Probability in the Engineering and Informational Sciences, **18**, 2004, 247–274. Printed in the U.S.A.

THE VALUE OF INFORMATION SHARING IN A TWO-STAGE SUPPLY CHAIN WITH PRODUCTION CAPACITY CONSTRAINTS

*THE INFINITE HORIZON CASE**

David Simchi-Levi

The Engineering Systems Division and the
Department of Civil and Environmental Engineering
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139
E-mail: dslevi@mit.edu

Yao Zhao

Department of Management Science and Information Systems
Rutgers University
Newark, New Jersey 07102-1895
E-mail: yaozhao@andromeda.rutgers.edu

We study the value of information sharing in a two-stage supply chain with a single manufacturer and a single retailer in an infinite time horizon, where the manufacturer has finite production capacity and the retailer faces independent demand. The manufacturer receives demand information even during periods of time in which the retailer does not order. Allowing for time-varying cost functions, our objective is to characterize the impact of information sharing on the manufacturer's cost and service level. We develop a new approach to characterize the induced Markov chains under cyclic order-up-to policy and provide a simple proof for the optimality of cyclic order-up-to policy for the manufacturer under the average cost criterion. Using extensive computational analysis, we quantify the impact of information

*Research supported in part by ONR contracts N00014-95-1-0232 and N00014-01-1-0146 and NSF contracts DMI-9732795, DMI-0085683, and DMI-0245352.

sharing on the manufacturer's performance in an infinite time horizon under both i.i.d. demand and independent but nonstationary demand.

1. INTRODUCTION

Information technology is an important enabler of efficient supply chain strategies. Indeed, much of the current interest in supply chain management is motivated by the possibilities introduced by the abundance of data and the savings inherent in sophisticated analysis of these data. For example, information technology has changed the way companies collaborate with suppliers and customers. These collaborations, which allow companies to share sensitive demand information with their suppliers in real time, have achieved huge success in practice in terms of inventory reduction, service level improvement, and quick response to market changes (Stein and Sweat [31]).

The benefits of sharing demand-related information among supply chain partners has been explored by many authors: for example, the Bullwhip effects reduction (Lee, Padmanabhan, and Whang [21] and Chen, Drezner, Ryan, and Simchi-Levi [5]) and the inventory cost reduction (Hariharan and Zipkin [13], Gallego and Ozer [10], Gavirneni, Kapuscinski, and Tayur [11], Aviv and Federgruen [2], Cachon and Fisher [3], Chen [4], and Simchi-Levi and Zhao [28]). These studies cover a broad range of production–distribution systems, such as single-supplier and single-retailer systems (Gavirneni et al. [11], Simchi-Levi and Zhao [28]), single-supplier, multi-retailer systems (Aviv and Federgruen [2], Cachon and Fisher [3]), and multistage series systems (Chen [4] and Chen et al. [5]). An excellent review of recent research can be found in Cachon and Fisher [3].

In this article, we focus on a single capacitated manufacturer serving a single retailer. The retailer faces independent demand and the objective is to analyze the impact of information sharing on the manufacturer's cost and service levels. The impact of information sharing on a capacitated manufacturer has been analyzed by several authors. Aviv and Federgruen [2] analyzed a single-supplier, multi-retailer system in which retailers faced random demand and shared inventories and sales data with the supplier. They analyzed the effectiveness of a Vendor Managed Inventory (VMI) program where sales and inventory data are used by the supplier to determine the timing and the amount of shipments to the retailers. For this purpose, they compared the performance of the VMI program with that of a traditional, decentralized system, as well as a supply chain in which information is shared continuously, but decisions are made individually (i.e., by the different parties). Their focus in the three systems analyzed was on minimizing long-run average cost. Aviv and Federgruen reported that information sharing reduces systemwide cost by 0–5%, whereas VMI reduces cost, relative to information sharing, by 0.4–9.5% and on average by 4.7%. They also showed that information sharing could be very beneficial for the supplier.

The work by Gavirneni et al. [11] analyzed a two-stage supply chain with a single capacitated supplier and a single retailer. In this periodic review model, the retailer made ordering decisions every period, using an (s,S) inventory policy, and transferred demand information every time an order decision was made, independent of whether an order was made. Assuming zero transportation lead time, they showed that the benefit (i.e., the supplier cost savings due to information sharing) increased with production capacity and it ranged from 1% to 35%.

The focus of this article is on the benefits of information sharing for a capacitated supplier in a two-stage supply chain with an infinite time horizon. In this model, the retailer shares demand information with the supplier even at times during which she does not make order decisions. This is clearly the case in many industries in which demand information can be shared very frequently (e.g., every second or every minute) while orders are placed less frequently (e.g., every day or every week). A recent article by Simchi-Levi and Zhao [28] studied this model in a finite time horizon setting. They show that by sharing demand information, the supplier can reduce its inventory cost substantially while maintaining the same service level to the retailer.

The optimal control of a periodic review production–inventory system in an infinite time horizon is a classical Markov decision problem with infinite state space and unbounded cost function. The literature on these problems is quite voluminous (see, e.g., Heyman and Sobel [15] and Puterman [24] for a general theory of the Markov decision process; Karlin [17], Zipkin [32], Sethi and Cheng [27], and Song and Zipkin [30] for systems with time-varying parameters; Federgruen and Zipkin [7,8] for systems with production capacity constraints). Finally, Aviv and Federgruen [1] and Kapuscinski and Tayur [16] have analyzed periodic review production–inventory systems with capacity constraints and time-varying parameters and their results are most relevant to our model. They show that a modified cyclic order-up-to policy (i.e., a modified order-up-to policy with periodically varying order-up-to levels) is optimal under both discounted and average cost criterion. This is done by verifying the optimality conditions developed by Sennott [26] and Federgruen, Schweitzer, and Tijms [6].

In the next section, we describe our model and identify the differences between our model and results and those of Aviv and Federgruen [1] and Kapuscinski and Tayur [16].

2. THE MODEL

Consider a single product, periodic review, two-stage production–inventory system with a single capacitated manufacturer and a single retailer facing independent demand and using an order-up-to inventory policy. In practice, many retailers place orders periodically (e.g., every week or every month). Thus, we assume that the retailer places an order every T time periods (e.g., 7 days) to raise his inventory position to a target level. The manufacturer receives demand information from the retailer every τ units of time $(\tau \leq T)$. For instance, the retailer places an order every week but provides demand information every day. We refer to the time between successive orders as the *ordering period* and the time between successive information sharing as the *information period*. Of course, in most supply chains, information can be shared almost continuously (e.g., every second) while decisions are made less frequently (e.g., every day). Thus, information periods really refers to the time interval between successive uses of the information provided.

The sequence of events in our model is as follows. At the beginning of an *ordering period*, the retailer reviews her inventory and places an order to raise the inventory position to the target inventory level. The manufacturer receives the order from the retailer, fills the order as much as she can from stock, and then makes a production decision. If the manufacturer cannot satisfy all of a retailer's order from stock, then the missing amount is backlogged. The backorder will not be delivered to the retailer until at least the beginning of the next ordering period. Finally, transportation lead time between the manufacturer and the retailer is assumed to be zero. Similarly, at the beginning of an *information period*, the retailer transfers POS (Point of Sales) data of previous information period to the manufacturer. Upon receiving this demand information, the manufacturer reduces this demand from her inventory position although she still holds the stock, and then she makes a production decision.

To simplify the analysis, we assume that demand is stationary across the ordering periods, whereas the demand distribution may vary across different information periods within one ordering period. For example, grocery retailers typically face much higher demand during weekends than week days. At the same time, these retailers are replenished once a week. Since grocery retailers cannot anticipate future shortages from the manufacturer, they are not able to inflate order-up-to levels to protect themselves from future shortages. Hence, we assume in our model that the retailer's target order-up-to level is constant for all the ordering periods; that is, every ordering period, the retailer raises her inventory position (outstanding order plus on-hand inventory minus backorders) to a constant level. Any unsatisfied demand at the retailer is backlogged; thus, the retailer transfers external demand of each ordering period to the manufacturer. The manufacturer has a production capacity limit (i.e., a limit on the amount the manufacturer can produce per unit of time) and the manufacturer runs her production line always at this capacity limit.

Let $N = T/\tau$ be an integer which represents the number of information periods in one ordering period. We index information periods within one ordering period $1, 2, \ldots, N$, where 1 is the first information period in the ordering period and N is the last. Let C denote production capacity in one information period τ, c is the production cost per item, and $0 < \beta < 1$ is the time discounted factor for each information period.

Since we calculate inventory holding cost for each information period, we let h be the inventory holding cost per unit product per information period. Evidently, one unit of product kept in inventory for n information periods, $n = N, N-1, \ldots, 1$, will incur a total inventory cost $h_n = h(1 + \beta + \cdots + \beta^{n-1})$.

Due to the finite production capacity, the manufacturer produces to stock. For every item backlogged at the end of each ordering period, the manufacturer pays π as a penalty cost. It is easy to see that the earlier the manufacturer produces during a single ordering period, the longer she will carry this inventory, thus the higher the inventory holding cost. On the other hand, postponing production too long may result in heavy penalty cost due to the finite production capacity.

We use D_i to denote the end-user demand in the information period i, $i = 1, \dots, N$. D_i is assumed to be independent, its distribution only depends on i, and its mean is defined as ED_i. To simplify the analysis, we assume that costs do not change from information period to information period; later, we will demonstrate that our results can be easily extended to include cases where costs change periodically within one ordering period but are the same across different ordering periods.

We start by considering a finite time horizon model with M ordering periods. Index the ordering periods from 0 to $M-1$, where the 0 ordering period is the first one and the $M-1$ ordering period is the last one. The finite horizon starts in the first ordering period at the beginning of the jth information period, $1 \le j \le N$. Consider the ith information period, $i = 1,2,\dots,N$, in ordering period m, $m = 0,1,\dots,M-1$. Of course, $mN+i \ge j$. We refer to this information period as the $mN+i$ information period. For instance, information period $mN+1$ is the first information period in the mth ordering cycle. We refer to this indexing convention as a forward indexing process.

Define x to be the inventory position at the beginning of the ith information period, and S_i to be the state space for x. Let $y-x$ be the amount produced in that information period, $y \in A_x$, and A_x be the set of feasible actions. Let $\xi(x,y,D)$ be the transition function and D be the demand. In our case, $A_x = [x, x+C]$ and $\xi(x,y,D) = y-D$.

It is easy to verify that $g_{mN+i}(x,y)$, $i = 1,2,\dots,N$, $m = 0,\dots,M-1$, the expected inventory and production cost in information period $mN+i$, given that the period starts with an inventory position x, and $y-x$ items are produced in that period, can be written as

$$g_{mN+i}(x,y) = \begin{cases} c(y-x) + E(L(y,D_N)), & i = N \\ c(y-x) + h_{N-i}(y-x) & \text{otherwise}, \end{cases} \tag{1}$$

where $E(\cdot)$ is the expectation with respect to D_N and

$$L(y,D) = h_N(y-D)^+ + \pi(D-y)^+. \tag{2}$$

Because the items produced in the ith information period are carried over until the end of the ordering period, an inventory holding cost equal to $h_{N-i}(y-x)$ is charged in this information period.

Since $g_{mN+i}(x,y)$ depends only on i and not m, we can write $g_{mN+i}(x,y) = r_i(x,y)$, $\forall m = 0,1,\dots,M-1$ and $\forall i = 0,1,\dots,N$. Furthermore, $r_i(x,y)$ can be written as the sum of a function of x ($\phi_i(x)$) and a function of $y(\varphi_i(y))$, where

$$\phi_i(x) = \begin{cases} -cx, & i = N \\ -(c+h_{N-i})x & \text{otherwise}, \end{cases} \tag{3}$$

$$\varphi_i(y) = \begin{cases} cy + E(L(y,D_N)), & i = N \\ (c+h_{N-i})y & \text{otherwise}. \end{cases} \tag{4}$$

Finally, we assume the salvage cost function $g_{MN+1} = 0$.

We define the optimal policy under the discounted cost criterion as follows. Let $\sigma = \{\sigma_1, \sigma_2, \ldots\}$ be any feasible policy where σ_k is a function depending on the initial inventory position in period k (i.e., $\sigma_k = \sigma_k(x)$ and $\sigma_k(x) \in A_x$ for all k). Let Π be the set of all feasible policies. Define the expected discounted cost from the ith information period in the first ordering period until the end of the horizon when $M \to \infty$, as

$$U_i^{\beta}(x,\sigma) = E\left(\lim_{M\to\infty} \sum_{k=i}^{MN} \beta^{k-i} g_k(x_k, \sigma_k(x_k)) | x_i = x \right), \tag{5}$$

where $E(\cdot)$ denotes the expectation with respect to demand in all information periods and x_i is the initial inventory position of the ith information period for $i = 1, 2, \ldots, N$. A policy $\sigma^* = \{\sigma_1^*, \sigma_2^*, \ldots\} \in \Pi$ is called optimal under the discounted cost criterion if for all $x \in S$ and i,

$$U_i^{\beta}(x,\sigma^*) = \inf_{\sigma\in\Pi} U_i^{\beta}(x,\sigma). \tag{6}$$

Similarly, the optimal policy in an infinite time horizon under the average cost criterion can be defined. Following Heyman and Sobel [14], the performance measure for any feasible policy $\delta = \{\delta_1, \delta_2, \ldots\} \in \Pi$ under the average cost criterion is defined as

$$G_i(x,\delta) = \lim_{M\to\infty} \sup \left(\frac{E\left\{ \sum_{k=i}^{MN} g_k(x_k, \delta_k(x_k)) | x_i = x \right\}}{MN - i + 1} \right). \tag{7}$$

A policy δ^* is optimal if it minimizes $G_i(x,\delta)$ for all x and i over Π.

It is easily seen that in the information sharing model, except for those periods in which orders are placed, the cost function $\varphi_i(y)$ tends to negative infinity as $y \to -\infty$ ($y \to -\infty$ implies that the manufacturer allows for infinite backorders). This is true because although demand is observed by the manufacturer in every information period, it only needs to be satisfied at the end of every ordering period. Thus, the penalty cost is charged only at the end of ordering periods, whereas the inventory holding cost is charged in every information period. Clearly, this issue was not addressed in the models analyzed by Aviv and Federgruen, and Kapuscinski and Tayur since they assumed $\varphi_i(y)$, for all i, were bounded from below.

To further explain the difference between this model and previous models, let us consider the following three cases. First, if the manufacturer only has inventory holding cost but no penalty cost for all periods, then the optimal policy is not finite; that is, the optimal policy does not have a finite order-up-to level because producing nothing in all periods is clearly the optimal policy. Second, if the manufacturer has both inventory holding cost and penalty cost for all periods, then we have the models studied by Aviv and Federgruen [1] and Kapuscinski and Tayur [16]. Finally, if in some periods the manufacturer only has inventory holding cost, whereas in other periods she has both holding and penalty cost, then it is not clear whether there exists an optimal policy with a finite order-up-to level.

Thus, the objective of this article is twofold: First, characterize the optimal policy for the information sharing model under both discounted and average cost criteria; second, quantify the benefits and identify the conditions under which information sharing is most beneficial for the manufacturer in an infinite time horizon; that is, characterize how frequently information should be shared and when it should be shared so that the manufacturer can maximize the potential benefits.

We show that a cyclic order-up-to policy is optimal under both discounted and average cost criteria. Although the optimal policy has the same structure as those of Aviv and Federgruen [1] and Kapuscinski and Tayur [16], our analysis of the average cost criterion, which is based on the vanishing discount method (Heyman and Sobel [14]), is quite different from theirs.

One of the major difficulties in proving the average cost criterion is to characterize the Markov chain induced by the optimal policy under the discounted cost criterion. To address this issue, we introduce a new method to construct and relate Markov chains in a common state space. Together with the Foster's criterion and the Lyapunov function, this method allows us to show that each of the Markov chains induced by any finite cyclic order-up-to policy has a single irreducible positive recurrent class and finite long-run average cost under certain nonrestrictive conditions (Section 3). Then, in Section 4, we provide a simple proof for the cyclic order-up-to policy to be optimal under the average cost criterion. Finally, extensive computational study is conducted in Section 5, using IPA (Infinitesimal Perturbation Analysis), to quantify the impacts of frequency and timing of information sharing, with a particular focus on the differences between finite and infinite time horizons, and nonstationary external demand.

3. PROPERTIES OF CYCLIC ORDER-UP-TO POLICY

In this section, we study the Markov processes associated with any cyclic order-up-to policy and identify conditions under which they are positive recurrent and have finite steady state average cost. The conditions are similar to those identified by Aviv and Federgruen, and Kapuscinski and Tayur, but the analysis is quite different.

At this point, it is appropriate to remind the reader of the definition of a cyclic order-up-to policy in a capacitated production system. In such a system, every time period the manufacturer produces up to the target inventory position if there is enough production capacity, otherwise the manufacturer produces as much as the production capacity allows. Of course, the target inventory position may vary periodically.

Consider the information sharing model with the cost function $r_i(x,y) \sim O(|x|^\rho) + O(|y|^\rho)$, where ρ is a positive integer. Define a cyclic order-up-to policy as a policy with different order-up-to levels for different information periods, but these levels are the same for the same information period in different ordering periods; that is, the order-up-to level in information period $mN + i$ is the same for all m but may be different for different i, $i = 1, 2, \ldots, N$.

Let D_i, $i = 1, \dots, N$, be the random variable representing demand in information period $mN + i$ for all m; demand is assumed to take discrete values. Consider a cyclic order-up-to policy with levels $a_1, a_2, \dots, a_N$, and define the shortfall processes $\{s_{mN+i}, m = 0,1,\dots\}$ for different $i = 1,\dots,N$ as $s_{mN+i} = a_i - y_{mN+i}$ if we are in period $mN + i$ and y_{mN+i} is the inventory position at the end of this period before demand is realized. Clearly, the shortfall may be negative. When it is positive, it represents the difference between what the manufacturer likes to produce and what the production capacity allows the manufacturer to produce. When the target inventory level, a_i, is lower than the initial inventory position at the beginning of the period, the shortfall is negative. The dynamics of the shortfall is

$$s_{mN+i+1} = \begin{cases} (a_{i+1} - a_i) + s_{mN+i} + D_i - C & \text{if } (a_{i+1} - a_i) + s_{mN+i} + D_i > C \\ 0 & \text{if } 0 \le (a_{i+1} - a_i) + s_{mN+i} + D_i \le C \\ (a_{i+1} - a_i) + s_{mN+i} + D_i & \text{if } (a_{i+1} - a_i) + s_{mN+i} + D_i < 0. \end{cases} \tag{8}$$

If excessive stock is returned when the inventory position is higher than the order-up-to level, then the dynamics of shortfall processes $\{s^r_{mN+i}, m = 0,1,\dots\}$ for $i = 1,\dots,N$ is $s^r_{mN+i+1} = (a_{i+1} - a_i + s^r_{mN+i} + D_i - C)^+$. We refer to this policy as an *order-up-to policy with returns*.

Proving that the shortfall process induced by a constant order-up-to policy has a finite steady state average cost is known to be difficult (see, e.g., the appendix of Kapuscinski and Tayur [15]). Proving that each of the shortfall processes induced by a general cyclic order-up-to policy has finite steady state average cost and a single irreducible positive recurrent class is even more difficult because of the more complicated transition matrices.

The following methods have been applied in the literature to address these issues. Aviv and Federgruen showed that if (1) $E(D_i^l) < \infty$ for all positive integers $l \le \rho + 1$ and $i = 1,\dots,N$ and (2) $E(\sum_{i=1}^N D_i) < NC$, then for any finite order-up-to policy, the shortfall process has a finite set of states such that it can be reached with finite expected cost from any starting state. Kapuscinski and Tayur's method has the following two steps: They first characterized the shortfall process under the order-up-to zero policy. Then, they related a cyclic order-up-to policy to an order-up-to zero policy and proved that in steady state, $E(|x_i|^\rho)$ and $E(|s_i|^\rho)$ are finite for a cyclic order-up-to policy under similar conditions as those proved by Aviv and Federgruen, namely (1) $E(D_i^{2\rho+2}) < \infty$ for $i = 1,\dots,N$ and (2) $E(\sum_{i=1}^N D_i) < NC$.

In this section, we introduce a new approach to prove that $\{x_{mN+i}, m = 0,1,\dots\}$, the inventory positions at the beginning of $mN + i$th information period, $\{y_{mN+i}, m = 0,1,\dots\}$, the inventory position at the end of period $mN + i$ but before demand is realized, and $\{s_{mN+i}, m = 0,1,\dots\}$, the shortfall in period $mN + i$, induced by any finite cyclic order-up-to policy give rise to Markov chains with a single irreducible and positive recurrent class and finite steady state average cost under the same condition as Aviv and Federgruen; that is, we require (1) $E(D_i^l) < \infty$ for all positive integers $l \le \rho + 1$ and $i = 1,\ldots,N$ and (2) $E(\sum_{i=1}^N D_i) < NC$.

This is done as follows. First, using the Foster's criterion and the Lyapunov function, we show that the Markov chain induced by an order-up-to constant policy has the following properties: (i) single irreducible and positive recurrent class and (ii) finite steady state average cost. The same properties hold true for the Markov chain induced by a cyclic order-up-to policy with returns. Then, we relate Markov chains induced by a cyclic order-up-to policy and a cyclic order-up-to policy with returns in a common state space. This allows us to prove, under mild conditions, that properties such as positive recurrence and finite steady state average cost hold true for systems using a cyclic order-up-to policy if and only if they hold true for systems using a similar policy with returns.

We begin our analysis by presenting new proofs for the positive recurrence and finite steady state average cost of the order-up-to zero policy. To simplify the analysis, let us assume $\Pr(D_i < 0) = 0, \forall i$.

PROPOSITION 3.1: *Given order-up-to zero policy, the inventory positions* $\{x_{mN+i}, m = 0,1,\ldots\}$ *and* $\{y_{mN+i}, m = 0,1,\ldots\}$ *and the shortfall process* $\{s_{mN+i}, m = 0,1,\ldots\}$ *for* $i = 1,\ldots,N$ *generate discrete-time Markov chains (DTMC) with single irreducible and positive recurrent class if* $\sum_{i=1}^N ED_i < NC$.

PROOF: See the appendix for details. ■

The proof that the steady state average cost associated with an order-up-to zero policy is finite is based on the following lemma.

LEMMA 3.2: *Consider an irreducible and aperiodic Markov chain* $\{X_n, n = 0,1,\ldots\}$ *with a single period cost function* $r(\cdot)$. $r(\cdot)$ *is continuous and bounded from below. Assume that there exists a Lyapunov function* $V(\cdot)$ *mapping the state space* S *to* $[0,\infty)$ *and a constant* η *such that*

$$E\{V(X_{n+1}) - V(X_n) | X_n = x\} \le -r(x) + \eta, \quad \forall x \in S. \tag{9}$$

Then, given an initial state x_0 *with* $V(x_0) < \infty$*, the Markov chain* X_n *has finite steady state average cost if* X_n *is positive recurrent.*

This lemma is a variation of the Foster's second criterion by Meyn and Tweedie [23; see Theorem 14.0.1 (f-regularity)]; thus, we omit the proof and refer to Simchi-Levi and Zhao [29] for technical details. The lemma implies the following.

LEMMA 3.3: *Consider the order-up-to zero policy. If* $E((D_i)^k) < \infty$ *for all integers* $0 \le k \le \rho + 1$, $\forall i$ *and* $\sum_{i=1}^N ED_i < NC$*, then in steady state, the following hold:*

(i) $E(|s_i|^\rho) < \infty$, $E(|x_i|^\rho) < \infty$ *and* $E(|y_i|^\rho) < \infty$ *for all* $i = 1,\ldots,N$.
(ii) For $0 < \beta < 1$, $E(\sum_{n=0}^\infty \beta^n |y_n|^\rho) < \infty$ *and* $E(\sum_{n=0}^\infty \beta^n |x_n|^\rho) < \infty$ *for any initial inventory position* x_0 *and initial information period* i.

PROOF: See the appendix for details. ■

Proposition 3.1 and Lemma 3.3 can be extended to any finite cyclic order-up-to policy with returns as follows.

PROPOSITION 3.4: *Consider any finite cyclic order-up-to policy with returns. If* $\sum_{i=1}^{N} ED_i < NC$ *and* $E((D_i)^k) \leq \infty$ *for any positive integer* $k \leq \rho + 1$ *and* $\forall i$, *then the following hold:*

(i) *Each shortfall process* $\{s^r_{mN+i}, m = 0,1,\ldots\}$, $i = 1,\ldots,N$, *gives rise to a Markov chain with single irreducible and positive recurrent class.*
(ii) $E(|x^r_i|^\rho) < \infty$ *and* $E(|y^r_i|^\rho) < \infty$ *for all i.*
(iii) *For* $0 < \beta < 1$, $E(\sum_{n=0}^{\infty} \beta^n |y^r_n|^\rho) < \infty$ *and* $E(\sum_{n=0}^{\infty} \beta^n |x^r_n|^\rho) < \infty$ *for any initial finite inventory position* x^r_0 *and initial information period i.*

PROOF: Assume $a_i, i = 1,\ldots,N$, to be the order-up-to levels. We only need to transform this policy to an order-up-to zero policy and then apply Proposition 3.1 and Lemma 3.3.

Consider y^r_{mN+i}, the inventory positions at the end of period $mN + i$ before demand is realized. For simplicity, we drop the subscript mN. The system dynamics is $y^r_{i+1} = \min\{y^r_i - D_i + C, a_{i+1}\} = a_{i+1} + (y^r_i - a_{i+1} - D_i + C)^-$, where $x^- = \min\{0, x\}$. Let $z'_i = y^r_i - a_i$ and $D'_i = D_i + (a_{i+1} - a_i)$; then, $z'_{i+1} = (z'_i - D'_i + C)^-$ and $\sum_{i=1}^{N} D'_i = \sum_{i=1}^{N} D_i < NC$. For z'_i and D'_i, this is order-up-to zero policy, where demand D'_i can be negative but bounded from below.

Notice that the shortfall processes associated with this order-up-to zero policy have state space $\{0,1,2,\ldots\}$ even if demand can be negative (due to return). Using proofs similar to those of Proposition 3.1 and Lemma 3.3, we can show that the same results hold for this order-up-to zero policy if there exists a positive constant d_i so that $\Pr\{D_i < -d_i\} = 0$, $\forall i = 1,2,\ldots,N$. ■

We now introduce a method to construct and relate Markov chains induced by a cyclic order-up-to policy and a cyclic order-up-to policy with returns in a common state space. Consider two inventory systems: One uses a cyclic order-up-to policy and the other uses a corresponding policy with returns. The key idea of the method is to characterize the gap between the inventory position processes without returns, y_{mN+i}, and with returns, y^r_{mN+i}, assuming that both inventory systems start with the same initial inventory level and face the same stream of demand.

LEMMA 3.5: *Consider two inventory systems with cyclic order-up-to levels* $a_1,\ldots,a_N$, *for* $N \geq 2$: *one system without returns and the second one with returns. If the two systems start with the same initial state* $x_0 \leq \max\{a_1,\ldots,a_n\}$ *and face the same stream of random demand, then the stochastic processes* $\{z_{mN+i} = y_{mN+i} - y^r_{mN+i}, m = 0,1,\ldots\}$ *have the following properties for all i:*

(i) $z_{mN+i} \geq 0$ *for all m.*
(ii) $z_{mN+i} \leq \max\{a_1,\ldots,a_N\} - \min\{a_1,\ldots,a_N\}$ *for all m.*

PROOF: The proof is by induction. Clearly, $z_1 = 0$. Consider period $mN + i$ and assume $0 \leq z_{mN+i} \leq \max\{a_1,\ldots,a_N\} - \min\{a_1,\ldots,a_N\}$. We distinguish between the following two cases. In the first case, $a_i \leq a_{i+1}$ and in the second case, $a_i > a_{i+1}$.

$a_i \leq a_{i+1}$: There are two subcases to consider. (1) x_{mN+i+1} and x^r_{mN+i+1} are no larger than a_{i+1}. In this case, $y_{mN+i+1} = \min\{x_{mN+i+1} + c, a_{i+1}\}$ and $y^r_{mN+i+1} = \min\{x^r_{mN+i+1} + c, a_{i+1}\}$ and, hence,

$$0 \leq y_{mN+i+1} - y^r_{mN+i+1} \leq x_{mN+i+1} - x^r_{mN+i+1} = y_{mN+i} - y^r_{mN+i}. \tag{10}$$

(2) $x_{mN+i+1} > a_{i+1} \geq x^r_{mN+i+1}$. This can only occur if $N > 2$ and $a_{i+1} < \max\{a_1,\ldots,a_N\}$. Clearly, $y_{mN+i+1} = x_{mN+i+1}$ and $y^r_{mN+i+1} = \min\{x^r_{mN+i+1} + c, a_{i+1}\}$, which implies that

$$0 \leq y_{mN+i+1} - y^r_{mN+i+1} \leq x_{mN+i+1} - x^r_{mN+i+1} = y_{mN+i} - y^r_{mN+i}. \tag{11}$$

$a_i > a_{i+1}$: In this case there are three possible subcases: (1) x_{mN+i+1} and x^r_{mN+i+1} are no larger than a_{i+1}; (2) $x_{mN+i+1} > a_{i+1} \geq x^r_{mN+i+1}$; (3) x_{mN+i+1} and x^r_{mN+i+1} are larger than a_{i+1}. The proof of the first two subcases is identical to the proof in the previous case. Consider subcase 3 and observe that in this case $y_{mN+i+1} = x_{mN+i+1}$ and $y^r_{mN+i+1} = a_{i+1}$ and, hence,

$$0 \leq y_{mN+i+1} - y^r_{mN+i+1} = x_{mN+i+1} - a_{i+1} \leq \max\{a_1,\ldots,a_N\} - \min\{a_1,\ldots,a_N\}. \tag{12}$$

■

Remark: If $x_0 > \max\{a_1,\ldots,a_N\}$, then the state space for $\{z_{mN+i}, m = 0,1,\ldots\}$ is $\{0,1,\ldots,x_0 - \min\{a_1,\ldots,a_N\}\}$, $\forall i$.

THEOREM 3.6: *Consider two arbitrary irreducible DTMCs $\{x_n, n = 0,1,\ldots\}$ and $\{y_n, n = 0,1,\ldots\}$ starting with the same initial state. If their difference process $\{z_n = x_n - y_n, n = 0,1,\ldots\}$ has finite state space S_z, then x_n is positive recurrent if and only if y_n is positive recurrent. Also, x_n has certain finite steady state moments if and only if the same steady state moments of y_n are finite.*

PROOF: First, we show that x_n is positive recurrence if y_n is positive recurrent. Define S_x and S_y to be the state space for $\{x_n, n = 0,1,\ldots\}$ and $\{y_n, n = 0,1,\ldots\}$, respectively. Assume that y_n is positive recurrent; using contradiction, we assume x_n is transient or nonrecurrent for all of its state. Since $y_n = x_n - z_n$ and $y_0 = x_0$, we have

$$\begin{aligned}
\Pr\{y_n = i|y_0\} = \Pr\{y_n = i|x_0\} &= \sum_{k\in S_z} \Pr\{x_n = i + k, z_n = k|x_0\} \\
&= \sum_{k\in S_z} \Pr\{z_n = k|x_n = i + k, x_0\}\Pr\{x_n = i + k|x_0\} \\
&\leq \sum_{k\in S_z} \Pr\{x_n = i + k|x_0\}.
\end{aligned} \tag{13}$$

Since S_z has only a finite number of states and x_n is transient or nonrecurrent, $\Pr\{y_n = i | y_0\} \to 0$ as $n \to \infty$ for $i \in S_y$. This contradicts the assumption that y_n is positive recurrent (see Kulkarni [20, p. 80, Theorem 3.4], and Kemeny, Snell, and Knapp [19, p. 36, Prop. 1-61]).

Second, assume that y_n has finite steady state moments $E(|y|^l)$ for $0 < l \le \rho$, where ρ is a positive integer. Consider

$$\begin{aligned}\sum_{i\in S_x} |i|^\rho \Pr\{x_n = i | x_0\} &= \sum_{i\in S_x} |i|^\rho \sum_{k\in S_z} \Pr\{y_n = i-k, z_n = k | x_0\} \\ &= \sum_{i\in S_x} |i|^\rho \sum_{k\in S_z} \Pr\{z_n = k | y_n = i-k, y_0\}\Pr\{y_n = i-k | y_0\} \\ &\le \sum_{k\in S_z}\sum_{i\in S_x} |i-k+k|^\rho \Pr\{y_n = i-k | y_0\}. \end{aligned} \tag{14}$$

Taking the limit $n \to \infty$ on both sides, we obtain

$$E(|x|^\rho) \le |S_z|(E(|y|^\rho) + c_1 E(|y|^{\rho-1}) + \ldots + c_\rho) < \infty, \tag{15}$$

where $c_1, \ldots, c_\rho$ are positive finite constants and $|S_z|$ is the size of the state space for z_n.

The necessary condition can be easily proved in a similar way. ■

Lemma 3.5 and Theorem 3.6 provide a method to simplify Markov chains with infinite state space and complicated dynamics. Finally, we have the following corollary.

Corollary 3.7: *Consider any finite cyclic order-up-to policy σ. If $\sum_{i=1}^N ED_i < NC$ and $E((D_i)^k) \le \infty$ for all i and positive integers k such that $k \le \rho + 1$, then the following hold:*

(i) Each shortfall process $\{s_{mN+i}, m = 0,1,\ldots\}$, $i = 1,\ldots,N$ gives rise to a Markov chain with single irreducible and positive recurrent class.
(ii) $E(|x_i|^\rho) < \infty$ and $E(|y_i|^\rho) < \infty$ for all i.
(iii) For $0 < \beta < 1$, $E(\sum_{n=0}^\infty \beta^n |x_n|^\rho) < \infty$ and $E(\sum_{n=0}^\infty \beta^n |y_n|^\rho) < \infty$ for any finite initial inventory position x_0 and initial information period i.
(iv) For $0 < \beta < 1$, we have

$$U_i^\beta(x,\sigma) = E\left(\lim_{M\to\infty} \sum_{k=i}^{MN} \beta^{k-i} g_k(x_k, \sigma_k(x_k)) | x_i = x\right) < \infty.$$

(v)

$$\lim_{M\to\infty} \frac{E\left\{\sum_{k=i}^{MN} g_k(x_k, y_k) | x_i = x\right\}}{MN - i + 1}$$

converges to a finite value independent of initial period i and initial state x.

PROOF: Parts (i)–(iv) are the direct results of Theorem 3.6, Lemma 3.5, and Proposition 3.4; thus they do not need a proof.

To prove part (v), notice that $\sum_{i=1}^{N} E(r_i(x, y)) < \infty$ in steady state is due to the assumption that $r_i(x, y) \sim O(|x|^{\rho}) + O(|y|^{\rho})$. Since the shortfall processes and thus y_{mN+i} and x_{mN+i+1} give rise to Markov chains with single irreducible and positive recurrent class, we let the steady state distribution for (x, y) in the ith information period be $p^i_{(x,y)}$, where $(x, y) \in \Omega^i$ (the feasible region of (x, y)), then $E(r_i(x, y)) = \sum_{(x,y)\in\Omega^i} p^i_{(x,y)} r_i(x, y)$ must converge since the summation is over at most a countable number of positive values and it is bounded from above. Finally, applying Proposition 1-61 (arithmetic average) of Kemeny et al. [19], the long-run average cost converges and equals the steady state average cost. ■

4. A MARKOV DECISION PROCESS

Our objective in this section is to discuss the discounted and average cost criterion and present finite optimal policies for the information sharing model. For the discounted cost criterion, we follow Aviv and Federgruen [1] and specify conditions under which the cyclic order-up-to policy is optimal and the optimal order-up-to levels are finite. Other methods can be found in Heyman and Sobel [14], Federgruen and Zipkin [8], and Kapuscinski and Tayur [16].

For the average cost criterion, Sennott [25] and Federgruen et al. [6] introduced conditions under which there exists an optimal policy for general Markov decision processes (MDPs). Thus, one way to characterize the optimal policy is to verify that these conditions hold in our case. In fact, Aviv and Federgruen [1] and Kapuscinski and Tayur [16] have applied this method in their analysis.

We apply a different approach based on the vanishing discount method. For this purpose, we observe that in our model, the optimal policy under the discounted cost criterion can be characterized by a finite set of critical numbers (i.e., order-up-to levels). Thus, to prove the existence of an optimal policy under the average cost criterion, we only need to specify conditions (see Theorem 4.3) for the optimal order-up-to levels under the discounted cost criterion, instead of conditions on the entire policy.

To specify these conditions, we apply a result from Aviv and Federgruen [1]; namely that under certain conditions, the optimal order-up-to levels in the discounted cost criterion are uniformly bounded for all $0 < \epsilon \leq \beta < 1$. We extend this result to the information sharing model (see Proposition 4.2), in which $\varphi_i(y)$ is unbounded from both above and below for some i. This result, together with the results that long-run average cost is finite (Corollary 3.7), implies that a cyclic order-up-to policy is optimal under the average cost criterion (Theorem 4.3). We start by presenting conditions under which the optimal policy in the discounted cost case is a finite cyclic order-up-to policy.

Following convention, we index periods in a reverse order starting at the end of the planning horizon. Let $m = 0$ be the last ordering period and $m = M - 1$ be the first ordering period. We set $i = N$ for the first information period and $i = 1$ for the last information period in any ordering period. Thus, period $mN + i$ represents the ith information period in the $(m-1)$st ordering period. Finally, let $g_{mN+i}(x, y) = r_i(x, y)$ denote the single period expected cost in this information period. We refer to this indexing as backward index and we will only use backward index in this section.

Let $U^\beta_{mN+i}(x)$ be the minimum expected total costs if there are $mN + i$ periods remaining in the planning horizon, starting with an initial state x. Let the salvage cost $U_0^\beta \equiv 0$, and, hence,

$$\begin{aligned} U^\beta_{mN+i}(x) &= \operatorname*{Min}_{y \in A_x} \{g_{mN+i}(x, y) + \beta E(U^\beta_{mN+i-1}(y - D_i))\} \\ &= \operatorname*{Min}_{y \in A_x} \{r_i(x, y) + \beta E(U^\beta_{mN+i-1}(y - D_i))\}, \end{aligned} \tag{16}$$

where $E(\cdot)$ is the expectation with respect to D_i. Observe that $r_i(x, y)$ can be written as the sum of a function of x ($\phi_i(x)$) and a function of y ($\varphi_i(y)$), where

$$\phi_i(x) = \begin{cases} -cx, & i = 1 \\ -(c + h_{i-1})x & \text{otherwise} \end{cases} \tag{17}$$

$$\varphi_i(y) = \begin{cases} cy + E(L(y, D_1)), & i = 1 \\ (c + h_{i-1})y, & \text{otherwise.} \end{cases} \tag{18}$$

Thus, the following recursion must hold:

$$\begin{aligned} U^\beta_{mN+i}(x) &= \phi_i(x) + V^\beta_{mN+i}(x), \\ V^\beta_{mN+i}(x) &= \operatorname*{Min}_{x \le y \le x+C} \{J^\beta_{mN+i}(y)\}, \\ J^\beta_{mN+i}(y) &= \varphi_i(y) + \beta E(U^\beta_{mN+i-1}(y - D_i)). \end{aligned} \tag{19}$$

Observe that in the very first information period (i.e., information period $(M - 1)N + N = MN$) of the entire planning horizon, we have to add $h_N x^+$ to $U^\beta_{MN}(x)$ to account for the holding cost of initial inventory.

The dynamic programming model has the following properties:

1. The cost function $r_i(x, y)$ for each information period $i = 1, \ldots, N$ is positive and convex in y. Thus, the following property, proved in Kapuscinski and Tayur [16], holds: $U^\beta_{mN+i}(x) \ge U^\beta_{(m-1)N+i}(x)$ for any x, $m = 1, \ldots, M$ and $i = 1, \ldots, N$.
2. $r_i(x, y) = \phi_i(x) + \varphi_i(y)$ for all i, and there exists a positive integer ρ so that $\phi_i(x) \sim O|x|^\rho$ and $\varphi_i(y) \sim O|y|^\rho$. From Proposition 3.7, the total expected discounted cost using cyclic order-up-to policy is finite if $\sum_{i=1}^N ED_i < NC$ and $E((D_i)^l) < \infty$ for any positive integer $l \le \rho + 1$, $\forall i$. This implies that $U^\beta_{mN+i}(x)$ converges pointwise to a finite value for any finite x and for all i (Heyman and Sobel [14, Thm. 8-13]). Let $U_i^\beta(x)$ denote the convergence point of $U^\beta_{mN+i}(x)$.

3. $J^{\beta}_{mN+i}(y)$ is convex in y, for all m and i. Furthermore, $\lim_{|y|\to+\infty} J_{mN+i}(y) \to +\infty$ for all m and i if $\beta^{N-1}\pi > c + h_{N-1}$. The main difficulty in proving that a cyclic order-up-to policy with finite target levels is optimal for this model is that the function $\varphi_i(y)$ may be unbounded from both above and below. To overcome this difficulty, we need to aggregate N consecutive information periods and identify conditions under which the cost function for all these periods tends to positive infinity as the action variables approach either positive or negative infinity. Following the analysis of Simchi-Levi and Zhao [31], a sufficient condition is $\beta^{N-1}\pi > c + h_{N-1}$. Intuitively, this condition implies that the discounted penalty cost has to be larger than the sum of the production cost and inventory holding cost for a single ordering period so that the manufacturer should produce even in the first information period given that the initial inventory position is sufficiently low.

THEOREM 4.1: *For the Markov decision process defined in Eq. (19), if*

(a) $\sum_{i=1}^{N} ED_i < NC$ *and* $E((D_i)^l) < +\infty$ *for any positive integer* $l \le \rho + 1$, $\forall i$,
(b) $\beta^{N-1}\pi > c + h_{N-1}$,

then the following hold:

(1) Order-up-to policy is optimal for any m and i.
(2) Optimal order-up-to levels y^*_{mN+i} *are bounded as* $m \to +\infty$.
(3) $J_{mN+i}(y)$ *converges to* $J_i(y)$ *for all* y, i, *and every limit of* y^*_{mN+i} *is a minimal point for* $J_i(x)$.
(4) Cyclic order-up-to policy is optimal under the discounted cost criterion.

PROOF: Since $J^{\beta}_{mN+i}(y)$ is a convex function of y and $\lim_{|y|\to+\infty} J^{\beta}_{mN+i}(y) \to +\infty$ for all m and i, the order-up-to policy is optimal for all m and i. Since $U^{\beta}_{mN+i}(x)$ is bounded from above for any m and finite x, order-up-to levels are finite as $m \to +\infty$ because of (b) (see the proof of Theorem 2 in Aviv and Federgruen [1]). Notice that $U^{\beta}_{mN+i}(x)$ is nondecreasing and converges to $U_i^{\beta}(x)$, which implies $J^{\beta}_{mN+i}(y)$ is nondecreasing and converges to, say, $J_i^{\beta}(y)$, due to the monotone convergence theorem. Hence, part 3 is true (also see the proof of Theorem 2 in Aviv and Federgruen [1]). Finally, parts 1–3 imply that cyclic order-up-to policy is optimal under the discounted cost criterion. ■

We now extend Theorem 3(b) of Aviv and Federgruen [1] to our model.

PROPOSITION 4.2: *For the Markov decision process defined by Eq. (1), if the conditions of Theorem 4.1 are satisfied, then the optimal order-up-to levels* $y_i^*, i = 1,2,\ldots,N$, *under the discounted cost criterion are uniformly bounded both from above and from below for any* $0 < \epsilon \le \beta < 1$.

Since the extension follows a similar proof technique to the one in Aviv and Federgruen, we omit the proof and refer to Simchi-Levi and Zhao [29] for technical details. We are ready to characterize the average cost criterion.

THEOREM 4.3: *Consider the information sharing model and assume the following:*

(a) Cyclic order-up-to policy is optimal under the discounted cost criterion.
(b) For all $0 < \epsilon \le \beta < 1$, the optimal order-up-to levels under the discounted cost criterion are uniformly bounded both from above and from below.
(c) The long-run average cost of any finite cyclic order-up-to policy converges to a finite value.

If these conditions are satisfied, then a cyclic order-up-to policy is optimal under the average cost criterion.

PROOF: Consider a sequence of $\beta_1, \beta_2, \ldots, \beta_j, \ldots \uparrow 1$ as $j \to \infty$. Since the optimal order-up-to levels under the discounted cost criterion are uniformly bounded from both above and below for all $0 < \epsilon \le \beta < 1$, there must exist a finite cyclic order-up-to policy f and a subsequence $j_n \to \infty$ so that f is the optimal policy for all $U_i^{\beta_{j_n}}(x)$, $\forall i = 1, 2, \ldots, N$.

Since we can show that for any finite cyclic order-up-to policy σ,

$$G_i(x,\sigma) = \lim_{M\to\infty} \frac{E\left\{\sum_{k=i}^{MN} g_k(x_k, \sigma_k(x_k)) | x_i = x\right\}}{MN - i + 1}$$

converges to a finite value which is independent of the initial period i and initial state x (see Corollary 3.7); then, by Tauberian theory (Heyman and Sobel [14, p. 172]),

$$\begin{aligned} G_i(x,f) &= \lim_{\beta_{j_n}\uparrow 1} (1-\beta_{j_n}) U_i^{\beta_{j_n}}(x) \\ &\le \lim_{\beta_j \uparrow 1} \sup((1-\beta_j) U_i^{\beta_j}(x)) \\ &\le G_i(x,\delta), \quad \forall \delta \in \Pi; x, i. \end{aligned} \tag{20}$$

The last inequality is justified by the Lemma A2 of Sennott [25]. ■

Finally, we combine Theorem 4.3, 4.1, Corollary 3.7 and Proposition 4.2, to get

COROLLARY 4.4: *In the information sharing model, if*

(a) $\sum_{i=1}^{N} ED_i < NC$ and $E((D_i)^l) < +\infty$ for any positive integer $l \le \rho + 1, \forall i$,
(b) $\beta^{N-1}\pi > c + h_{N-1}$,

then cyclic order-up-to policy is optimal under the average cost criterion.

5. COMPUTATIONAL RESULTS

In this section, we report on an extensive computational study conducted to develop insights about the benefits of information sharing. Our goal is to determine situations in which information sharing provides significant cost savings relative to supply chains with no information sharing in the infinite time horizon. For this purpose, we first examine the impact of the production capacity and the frequency and timing of information sharing on cost savings. This is followed by a systematic comparison between finite and infinite time horizons for i.i.d. external demand. Then, we study the impact of independent but nonstationary demand on the benefits of information sharing. In the computational study, we focus on the average cost criterion and compute the optimal order-up-to levels and cost by employing IPA (Fu [9] and Glasserman and Tayur [12]).

In the model with no information sharing, we assume that the retailer only places orders to the manufacturer at the end of each ordering period. Since the demand is backlogged and the retailer uses an order-up-to policy with constant order-up-to level, the order placed by the retailer is equal to the total demand in one ordering period. Furthermore, we assume that the manufacturer knows the retailer's ordering policy and, therefore, the demand distribution of one ordering period. Finally, we assume that the manufacturer has the same production capacity per information period and charges the same inventory holding cost per item per information period in the no-information-sharing model as in the information sharing model.

The model with no information sharing can be considered as a special case of the information sharing model. Indeed, consider an instance of the model with no information and construct an information sharing model in which demand in every information period within an ordering period is exactly zero except in the last information period. Demand in this information period equals the total demand during that ordering period. This information sharing model has the same dynamic programming formulation as the model with no information sharing. Thus, the dynamic program designed to solve the information sharing model can be applied to solve the model in which information is not shared. Finally, a finite cyclic order-up-to policy is optimal for the model with no information sharing under both the discounted and the average cost criterion.

In all of the numerical studies, we set the production cost $c = 0$ and focus on holding and penalty costs. The initial inventory position, x, at the beginning of the first ordering period is set to be zero without loss of generality.

5.1. i.i.d. Demand

In this subsection, external demand is assumed to be i.i.d. To identify situations in which the manufacturer can achieve significant benefits from information sharing and to compare the cost savings between finite and infinite time horizons, we examine the cases with variation of the following parameters: production capacity, the number of information periods in one ordering period, and the time when information is shared.

5.1.1. The effect of production capacity. To explore the impact of production capacity on the benefit of information sharing in an infinite time horizon, we illustrate in Figure 1 the percentage cost savings from information sharing relative to no information sharing as a function of the production capacity. To compare with the finite time horizon, we also show in Figure 1 the percentage cost savings from

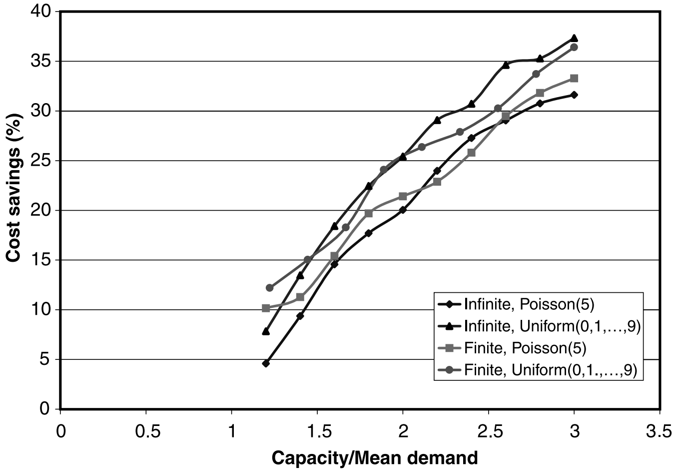

FIGURE 1. The impact of the production capacity, i.i.d. demand.

information sharing when the planning horizon includes only one ordering period (Simchi-Levi and Zhao [28]).

The demand distribution of one information period are Poisson(5) and Uniform(0,1,...,9), and there are four information periods in each ordering period. The inventory holding cost per ordering period is set to be a constant \$4 per unit product for all cases. Thus, the inventory holding cost per information period is \$1 per unit. For each demand distribution and each capacity level, we consider the cases where the ratio of penalty cost to holding costs in one ordering period is \$4.75.

The computational study reveals that, as production capacity increases, the cost saving percentage increases in both finite and infinite time horizons. Indeed, it increases from about 5% to about 35% as capacity over mean demand varies from 1.2 to 3. This is quite intuitive, since as production capacity increases, the optimal policy would postpone production as much as possible and take advantage of all information available prior to the time production starts. Similarly, if the production capacity is limited, then information is not very beneficial because the production quantity is mainly determined by capacity, not realized demand.

We would like to point out that this result is valid only when production capacity is finite. Indeed, if production capacity is large enough (e.g., infinite), then the manufacturer's optimal production policy is obviously to produce to order. Thus, in this case the manufacturer does not benefit from information sharing. Of course, in practice, many manufacturers have limited production capacity and thus they produce to stock. Hence, we limit our computational study to situations in which the ratio of production capacity to mean demand in one ordering period is no more than 3.

Second, we observe that the differences of the percentage cost savings between one ordering period and infinite time horizon is quite small, even if we use different computational methods (i.e., IPA for an infinite time horizon and dynamic programming for the one ordering period case). Finally, similar to the finite horizon case (Simchi-Levi and Zhao [28]), our computational study reveals that information sharing and no information sharing have almost identical fill rates.

5.1.2. The effect of the frequency and timing of information sharing. To understand the impact of the frequency of information sharing in the infinite time horizon case, in Figure 2 we display the percentage cost savings from information sharing as a function of the number of information periods in one ordering period, for both finite and infinite time horizon cases. The number of information periods, N, was two, four, six, and eight, whereas the length of the ordering period was assumed to be constant in all cases. The demand distribution during the entire ordering period is assumed to be Poisson with parameter $\lambda = 24$; hence, demand in a single information period is Poisson with parameter λ/N. Similarly, the inventory holding cost per ordering period is set up to be a constant \$4 per unit product. Thus, the inventory holding cost per information period is $4/N$, where N is the number of information periods within one ordering period. Total production capacity in the entire ordering period is kept constant and is equally divided among the different information periods. Finally, the ratio of penalty to holding costs is set to be 4.75 in all cases.

Figure 2 implies that as the number of information periods increases, the percentage savings increase. However, the marginal benefit is a decreasing function of the number of information periods. Specifically, the additional benefit achieved by increasing the number of information periods from four to eight is relatively small. Finally, the difference between the cost savings obtained in finite and infinite time horizons is relatively small.

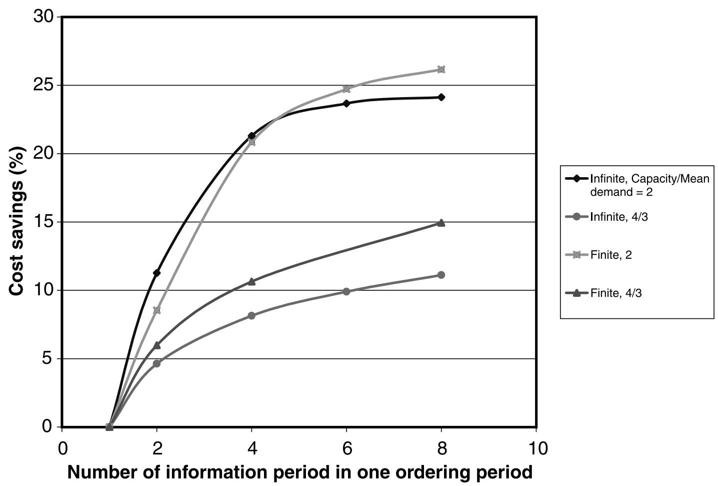

FIGURE 2. The impact of the frequency of information sharing.

To understand the impact of the time when information is shared and the optimal timing of information sharing in infinite time horizon, we analyze the following case in which the retailer only shares demand information once with the manufacturer in one ordering period. In this case, we equally divide 1 ordering period into 10 intervals and compute the total cost for the manufacturer when the retailer shares demand information with her at one of these intervals.

Figure 3 presents the total cost of the manufacturer as a function of time when information is shared. Figure 3 provides normalized cost as a function of normalized time; that is, time is normalized and is measured from zero to one, while cost is normalized by $h_N NED$, where ED is the expected demand in one information period. Thus, 0 in the x coordinate implies that information is shared at the beginning of an ordering period, and 1 means that information is shared at the end of an ordering period and hence cannot be used. Demand distribution is assumed to be Poisson(24) and the ratio of penalty to holding costs is 4.

Figure 3 implies that the optimal timing for information sharing is in the second half of the order interval. This is true both in the finite and infinite horizon models. Intuitively, when capacity is very large, it is appropriate to postpone the timing of information sharing to the last production opportunity in this ordering period; interestingly, this is also the right thing to do when capacity is tightly constrained (i.e., postpone the timing of information sharing until the last production opportunity). One possible explanation is that when capacity is very tight, the manufacturer needs to build as much inventory as she can until the last production opportunity, when she can review demand information and adjust production quantity.

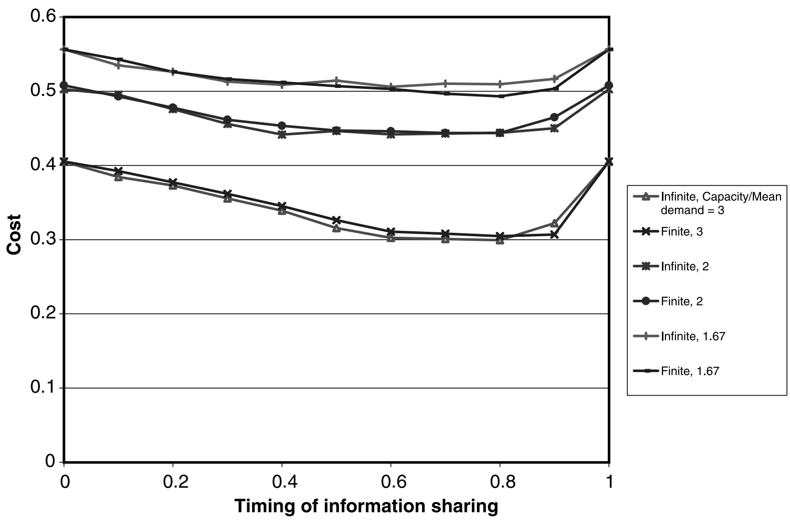

FIGURE 3. The impact of the timing of information sharing for different levels of production capacity, i.i.d. demand.

5.2. Nonstationary Demand

In this subsection, we study the impact of independent but nonstationary demand on the benefits of information sharing in an infinite time horizon case. For this purpose, we first compare the percentage cost savings obtained in two systems: one facing i.i.d. demand and the other facing independent, nonstationary demand distributions. Then, we examine how nonstationary demand affects the optimal timing of information sharing.

To generate the nonstationary demand process, we modeled demand as a nonhomogenous Poisson process and considered the following two scenarios: piecewise increasing rate and piecewise decreasing demand rate.

In particular, demand follows Poisson$(\lambda(t))$ distribution for $t \in [0, T]$. The demand rate $\lambda(t)$ is a piecewise constant simple function, and $\int_0^T \lambda(t)\,dt = \Lambda$, where Λ is the total average demand in one ordering period. Assuming that there are N information periods in one ordering period and $\lambda(t)$ is a constant in every information period, we can write it as $\lambda(n), n = 1, \ldots, N$.

Define $d\lambda$ to be the difference between the highest demand rate and the lowest demand rate in one ordering period and let $\Delta\lambda = [N/(N-1)]\,d\lambda$. Then, we determine the increasing $\lambda(n)$ as follows:

$$\begin{aligned}
\lambda(1) &= \frac{\Lambda}{N} - \frac{\Delta\lambda}{2} + \frac{\Delta\lambda}{2N}, \\
\lambda(2) &= \lambda(1) + \frac{\Delta\lambda}{N}, \\
&\vdots \\
\lambda(n+1) &= \lambda(n) + \frac{\Delta\lambda}{N}, \qquad n = 2, \ldots, N-1. \qquad \textbf{(21)}
\end{aligned}$$

The decreasing $\lambda(n)$ can be determined by reversing the index of information periods. It is easily seen that $\lambda(n)$ satisfies the condition $\int_0^T \lambda(t)\,dt = \Lambda$. To create a comparable i.i.d. demand distribution, we let the demand distribution be Poisson(Λ/N) in every information period.

To study the impact of nonstationary demand on cost savings, we show in Figure 4 the percentage cost saving from information sharing as a function of production capacity with increasing, decreasing, and i.i.d. demand. In all cases, $N = 4$ and each ordering period is equally divided into N information periods. We choose $\Lambda = 20$, $\Delta\lambda = 6$, and the ratio of penalty to inventory holding cost in one ordering period to be 4.75.

Figure 4 illustrates that nonstationary demand has a significant impact on the benefit from information sharing. The percentage cost saving is the smallest when the demand rate is increasing, and it is the highest when the demand rate is decreasing for all production capacity levels. For instance, the difference between the percentage cost saving of increasing and decreasing demand rates is about 15% when the ratio of capacity to mean demand is 2.

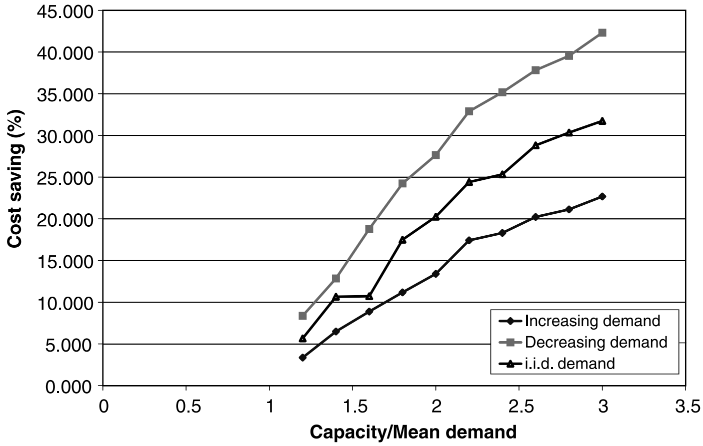

FIGURE 4. The impact of production capacity, non-i.i.d. demand.

This is quite intuitive because in the case of decreasing demand rate, realized demand in the first few information periods, on average, accounts for a larger portion of the total demand in one ordering period relative to increasing demand rate. Indeed, an extreme example of increasing demand rate is that demand is exactly zero in all information periods except the last one. In the last information period, demand equals the total demand in one ordering period. As we show in Section 5, this example is equivalent to a model with a no-information-sharing case. Similarly, an extreme example of the decreasing demand rate is that demand is exactly zero in all information periods except for the first one. In the first information period, demand equals the total demand in one ordering period and, hence, information has a significant impact on the manufacturer's cost.

To study the impact of nonstationary demand on the optimal timing of information sharing, we present in Figure 5 the manufacturer's cost as a function of the time when information is shared. The ratio of production capacity to mean demand in one ordering period is 2. We use the same settings as in Section 5.1 except that we allow for nonstationary demand distributions, where $\Lambda = 30$ and $\Delta\lambda = 4$.

In Figure 6, we demonstrate the optimal timing of information sharing as a function of the production capacity in the cases of increasing and decreasing demand rates.

Figures 5 and 6 suggest the following observations:

- Information sharing is most beneficial in the case of decreasing demand rate and it is least beneficial in the case of increasing demand rate. Given that demand information is shared only once in one ordering period, the percentage cost saving in the case of decreasing demand rate can be as much as 17%, whereas it can only be as much as 11.3% in the case of increasing demand rate.

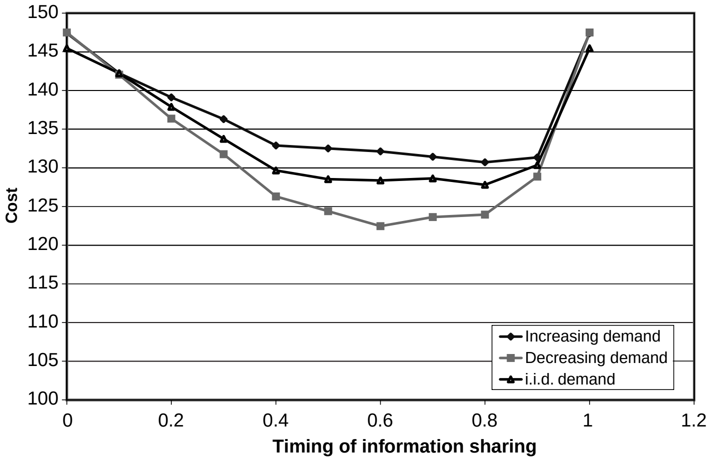

FIGURE 5. The impact of the timing of information sharing, non-i.i.d. demand.

- Figures 5 and 6 show that demand information should always be shared earlier in the case of decreasing demand rate relative to increasing demand rate for different ratios of production capacity to mean demand. For example, when the ratio of capacity to mean demand equals 2 (Fig. 5), it is optimal to transfer demand information as early as 0.6 (in a scaled time horizon) in the

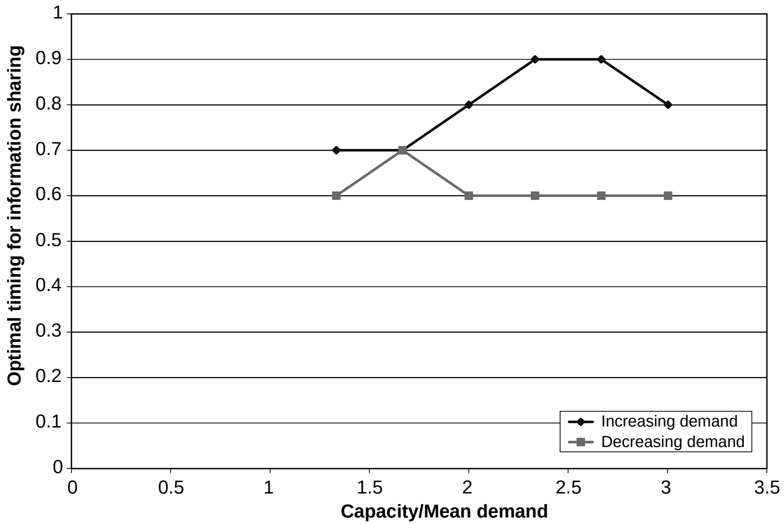

FIGURE 6. The impact of non-i.i.d. demand on the optimal timing of information sharing.

case of decreasing demand rate, whereas in the case of increasing demand rate, the optimal timing of information sharing is as late as 0.8.

- This is quite intuitive because when demand rate is decreasing, most of the demand faced by the retailer is realized close to the beginning of the ordering period. On the other hand, when the demand rate is increasing, most demands are realized close to the end of the ordering period. Thus, it is better to share information earlier and leave more time for production in the decreasing demand rate case.

6. CONCLUSION

In this article, we analyze the value of information sharing in a two-stage supply chain with a single manufacturer and a single retailer. The manufacturer has finite production capacity and she receives demand information from the retailer even during periods of time in which the retailer does not make ordering decisions. A similar model is studied by Simchi-Levi and Zhao [28] in the finite time horizon case; the current article extends the analysis to the infinite time horizon case.

For this purpose, we first show that for any finite cyclic order-up-to policy, the associated inventory positions and shortfalls give rise to Markov chains with a single irreducible, positive recurrent class and a finite steady state average cost. The proof is based on the Foster's criterion and Lyapunov function, as well as a new method to relate Markov chains. This, together with the observation that the optimal policy under the discounted cost criterion can be characterized by a finite set of critical numbers, enables us to provide a simple proof for the optimality of cyclic order-up-to policy under the average cost criterion.

Interestingly, this approach can be applied to prove the existence of an optimal policy under the average cost criterion for other MDPs with unbounded costs. For this purpose, the optimal policy in these MDPs under the discounted cost criterion must be characterized by a finite set of critical numbers—for instance, inventory models with setup cost where the optimal policy under the discounted cost criterion is an (s, S) policy.

Using an extensive computational study, we demonstrate the potential benefits of information sharing on the manufacturer's cost and service level. In particular, we observe that the percentage cost savings due to information sharing increases as production capacity increases. This is true in both finite and infinite time horizon models. We also observe that nonstationary demand may have a substantial impact on both the benefits from information sharing and the optimal timing of information sharing. For instance, if the demand rate is increasing, the benefit from information sharing is not as high as that of decreasing demand rate.

References

1. Aviv, Y. & Federgruen, A. (1997). Stochastic inventory models with limited production capacities and periodically varying parameters. *Probability in the Engineering and Informational Sciences* 11: 107–135.

2. Aviv, Y. & Federgruen, A. (1998). The operational benefits of information sharing and vendor managed inventory (VMI) programs. Working Paper. Washington University, St. Louis.
3. Cachon, G.P. & Fisher, M. (2000). Supply chain inventory management and the value of shared information. *Management Science* 46: 1032–1050.
4. Chen, F. (1998). Echelon reorder points, installation reorder points, and the value of centralized demand information. *Management Science* 44: 221–234.
5. Chen, F., Drezner, I., Ryan, J.K., & Simchi-Levi, D. (2000). Quantifying the bullwhip effect in a simple supply chain: The impact of forecasting, lead time, and information. *Management Science* 46: 436–443.
6. Federgruen, A., Schweitzer, P., & Tijms, H. (1983). Denumerable undiscounted decision processes with unbounded rewards. *Mathematics of Operations Research* 8: 298–314.
7. Federgruen, A. & Zipkin, P. (1986). An inventory model with limited production capacity and uncertain demands 1: The average-cost criterion. *Mathematics of Operations Research* 11: 193–207.
8. Federgruen, A. & Zipkin, P. (1986). An inventory model with limited production capacity and uncertain demands 2: The discounted-cost criterion. *Mathematics of Operations Research* 11: 208–215.
9. Fu, M.C. (1994). Optimization via simulation: A review. *Annals of Operations Research* 53: 199–247.
10. Gallego, G. & Ozer, O. (2001). Integrating replenishment decisions with advanced demand information. *Management Science* 47: 1344–1360.
11. Gavirneni, S., Kapuscinski, R., & Tayur, S. (1999). Value of information in capacitated supply chains. *Management Science* 45: 16–24.
12. Glasserman, P. & Tayur, S. (1995). Sensitivity analysis for base-stock levels in multiechelon production-inventory systems. *Management Science* 41: 263–281.
13. Hariharan, R. & Zipkin, P. (1995). Customer-order information, lead times, and inventories. *Management Science* 41: 1599–1607.
14. Heyman, D. & Sobel, M. (1984). *Stochastic models in operations research*, Vol. 2. New York: McGraw-Hill.
15. Kapuscinski, R. & Tayur, S. (1995). A capacitated production–inventory model with periodic demand. GSIA Working Paper. Carnegie Mellon University, Pittsburgh, PA.
16. Kapuscinski, R. & Tayur, S. (1998). A capacitated production-inventory model with periodic demand. *Operations Research* 46: 899–911.
17. Karlin, S. (1960). Dynamic inventory policy with varying stochastic demands. *Management Science* 6: 231–258.
18. Karr, A.F. (1993). *Probability*. New York: Springer-Verlag.
19. Kemeny, J.G., Snell, J.L., & Knapp, A.W. (1966). *Denumerable Markov chains*. New York: D. Van Nostrand.
20. Kulkarni, V.G. (1995). *Modeling and analysis of stochastic systems*. New York: Chapman & Hall.
21. Lee, H., Padmanabhan, P., & Whang, S. (1997). Information distortion in a supply chain: The Bullwhip effect. *Management Science* 43: 546–558.
22. Lee, H., So, K.C., & Tang, C.S. (2000). The value of information sharing in a two-level supply chain. *Management Science* 46: 626–643.
23. Meyn, S.P. & Tweedie, R.L. (1993). *Markov chains and stochastic stability*. London: Springer-Verlag.
24. Puterman, M.L. (1994). *Markov decision processes: Discrete stochastic dynamic programming*. New York: Wiley.
25. Sennott, L.I. (1989). Average cost optimal stationary policies in infinite state Markov decision processes with unbounded costs. *Operations Research* 37: 261–272.
26. Sennott, L.I. (1999). *Stochastic dynamic programming and the control of queuing systems*. New York: Wiley.
27. Sethi, S.P. & Cheng, F. (1997). Optimality of (s,S) policies in inventory models with Markovian demand. *Operations Research* 45: 931–939.
28. Simchi-Levi, D. & Zhao, Y. (to appear). The value of information sharing in a two-stage supply chain with production capacity constraints. *Naval Research Logistics*; available upon request from yaozhao@andromeda.rutgers.edu.

29. Simchi-Levi, D. & Zhao, Y. (2001). The value of information sharing in a two-stage supply chain with production capacity constraints: The infinite horizon case. Working Paper. Rutgers University, Newark, NJ; available upon request from yaozhao@andromeda.rutgers.edu.
30. Song, J.S. & Zipkin, P. (1993). Inventory control in a fluctuating environment. *Operations Research* 41: 351–370.
31. Stein, T. & Sweat, J. (1998). Killer supply chains. *Information Week*. Obtained at www.information week.com/708/08iukil.htm.
32. Zipkin, P. (1989). Critical number policies for inventory models with periodic data. *Management Science* 35: 71–80.

APPENDIX

Proof of Proposition 3.1

Let us first consider the shortfall process $\{s_{mN+n}, m = 0,1,\ldots\}$ for $n = 1,\ldots,N$. We start our proof by assuming i.i.d. demand D with mean ED. Without loss of generality, we assume initial state $x_0 \leq 0$, since states larger than zero are transient. Since

$$\Pr\{D = i\}\begin{cases} \geq 0, & \forall i \geq 0 \\ = 0 & \text{otherwise,} \end{cases}$$

the shortfall process $\{s_t, n = 0,1,\ldots\}$ has state space $S = \{0,1,\ldots\}$ and transition function $s_t = (s_{t-1} + D - C)^+$. Thus, the transition matrix is

$$P = \begin{bmatrix} \Pr\{D \leq C\} & \Pr\{D = C+1\} & \Pr\{D = C+2\} & \ldots \\ \Pr\{D \leq C-1\} & \Pr\{D = C\} & \Pr\{D = C+1\} & \ldots \\ \Pr\{D \leq C-2\} & \Pr\{D = C-1\} & \Pr\{D = C\} & \ldots \\ \vdots & & & \\ \Pr\{D = 0\} & \Pr\{D = 1\} & \Pr\{D = 2\} & \ldots \\ 0 & \Pr\{D = 0\} & \Pr\{D = 1\} & \ldots \\ \vdots & & & \end{bmatrix}. \tag{A.1}$$

Let $d(i) = E(s_{t+1} - s_t | s_t = i) = \sum_{j\in S}(j-i)P_{ij}$. If $i \geq C$,

$$\begin{aligned} d(i) &= \sum_{j=0}^{\infty}[(j+i-C) - i]\Pr\{D = j\} \\ &= \sum_{j=0}^{\infty}(j-C)\Pr\{D = j\} \\ &= ED - C. \end{aligned} \tag{A.2}$$

If $i < C$,

$$\begin{aligned} d(i) &= -i\,\Pr\{D \le C-i\} + \sum_{j=1}^{\infty}(j-i)\Pr\{D = C+j-i\} \\ &= (C-i)\Pr\{D < C-i\} + \sum_{j=0}^{\infty}(C+j-i)\Pr\{D = C+j-i\} - C \\ &\le (C-i)\Pr\{D < C-i\} + ED - C \\ &< \infty. \end{aligned} \tag{A.3}$$

So from Pakes' lemma (Kulkarni [20]), the shortfall process of order-up-to zero policy with i.i.d demand is positive recurrent if $ED < C$.

Assume that demands in different periods are independent of each other; the result can be easily extended to systems with periodical demand $D_1, D_2, \ldots, D_N$ for any finite N (see Simchi-Levi and Zhao [29] for technical details).

The relationship between y_t and s_t is $s_t = -y_t$, and we also have $x_{t+1} = \min\{0, x_t + C\} - D_t$. The proofs of positive recurrence for x_t and y_t are similar.

Proof of Lemma 3.3

To prove the *first part* of the lemma, it is sufficient to show the existence of a Lyapunov function V satisfying the requirement of Lemma 3.2.

We start by analyzing *i.i.d. demand* D with mean ED. Let $V(x) = q_\rho(x+C)^{\rho+1}$ and $q_\rho = 1/(\rho+1)(C-ED)$. Clearly, $V(\cdot)$ maps the state space of the shortfall $S = \{0,1,2,\ldots\}$ to R^+.

If $i \ge C$,

$$\begin{aligned} E(V(s_{n+1}) - V(s_n)|s_n = i) &= q_\rho \sum_{j=0}^{\infty}[(j+i)^{\rho+1} - (i+C)^{\rho+1}]\Pr\{D = j\} \\ &= q_\rho \sum_{k=0}^{\rho+1}\binom{\rho+1}{k}(m_k - C^k)i^{\rho+1-k}, \end{aligned} \tag{A.4}$$

where $m_k = \sum_{j=0}^{\infty} j^k \Pr\{D = j\}$ is the kth moment of demand. Further expanding the equation, we obtain

$$\begin{aligned} &E(V(s_{n+1}) - V(s_n)|s_n = i) \\ &\quad = -i^\rho + q_\rho\left[\binom{\rho+1}{2}(m_2 - C^2)i^{\rho-1} + \cdots + (m_{\rho+1} - C^{\rho+1})\right]. \end{aligned} \tag{A.5}$$

If $i < C$,

$$\begin{aligned} &E(V(s_{n+1}) - V(s_n)|s_n = i) \\ &\quad = q_\rho\left[(C^{\rho+1} - (i+C)^{\rho+1})\Pr\{D < C-i\} + \sum_{j=C-i}^{\infty}[(j+i)^{\rho+1} - (i+C)^{\rho+1}]\Pr\{D = j\}\right] \\ &\quad = q_\rho\left[C^{\rho+1}\Pr\{D < C-i\} + \sum_{j=C-i}^{\infty}(j+i)^{\rho+1}\Pr\{D = j\} - (i+C)^{\rho+1}\right] \\ &\quad = q_\rho\left[\sum_{j=0}^{C-i-1}[C^{\rho+1} - (j+i)^{\rho+1}]\Pr\{D = j\} + \sum_{j=0}^{\infty}[(j+i)^{\rho+1} - (i+C)^{\rho+1}]\Pr\{D = j\}\right]. \end{aligned} \tag{A.6}$$

To summarize, in both cases,

$$E(V(s_{n+1}) - V(s_n)|s_n = i)$$
$$= q_\rho g_\rho(C,i) - i^\rho + q_\rho\left[\binom{\rho+1}{2}(m_2 - C^2)i^{\rho-1} + \cdots + (m_{\rho+1} - C^{\rho+1})\right], \quad \textbf{(A.7)}$$

where

$$g_\rho(C,i) = \begin{cases} 0, & i \geq C \\ \sum_{j=0}^{C-i-1} [C^{\rho+1} - (j+i)^{\rho+1}] \Pr\{D=j\}, & 0 \leq i < C. \end{cases} \quad \textbf{(A.8)}$$

Define the single-period cost function as

$$r_\rho(x) = x^\rho - q_\rho\left[\binom{\rho+1}{2}(m_2 - C^2)x^{\rho-1} + \cdots + \binom{\rho+1}{\rho}(m_\rho - C^\rho)x\right], \quad \textbf{(A.9)}$$

and since $m_k < +\infty$ for all positive integer $k \leq \rho + 1$, then from Lemma 3.2 and Proposition 3.1, steady-state average cost is finite for the shortfall s_n with single-period cost function $r_\rho(x)$.

In fact, if the single period cost function is $r_l(x), \forall 0 < l < \rho$, where

$$r_l(x) = x^l - q_l\left[\binom{l+1}{2}(m_2 - C^2)x^{l-1} + \cdots + \binom{l+1}{l}(m_l - C^l)x\right], \quad \textbf{(A.10)}$$

the same analysis shows that the corresponding steady state average cost is finite for the shortfall s_n.

Finally, our objective is to show that if the single period cost function is x^l, then the steady state average cost of the shortfall is finite. For this purpose, we use induction on l. The case $l = 1$ is obvious since we already know that for $r_1(x) = x$, the steady-state average cost is finite. By induction on l and the fact that steady-state average cost of the shortfall is finite for $r_l(x)$, $0 < l < \rho$, we obtain our result.

We can extend the result to independent demand with periodically *varying* distributions $D_1, D_2, \ldots, D_N$ in a similar way by defining $V(x) = Q_\rho(x + NC)^{\rho+1}$ and $Q_\rho = 1/(\rho+1)(NC - \sum_{n=1}^N ED_n)$. We omit the proof and refer the readers to Simchi-Levi and Zhao [29] for technical details.

We now prove the *second part* of the lemma. Since s_t is nonnegative, the monotone convergence theorem implies that

$$E\left(\sum_{t=1}^{\infty} \beta^t |s_t|^n\right) = \sum_{t=1}^{\infty} \beta^t E|s_t|^n. \quad \textbf{(A.11)}$$

Because of the first part of this lemma and $0 < \beta < 1$, this summation is for a power series with positive and bounded coefficients, so it is finite.

Since the inventory position processes $y_t = -s_t$ and $x_{t+1} = y_t - D_t$, it is easy to show that the same arguments hold for y_t and x_t. ■