + \mu + \sum_{j=1}^{i-1} v_j y'_j \right) / \Lambda, \right. \\ & \left. \left(\lambda + \mu + \sum_{j=1}^i v_j y'_j \right) / \Lambda \right] \\ & \text{for } i = 1, \dots, k-1 \\ (x' - 1, \mathbf{y}' - \mathbf{e}_k) & \text{if } a \in \left(\left(\lambda + \mu + \sum_{j=1}^{k-1} v_j y'_j \right) / \Lambda, \right. \\ & \left. \left(\lambda + \mu + \sum_{j=1}^k v_j y'_j \right) / \Lambda \right] \\ (x', \mathbf{y}') & \text{otherwise.} \end{cases}$$

For the fixed value $(x, \mathbf{y}) \in \mathcal{S}_m$, define

$$\begin{aligned} (X_0, \mathbf{Y}_0) &:= (x, \mathbf{y}) \\ (X_n, \mathbf{Y}_n) &:= f((X_{n-1}, \mathbf{Y}_{n-1}), A_n) \quad n = 1, 2, \dots \end{aligned} \quad (\text{A3})$$

and $(X(t), \mathbf{Y}(t)) := (X_{N(t)}, \mathbf{Y}_{N(t)})$. Note that $\{(X(t), \mathbf{Y}(t)) : t \geq 0\}$ has the distribution of the CTMC induced by π^+ , as desired.

For $n = 1, 2, \dots$ define

$$\begin{aligned} R_n &:= |\{j \in \{1, \dots, n\} : (X_j, \mathbf{Y}_j) = (X_{j-1} + 1, \mathbf{Y}_{j-1})\}| \\ U_n &:= |\{j \in \{1, \dots, n\} : (X_j, \mathbf{Y}_j) = (X_{j-1}, \mathbf{Y}_{j-1} + \mathbf{e}_1)\}| \\ D_n &:= |\{j \in \{1, \dots, n\} : (X_j, \mathbf{Y}_j) = (X_{j-1} - 1, \mathbf{Y}_{j-1} - \mathbf{e}_k)\}| \end{aligned}$$

where here $|\cdot|$ is set cardinality. When $k = 1$, it is possible to visualize R_n, U_n , and D_n by graphing the transitions of $\{(X(t), \mathbf{Y}(t))\}$ on a two-dimensional grid. Then, R_n is how many of the first n transitions are “to the right,” U_n is how many of the first n transitions are “up,” and D_n is how many of the first n transitions are “diagonal.” From the above definitions, note that

$$X_n = X_0 + R_n - D_n \quad \text{and} \quad \bar{\mathbf{Y}}_n = \bar{\mathbf{Y}}_0 + U_n - D_n. \quad (\text{A4})$$

(Again, $\bar{\mathbf{Y}}_n = \sum_{i=1}^k Y_{n,i}$.) For $n = 1, 2, \dots$ also define

$$\tilde{U}_n := |\{j \in \{1, \dots, n\} : A_j \in (\mu/\Lambda, (\lambda + \mu)/\Lambda)\}|.$$

By construction, we have

$$U_n \leq \tilde{U}_n \quad \text{and} \quad D_n \leq \bar{\mathbf{Y}}_0 + U_n. \quad (\text{A5})$$

Next we construct another process that is coupled with $\{(X(t), \mathbf{Y}(t))\}$. Consider the function $g : \mathbb{Z}_+ \times [0, 1] \rightarrow \mathbb{Z}_+$ given by

$$g(z, a) := \begin{cases} z + 1 & \text{if } a \in [0, (\lambda + \mu)/\Lambda] \\ z & \text{otherwise.} \end{cases}$$

Define

$$\begin{aligned} Z_0 &:= |x| + \bar{\mathbf{y}} \\ Z_n &:= g(Z_{n-1}, A_n) \quad n = 1, 2, \dots, \end{aligned}$$

where (x, \mathbf{y}) is as in (A3) and $Z(t) := Z_{N(t)}$. Observe that

$$Z_n = Z_0 + R_n + \tilde{U}_n = |x| + \bar{\mathbf{y}} + R_n + \tilde{U}_n. \quad (\text{A6})$$

Combining (A3)–(A6), we get

$$\begin{aligned} |X_n| &= |X_0 + R_n - D_n| \leq |X_0| + R_n + D_n \\ &\leq |X_0| + R_n + \bar{\mathbf{Y}}_0 + U_n \\ &\leq |X_0| + R_n + \bar{\mathbf{Y}}_0 + \tilde{U}_n \\ &= |x| + R_n + \bar{\mathbf{y}} + \tilde{U}_n \\ &= Z_n. \end{aligned}$$

Hence, $|X(t)| \leq Z(t)$. Next, define $c^\dagger(x) := (h + b)x$. Note that $c(x) \leq c^\dagger(|x|)$ and that $c^\dagger(\cdot)$ is increasing on \mathbb{Z}_+ . Therefore,

$$\begin{aligned} E_{(x, \mathbf{y})}^{\pi^+} \int_{t=0}^{\infty} e^{-\beta t} c(X(t)) dt &= \mathbb{E} \int_{t=0}^{\infty} e^{-\beta t} c(X(t)) dt \\ &\leq \mathbb{E} \int_{t=0}^{\infty} e^{-\beta t} c^\dagger(|X(t)|) dt \leq \mathbb{E} \int_{t=0}^{\infty} e^{-\beta t} c^\dagger(Z(t)) dt, \end{aligned}$$

where \mathbb{E} is expectation on the probability space upon which $\{A_i\}$ and $\{E_i\}$ are defined and where the initial state is (x, \mathbf{y}) . Hence, $\mathbb{E} \int_{t=0}^{\infty} e^{-\beta t} c^\dagger(Z(t)) dt$ is an upper bound on $v_n^*(x, \mathbf{y})$.

Regardless of m , the process $\{Z(t)\}$ has the following “dynamics”: $Z(0) = |x| + \bar{\mathbf{y}}$ and $Z(\cdot)$ remains in state (say) $z \in \mathbb{Z}_+$ an exponential amount of time with mean $1/(\lambda + \mu)$ before moving to state $z + 1 \in \mathbb{Z}_+$. The latter fact can be verified by conditioning on the geometric number of transitions made from z back to z by the embedded process $\{Z_n\}$. Direct calculations using value iteration [to compute the expected discounted cost accrued by $\{Z(t)\}$ through the time of its (say) j th jump to the right] and induction show that

$$\begin{aligned} \mathbb{E} \int_{t=0}^{\infty} e^{-\beta t} c^\dagger(Z(t)) dt &= \frac{(h + b)(|x| + \bar{\mathbf{y}})}{\lambda + \mu + \beta} \sum_{i \geq 0} \left(\frac{\lambda + \mu}{\lambda + \mu + \beta} \right)^i \\ &\quad + \frac{h + b}{\lambda + \mu + \beta} \sum_{i \geq 1} i \left(\frac{\lambda + \mu}{\lambda + \mu + \beta} \right)^i \\ &= \frac{(h + b)(|x| + \bar{\mathbf{y}})}{\beta} + \frac{(h + b)(\lambda + \mu)}{\beta^2} < \infty, \end{aligned}$$

regardless of m . This completes the proof. \square

REFERENCES

- [1] P. Brémaud, Markov chains: Gibbs fields, Monte Carlo simulation, and queues, Springer-Verlag, New York, 1999.
- [2] J.A. Buzacott and J.G. Shanthikumar, Stochastic models of manufacturing systems, Prentice-Hall, Upper Saddle River, NJ, 1993.
- [3] J.A. Buzacott and J.G. Shanthikumar, Safety stock versus safety time in MRP controlled production systems, *Manage Sci* 40 (1994), 1678–1689.
- [4] F. Chen and J.-S. Song, Optimal policies for multiechelon inventory problems with Markov-modulated demand, *Oper Res* 49 (2001), 226–234.
- [5] F. de Véricourt, F. Karaesmen, and Y. Dallery, Optimal stock allocation for a capacitated supply system, *Manage Sci* 48 (2002), 1486–1501.
- [6] G.A. DeCroix and V.S. Mookerjee, Purchasing demand information in a stochastic-demand inventory system, *Eur J Oper Res* 102 (1997), 36–57.
- [7] I. Duenyas and W.J. Hopp, Quoting customer lead times, *Manage Sci* 41 (1995), 43–57.

- [8] G. Gallego and Ö. Özer, Integrating replenishment decisions with advance order information, *Manage Sci* 47 (2001), 1344–1360.
- [9] G. Gallego and Ö. Özer, “Optimal use of demand information in supply chain management,” *Supply chain structures: Coordination, information and optimization*, J. Song and D. Yao (Editors), Kluwer Academic Publishers, Norwell, MA, 2002, pp. 119–160.
- [10] G. Gallego and Ö. Özer, Optimal replenishment policies for multi-echelon inventory problems under advance order information, *Manuf Serv Oper Manage* 5 (2003), 157–175.
- [11] S. Gavirneni, R. Kapuscinski, and S. Tayur, Value of information in capacitated supply chains, *Manage Sci* 45 (1999), 16–24.
- [12] J.-P. Gayon, S. Benjaafar, and F. de Véricourt, Using imperfect advance demand information in production-inventory systems with multiple customer classes, *Manuf Serv Oper Manage* 11 (2009), 128–143.
- [13] S.C. Graves, H.C. Meal, S. Dasu, and Y. Qiu, “Two-stage production planning in a dynamic environment,” *Multi-stage production planning and control*, S. Axsäter, C. Schneeweiss, and E. Silver (Editors), Springer-Verlag, Berlin, 1986.
- [14] R. Güllü, On the value of information in dynamic production/inventory problems under forecast evolution, *Nav Res Logist* 43 (1996), 289–303.
- [15] X. Guo and O. Hernández-Lerma, Continuous-time controlled Markov chains with discounted rewards, *Acta Applicandae Mathematicae* 79 (2003), 195–216.
- [16] A.Y. Ha, Inventory rationing in a make-to-stock production system with several demand classes and lost sales, *Manage Sci* 43 (1997), 1093–1103.
- [17] R. Hariharan and P. Zipkin, Customer-order information, leadtimes, and inventories, *Manage Sci* 41 (1995), 1599–1607.
- [18] D.C. Heath and P.L. Jackson, Modeling the evolution of demand forecasts with application to safety-stock analysis in production/distribution systems, *IIE Trans* 26 (1994), 17–30.
- [19] W.J. Hopp and M.R. Sturgis, A simple, robust leadtime-quoting policy, *Manuf Serv Oper Manage* 3 (2001), 321–336.
- [20] F. Karaesmen, J.A. Buzacott, and Y. Dallery, Integrating advance order information in make-to-stock production systems, *IIE Trans* 34 (2002), 649–662.
- [21] F. Karaesmen, G. Liberopoulos, and Y. Dallery, The value of advance demand information in production/inventory systems, *Ann Oper Res* 126 (2004), 135–158.
- [22] G. Liberopoulos, A. Chronis, and S. Koukournialos, “Base stock policies with some unreliable advance demand information,” *Proceeding of the 4th Aegean International Conference on Analysis of Manufacturing Systems*, Samos Island, Greece, 2003, pp. 77–86.
- [23] S. Lippman, Applying a new device in the optimization of exponential queueing systems, *Oper Res* 23 (1975), 687–710.
- [24] Ö. Özer, Replenishment strategies for distribution systems under advance demand information, *Manage Sci* 49 (2003), 255–272.
- [25] Ö. Özer and W. Wei, Inventory control with limited capacity and advance demand information, *Oper Res* 52 (2004), 988–1000.
- [26] M.L. Puterman, *Markov decision processes: Discrete stochastic dynamic programming*, Wiley, New York, 1994.
- [27] L.B. Schwarz, N.C. Petrucci, and K. Wee, “The value of advance-order information and the implications for managing the supply chain: An information/control/buffer portfolio perspective,” Working paper, Purdue University, 1997.
- [28] S.P. Sethi, H. Yan, and H. Zhang, Peeling layers of an onion: Inventory model with multiple delivery modes and forecast updates, *J Optim Theory Appl* 108 (2001), 253–281.
- [29] U.W. Thonemann, Improving supply-chain performance by sharing advance demand information, *Eur J Oper Res* 142 (2002), 81–107.
- [30] K. van Donselaar, L.R. Kopczak, and M. Wouters, The use of advance demand information in a project-based supply chain, *Eur J Oper Res* 130 (2001), 519–528.
- [31] M.H. Veatch and L.M. Wein, Monotone control of queueing networks, *Queueing Syst* 12 (1992), 391–408.
- [32] M.H. Veatch and L.M. Wein, Scheduling a make-to-stock queue: Index policies and hedging points, *Oper Res* 44 (1996), 634–647.
- [33] T. Wang and B.L. Toktay, Inventory management with advance demand information and flexible delivery, *Manage Sci* 54 (2008), 716–732.
- [34] K. Zhu and U.W. Thonemann, Modeling the benefits of sharing future demand information, *Oper Res* 52 (2004), 136–147.
- [35] P.H. Zipkin, *Foundations of inventory management*, McGraw-Hill, New York, 2000.