Approximation Algorithms for Perishable Inventory Systems with Setup Costs

Huanan Zhang, Cong Shi, Xiuli Chao
Industrial and Operations Engineering, University of Michigan, MI 48105, {zhanghn, shicong, xchao}@umich.edu

We develop the first approximation algorithm for periodic-review perishable inventory systems with setup costs. The ordering lead time is zero. The model allows for correlated demand processes which generalize the well-known approaches to model dynamic demand forecast updates. The structure of optimal policies for this fundamental class of problems is not known in the literature. Thus, finding provably near-optimal control policies has been an open challenge. We develop a randomized proportional-balancing policy (RPB) that can be efficiently implemented in an online manner, and show that it admits a worst-case performance guarantee between 3 and 4. The main challenge in our analysis is to compare the setup costs between RPB and the optimal policy in the presence of inventory perishability, which departs significantly from the previous works of Chao et al. (2015b,a). The numerical results show that the average performance of RPB is good (within 1% of optimality under i.i.d. demands and within 7% under correlated demands).

Key words: inventory, perishable products, setup costs, randomized algorithms, worst-case analysis
Received October 2014; revisions received March 2015, October 2015, December 2015; accepted January 2016.

1. Introduction

We study the periodic-review perishable inventory systems with setup costs under a general class of associated demand processes (see §2). Undoubtedly, perishable products are an indispensable part of our lives. For example, perishable products such as meat, fruit, vegetable, dairy products, and frozen foods constitute the majority of supermarket sales. Moreover, virtually all pharmaceuticals belong to the category of perishable products. Blood bank provides another salient example where whole blood has finite lifetimes (see, e.g., Prastacos (1984)). In general, the analysis of perishable products is much harder than that of non-perishables.

The work on classical perishable inventory systems dates back to Nahmias (1975) and Fries (1975) who characterized the structure of the optimal ordering policy with i.i.d. demands and no setup costs for backlogging and lost-sales models, respectively. We refer interested readers to Karaesmen et al. (2011) and Chao et al. (2015b,a) for a review of the field. Unfortunately, almost all the papers in the existing perishable inventory literature do not model positive setup costs (mainly due to its mathematical intractability), despite the fact that fixed cost is often unavoidable in many practical settings. As a result, little is known about the structure of the optimal policies for this class of problems with setup costs. Thus, finding provably near-optimal control policies has been an open challenge. In this paper, we propose an efficient randomized proportional-balancing policy that admits a worst-case performance guarantee between 3 and 4. We also demonstrate via numerical experiments that the proposed policy performs consistently well.

Relevant literature. Nahmias (1978) was the first to analyze perishable inventory problems with positive setup costs. He showed that the cost function is in general not K-convex and that (s, S) policy may not be optimal. Indeed, the optimal control policy for this class of problems is very complicated, and there is no known structural characterization for it. Nevertheless, Nahmias (1978) demonstrated, computationally, that (s, S) type policies perform well under i.i.d. demands. Lian and Liu (1999) also considered a periodic-review model with positive setup costs, and constructed a multi-dimensional Markov chain to model and analyze the inventory-level process. Lian et al. (2005) used the same costs and replenishment assumptions as Lian and Liu (1999), and constructed a multi-dimensional Markov chain to model the inventory-level process to derive cost expressions. They numerically showed that the variability in the lifetime distribution can have a significant effect on system performance. However, the aforementioned heuristic policies can only be applied to i.i.d. demand processes, and furthermore they lack any performance guarantees.

Main results and contribution. The contribution of this note is to present the first approximation algorithm for periodic-review perishable inventory system with setup costs that admits a worst-case performance guarantee between 3 and 4. The result holds not only for independent demand processes, but also for a general class of correlated processes, which we call *associated* demand processes. The proposed policy will be referred to as a *randomized proportional-balancing policy* (RPB). As seen from our literature review, this class of problems is fundamental in perishable inventory literature that has challenged researchers for decades; yet little is known about both the structure of optimal policies and the design of provably-good heuristic policies.

Our approach builds on Chao et al. (2015b,a) on perishable inventory systems without setup costs. In particular, we adopt their marginal cost accounting schemes of computing marginal holding, backlogging and outdating costs. The main idea underlying this marginal cost accounting approach is to decompose the total cost in terms of the marginal costs of individual decisions (rather than the conventional per-period cost accounting). However, the nonlinear setup costs make the balancing of the various costs much more complicated. Thus, we construct a randomized algorithm to tackle this problem, striking a right balance between different cost components. The worst-case performance analysis is inevitably much more sophisticated, due to inventory perishability. The key technical challenge is to amortize the setup costs of RPB against the optimal policy. We note that Levi and Shi (2013) on stochastic lot-sizing problems had the same challenge but their partition of periods is rather straightforward without perishability, while our problem needs to consider the inventory age information that is essential in the analysis of our system. (To this end, we also provide a detailed discussion in §4 on why the previous partition in Levi and Shi (2013) fails to work in our perishable inventory systems.) This additional information on inventory age enables us to amortize the setup cost incurred by RPB against that of the optimal policy. On the technical level, we provide a novel sample-path argument that shows that the optimal policy has to place an order between two consecutive *problematic* periods in which the relationship between the trimmed ending inventory levels of RPB and the optimal policy is unclear (due to randomized decisions). This is in sharp contrast with the averaging argument used in Levi and Shi (2013). Moreover, Levi and Shi (2013) need not consider the age information in their analysis. Our construction provides the right and delicate framework to use the age information to analyze the system, which could be useful in analyzing other more complex perishable inventory systems (e.g., with positive lead times and/or finite ordering capacities).

We also demonstrate through an extensive computational study that our proposed algorithm performs quite well empirically (with average error under 7% and maximum error of 11.28% under correlated demands). Our algorithm has also comparable (if not, better) performance with that of Nahmias (1978) under i.i.d. demands, in which both algorithms perform very close to optimal.

Structure and general notation. The rest of this paper is organized as follows. In §2, we formally describe the periodic-review perishable inventory systems with setup costs. In §3, we introduce the randomized proportional-balancing policy (RPB). In §4, we carry out a worst-case performance analysis of RPB. In §5, we demonstrate the empirical performance of RPB.

Throughout this paper, we often distinguish between a random variable and its realizations using capital and lower-case letters, respectively. For any real numbers x and y, we denote $x^+ = \max\{x, 0\}$, $x \vee y = \max\{x, y\}$, and $x \wedge y = \min\{x, y\}$. In addition, for a sequence $x_1, x_2, \ldots$ and any integers t and s with $t \leq s$, we denote $x_{[t,s]} = \sum_{j=t}^{s} x_j$ and $x_{[t,s)} = \sum_{j=t}^{s-1} x_j$. The indicator function $\mathbb{1}(A)$ takes value 1 if A is true and 0 otherwise, and "$\triangleq$" stands for "defined as".

2. Perishable Inventory Systems with Setup Costs

We formally describe the stochastic periodic-review perishable inventory system with setup costs over a planning horizon T (possibly infinite), indexed by $t = 1, \ldots, T$. The product lifetime is $m \geq 2$, i.e., items perish after staying in inventory for m periods if not consumed. The ordering lead time is 0 (see, e.g., Karaesmen et al. (2011)).

Cost structure. The unit holding and shortage costs are denoted by h and b, respectively, and the unit outdating cost is θ. Without loss of generality, we assume the unit purchasing cost $c=0$ (see a detailed cost transformation in Chao et al. (2015b)). In addition, there is a setup cost K that is incurred whenever an order is placed. The discount factor is $\alpha \in [0,1]$ (but α is strictly less than 1 if $T=\infty$), and the firm's objective is to minimize the expected total discounted cost. (Our model and results can allow for non-stationary setup costs K_t as long as they satisfy $\alpha K_{t+1} \le K_t$.)

Demand structure. We adopt the same demand structure as that of Chao et al. (2015a), where the demands $D_1, \ldots, D_T$ are *associated* (which is defined formally below). At the beginning of each period t, the manager observes an *information set* denoted by $\mathscr{F}_t$. The information set $\mathscr{F}_t$ contains the past information accumulated up to the beginning of period t, including the realized demands $d_1, \ldots, d_{t-1}$ in the first $t-1$ periods and possibly some other exogenous information (e.g., state of economy, weather, etc). The information set $\{\mathscr{F}_t \mid t \in [1,T]\}$ form the filtration over a probability space $(\Omega, \mathscr{F}, P)$. We assume that the conditional expectations of all relevant quantities, given $\mathscr{F}_t$, are well-defined. The formal definition of *associated* demand process is as follows.

Definition 1 (Chao et al. (2015a)). A stochastic demand process $\{D_t;\, t=1,2,\ldots\}$ with filtration $\{\mathscr{F}_t, t \ge 1\}$ is called associated if, conditional on $\mathscr{F}_t$, the random demand vector $\mathbf{D}_t = (D_t, D_{t+1}, \ldots, D_T)$ satisfies $\mathbb{E}[f(\mathbf{D}_t)g(\mathbf{D}_t)] \ge \mathbb{E}[f(\mathbf{D}_t)]\ \mathbb{E}[g(\mathbf{D}_t)]$ for all non-decreasing (or non-increasing) functions f and g for which the expectations $\mathbb{E}[f(\mathbf{D}_t)]$, $\mathbb{E}[g(\mathbf{D}_t)]$, $\mathbb{E}[f(\mathbf{D}_t)g(\mathbf{D}_t)]$ exist.

The class of associated processes includes not only independent demand processes, but also most time-series demand models such as autoregressive (AR) and autoregressive moving average (ARMA) demand models (Box et al. (2008)), multiplicative auto-regression model (Levi et al. (2008)), demand forecast updating models such as martingale models for forecast evolution (MMFE) (Heath and Jackson (1994)), demand processes with advance demand information (ADI) (Gallego and Özer (2001)), as well as economic-state driven demand processes such as Markov modulated demand processes with stochastically monotone transition matrix, among others. We refer interested readers to Chao et al. (2015a) for more discussions.

System dynamics. In perishable inventory systems, any inventory unit that stays in the system for m periods without meeting the demand expires and exits the system. Thus, we use a vector to keep track of the inventory age information, which results in a multidimensional state space.

At the beginning of each period t, $t=1,2,\ldots,T$, we are endowed with the realized information f_t. (Here we abuse the notation to use f_t to denote the realized information up to time t, which should be distinguished with the filtration $\mathscr{F}_t$.) The starting inventory at period t is $\mathbf{x}_t = (x_{t,1}, \ldots, x_{t,m-1})$, where for $i = 1, \ldots, m-2$, $x_{t,i}$ is the inventory level of product whose remaining lifetime is i periods, and $x_{t,m-1}$ is the number of product whose remaining lifetime is $m-1$ minus the number of backlogs. After receiving the order q_t in period t, the random demand D_t will be realized (denote its realization by d_t) and satisfied to the maximum extent by FIFO issuing policy, i.e., the oldest inventory meets demand first. Following the convention by Nahmias (1975), we assume the inventory units that will perish at the end of this period also incur a holding cost. The discounted one-period cost is $\alpha^{t-1}\Big(K\cdot\mathbb{1}(q_t>0)+h(y_t-d_t)^+ + b(d_t-y_t)^+ + \theta e_t\Big)$, where $y_t=\sum_{i=1}^{m-1} x_{t,i}+q_t$ is the total inventory level (after ordering) in period t, and $e_t=(x_{t,1}-d_t)^+$ is the outdating inventory. Then, the system proceeds to period $t+1$ with $\mathbf{x}_{t+1}$ given by

$$x_{t+1,j}=\left(x_{t,j+1}-\left(d_t-\sum_{i=1}^{j}x_{t,i}\right)^+\right)^+, \text{ for } 1\le j\le m-2, \tag{1}$$

$$x_{t+1,m-1}=q_t-\left(d_t-\sum_{i=1}^{m-1}x_{t,i}\right)^+. \tag{2}$$

For simplicity we assume that the system is initially empty. Then the total expected discounted cost for any given FIFO issuance policy P can be written as

$$\mathscr{C}(P)=\mathbb{E}\left[\sum_{t=1}^{T}\alpha^{t-1}\Big(K\cdot\mathbb{1}(Q_t>0)+h(Y_t-D_t)^+ + b(D_t-Y_t)^+ + \theta E_t\Big)\right]. \tag{3}$$

Note that, the quantities Q_t, Y_t, E_t and X_t depend on the policy P, and we shall make the dependency explicit whenever necessary, i.e., writing them as Q_t^P, Y_t^P, E_t^P and X_t^P.

3. Randomized Proportional-Balancing Policy

We present the randomized proportional-balancing policy (denoted by RPB) and its worst-case performance guarantee. We start by reviewing the nested marginal holding and outdating cost accounting scheme used in Chao et al. (2015b,a). The main idea is to decompose the total cost in terms of the marginal costs of individual decisions (first used in Levi et al. (2007)). That is, we associate the decision in period t with its affiliated cost contributions to the system. These costs are only affected by future demands but not by future decisions. These marginal costs may include costs (associated with the decision) incurred in both the current and subsequent periods.

3.1. Review of the marginal cost accounting scheme

Under FIFO issuing policy, after a unit is ordered, the number of periods that it stays in the system is not affected by future decisions but only affected by future demands. In other words, the total expected holding and outdating cost of any unit is determined in the period in which that unit is ordered.

Marginal holding and outdating costs. Suppose an order quantity $q_t^P \geq 0$ is placed by a policy P in period t. Then the total expected marginal holding and outdating cost of those q_t^P units can be computed via a recursive equation. For $t \leq s < t+m$, let $B_{[t,s)}(\mathbf{x}_t^P)$ denote the total number of outdated units from period t up to period $s-1$ when the starting inventory at the beginning of period t is $\mathbf{x}_t^P$. With the convention that $B_{[t,t)}(\mathbf{x}_t^P) \equiv 0$, we can write $B_{[t,s)}(\mathbf{x}_t^P)$ recursively by

$$B_{[t,s)}(\mathbf{x}_t^P) = \max\left\{ \sum_{j=1}^{s-t} x_{t,j}^P - D_{[t,s)},\ B_{[t,s-1)}(\mathbf{x}_t^P) \right\}, \quad t < s \leq t+m-1. \tag{4}$$

Using $B_{[t,s)}(\mathbf{x}_t^P)$, the nested marginal holding cost H_t^P can be written as, for $1 \leq t \leq T$,

$$H_t^P = H_t^P(q_t^P; \mathbf{x}_t^P) \triangleq h \sum_{s=t}^{(t+m-1)\wedge T} \alpha^{s-1} \Big(q_t^P - \big(D_{[t,s]} + B_{[t,s)}(\mathbf{x}_t^P) - \sum_{j=1}^{m-1} x_{t,j}^P \big)^+ \Big)^+, \tag{5}$$

and the nested marginal outdating cost Θ_t^P can be written as, for $1 \leq t \leq T-m+1$,

$$\Theta_t^P = \Theta_t^P(q_t^P; \mathbf{x}_t^P) \triangleq \alpha^{t+m-2} \theta E_{t+m-1} = \alpha^{t+m-2} \theta \Big(q_t^P + \sum_{j=1}^{m-1} x_{t,j}^P - B_{[t,t+m-1)}(\mathbf{x}_t^P) - D_{[t,t+m-1]} \Big)^+, \tag{6}$$

and $\Theta_t^P \equiv 0$ for $t = T-m+2, \ldots, T$ as the ordered units do not expire within the planning horizon. It is clear that both $H_t^P(\cdot)$ and $\Theta_t^P(\cdot)$ are increasing in q_t^P.

Backlogging cost and setup cost. It is clear that in period t, no future marginal backlogging cost is caused by the current order q_t^P, since any under-ordering can be corrected by subsequent orders. As a result, the marginal backlogging cost in period t is the same as the conventional backlogging cost in period t, which can be written as

$$\Pi_t^P = \Pi_t^P(q_t^P; \mathbf{x}_t^P) \triangleq \alpha^{t-1}\, b \left(D_t - \sum_{i=1}^{m-1} x_{t,i}^P - q_t^P \right)^+. \tag{7}$$

Moreover, the setup cost in any period t is also only affected by q_t^P, which is $\alpha^{t-1} K \cdot \mathbb{1}(q_t^P > 0)$.

System total cost. Since the system is initially empty, the expected total system cost $\mathscr{C}(P)$ of policy P can be obtained by summing (5), (6), (7) and the setup costs over t from 1 to T, and then taking expectations. Thus, (3) can be rewritten as

$$\mathscr{C}(P) = \mathbb{E}\left[\sum_{t=1}^{T} \Big(H_t^P(q_t^P; \mathbf{x}_t^P) + \Pi_t^P(q_t^P; \mathbf{x}_t^P) + \Theta_t^P(q_t^P; \mathbf{x}_t^P) + \alpha^{t-1} K \cdot \mathbb{1}(q_t^P > 0) \Big) \right], \tag{8}$$

and the system dynamics are governed by (1) and (2).

3.2. RPB policy

To describe the randomized proportional-balancing (RPB) policy, we modify the definition of the information set f_t to also include the implemented decisions of the randomized policy up to period $t-1$. To determine whether or not to place an order, and how much to order, in period t, RPB computes the following quantities:

(a) Compute $\hat{q}_t$ that balances the conditional expected marginal holding and outdating costs of these units against the conditional expected backlogging cost in period t. That is, $\hat{q}_t$ solves

$$\tfrac{mh+\theta}{2(m-1)h+\theta}\mathbb{E}\left[H_t^{RPB}(\hat{q}_t)+\Theta_t^{RPB}(\hat{q}_t)\mid f_t\right]=\mathbb{E}\left[\Pi_t^{RPB}(\hat{q}_t)\mid f_t\right]\triangleq\theta_t, \tag{9}$$

where $\hat{q}_t$ is referred to as the *proportional balancing quantity*, and θ_t the *proportional balancing cost*. The solution to (9) is unique and can be computed efficiently via bisection search.

(b) Compute the (proportional) balancing-K-quantity $\tilde{q}_t$ that solves

$$\tfrac{mh+\theta}{2(m-1)h+\theta}\mathbb{E}\left[H_t^{RPB}(\tilde{q}_t)+\Theta_t^{RPB}(\tilde{q}_t)\mid f_t\right]=K.$$

That is, ordering $\tilde{q}_t$ that balances a proportion of the marginal holding and outdating costs with setup cost K.

(c) Compute $\mathbb{E}[\Pi_t^{RPB}(\tilde{q}_t)\mid f_t]$, i.e., the conditional expected backlogging cost in period t if one orders the balancing-K-quantity $\tilde{q}_t$ in period t.

(d) Compute $\mathbb{E}[\Pi_t^{RPB}(0)\mid f_t]$, i.e., the conditional expected backlogging cost in period t resulting from not ordering in period t.

Description of RPB policy. Let P_t denote the probability that RPB policy places an order in period t which is a-priori random, and let $p_t=(P_t\mid f_t)$. Given f_t, the RPB policy determines the ordering decisions according to two cases below:

Case (I). If the proportional balancing cost exceeds K, i.e., $\theta_t\geq K$, then the RPB policy orders the balancing quantity $q_t^{RPB}=\hat{q}_t$ in period t with probability $p_t=1$.

Case (II). If the proportional balancing cost is less than K, i.e., $\theta_t<K$, then the RPB policy orders the balancing-K-quantity (i.e., $q_t^{RPB}=\tilde{q}_t$) in period t with probability p_t and orders nothing with probability $1-p_t$, where the probability p_t is computed by solving equation

$$p_tK=p_t\cdot\mathbb{E}[\Pi_t^{RPB}(\tilde{q}_t)\mid f_t]+(1-p_t)\cdot\mathbb{E}[\Pi_t^{RPB}(0)\mid f_t]. \tag{10}$$

This completes the description of the RPB policy. We offer some intuitive explanations as follows. In Case (I), that is when $\theta_t\geq K$, the fixed ordering cost K is not dominant compared to the other cost components, hence the policy only strives to achieve (the right) balance between marginal holding, outdating, and shortage costs. In Case (II), the underlying reasoning behind the choice of the particular randomization in (10) is that, the policy attempts to balance between the three cost components, namely, holding and outdating cost, backlogging cost and setup cost associated with period t. In particular, since we order the balancing-K-quantity with probability p_t and do not order anything with probability $1-p_t$, the conditional expected proportional holding and outdating cost for this case is

$$\begin{aligned}
&\tfrac{mh+\theta}{2(m-1)h+\theta}\mathbb{E}[H_t^{RPB}(q_t^{RPB})+\Theta_t^{RPB}(q_t^{RPB}) \mid f_t] \\
&= \tfrac{mh+\theta}{2(m-1)h+\theta}\left\{p_t\mathbb{E}[H_t^{RPB}(\tilde{q}_t)+\Theta_t^{RPB}(\tilde{q}_t) \mid f_t]+(1-p_t)\mathbb{E}[H_t^{RPB}(0)+\Theta_t^{RPB}(0) \mid f_t]\right\} \\
&= p_t K.
\end{aligned} \tag{11}$$

By the selection of p_t in (10), the conditional expected backlogging cost is

$$\mathbb{E}[\Pi_t^{RPB}(q_t^{RPB}) \mid f_t] = p_t\mathbb{E}[\Pi_t^{RPB}(\tilde{q}_t) \mid f_t]+(1-p_t)\mathbb{E}[\Pi_t^{RPB}(0) \mid f_t] = p_t K. \tag{12}$$

Since p_t is the ordering probability in Case (II), the expected fixed ordering cost is also $p_t K$.

It can be shown the p_t that solves (10) is

$$0 \le p_t = \frac{\mathbb{E}[\Pi_t^{RPB}(0) \mid f_t]}{K-\mathbb{E}[\Pi_t^{RPB}(\tilde{q}_t) \mid f_t]+\mathbb{E}[\Pi_t^{RPB}(0) \mid f_t]} < 1, \tag{13}$$

where the above inequalities follow from the fact that $\theta_t < K$ and $\tilde{q}_t > \hat{q}_t$, which implies that $\mathbb{E}[\Pi_t^{RPB}(\tilde{q}_t) \mid f_t] < \mathbb{E}[\Pi_t^{RPB}(\hat{q}_t) \mid f_t] = \theta_t < K$. We remark that when the setup cost $K=0$, the RPB policy reduces exactly to the PB policy proposed in Chao et al. (2015b), i.e., RPB is a generalization of PB in the presence of setup costs.

The following is the main theoretical result of this paper.

THEOREM 1. *When the demand process is associated, the RPB policy for perishable inventory system with setup cost and $m \ge 2$ periods of product lifetime has a worst-case performance guarantee of $\left(3+\frac{(m-2)h}{mh+\theta}\right)$, i.e., for any instance of the problem, the expected cost of the RPB policy is at most $\left(3+\frac{(m-2)h}{mh+\theta}\right)$ times the expected cost of an optimal policy.*

The balancing ratio on the left hand side of (9) is chosen such that the resulting RPB policy admits the tightest worst-case performance guarantee. Suppose an arbitrary balancing ratio $\beta \in (0,1]$ is used to construct RPB. Then one can show that it admits a worst-case performance guarantee of $(2\beta+1)/\min\{\beta,\beta_0\}$, where $\beta_0 = \frac{mh+\theta}{2(m-1)h+\theta}$ (which is obtained from (22) in our worst-case analysis). This worst-case performance guarantee is minimized when $\beta=\beta_0$.

Remark on discrete demand and order quantities: If the demand and order quantities are discrete, then we can always write $\hat{q}_t = \hat{\lambda}_t \hat{q}_t^1 + (1-\hat{\lambda}_t)\hat{q}_t^2$, where $\hat{q}_t^1$ and $\hat{q}_t^2 = \hat{q}_t^1 + 1$ are consecutive integers with $\hat{q}_t^1 \le \hat{q}_t < \hat{q}_t^2$ and $0 < \hat{\lambda}_t \le 1$. Similarly, we write $\tilde{q}_t = \tilde{\lambda}_t \tilde{q}_t^1 + (1-\tilde{\lambda}_t)\tilde{q}_t^2$, where $\tilde{q}_t^1$ and $\tilde{q}_t^2 = \tilde{q}_t^1 + 1$ are consecutive integers with $\tilde{q}_t^1 \le \tilde{q}_t < \tilde{q}_t^2$ and $0 < \tilde{\lambda}_t \le 1$. In Case (I), RPB orders either $\hat{q}_t^1$ units (with probability $\hat{\lambda}_t$) or $\hat{q}_t^2$ units (with probability $1-\hat{\lambda}_t$). In Case (II), RPB orders either $\tilde{q}_t^1$ units (with probability $p_t\tilde{\lambda}_t$), or $\tilde{q}_t^2$ units (with probability $p_t(1-\tilde{\lambda}_t)$), or zero unit (with probability $1-p_t$). This randomized procedure will not affect our worst-case performance guarantee, and we refer readers to the detailed discussions in §4.3 of Levi et al. (2007) and also §6 of Shi et al. (2014).

4. Worst-Case Performance Analysis

We carry out a worst-case performance analysis of RPB policy. For simplicity, we provide the proofs with a discount factor $\alpha = 1$ (albeit the analysis holds under general $\alpha \in [0,1]$).

We first describe a concept called "*trimmed inventory level*" (see Chao et al. (2015a)). The trimmed inventory level, denoted by $Y_{t,s}$ for any $s \ge t \ge 1$, is defined as the inventory at the beginning of period s for the products that are ordered in period t or earlier. Equivalently, $Y_{t,s}$ is the inventory level Y_s (after ordering) at the beginning of period s less the total order quantity in periods $t+1,\dots,s$, i.e., $Y_{t,s} = Y_s - \sum_{s'=t+1}^{s} Q_{s'}$. By definition, $Y_{t,t} = Y_t$ and we also have

$$Y_{t,s} = Y_{t,s'} - D_{[s',s)} - E_{[s',s)}, \quad t \le s' < s \le t+m-1, \tag{14}$$

$$Y_{t,s} = Y_{t,t+m-1} - D_{[t+m-1,s)}, \qquad s > t+m-1. \tag{15}$$

Note that $Y_{t,s}$ can be negative when s is large enough. These trimmed inventory levels serve as a generalization of the traditional inventory levels, as they provide critical partial information on the age of the on-hand products. Due to the nature of perishable systems, it is impossible to quantify the effect of the decision made in the current period t on future costs only through the traditional total inventory level Y_t. The trimmed inventory levels provide a tractable way to analyze this effect, and also provide the *right* framework for coupling the marginal holding, outdating and setup costs in different systems. Furthermore, for any realization f_t, if $p_t < 1$, then we denote $\overline{\text{RPB}}$ as the policy that orders exactly the same as RPB before period t, but orders $\tilde{q}_t$ at t (instead of using a randomized decision between 0 and $\tilde{q}_t$), and orders 0 afterwards.

We now compare the costs of RPB policy against that of an optimal policy, denoted by OPT. For each realization f_T, we partition all the periods into the following sets:

$$\mathscr{T}_{1H} = \left\{t : p_t = 1 \text{ and } y_t^{OPT} > y_t^{RPB}\right\}; \tag{16}$$

$$\mathscr{T}_{1\Pi} = \left\{t : p_t = 1 \text{ and } y_t^{OPT} \le y_t^{RPB}\right\}; \tag{17}$$

$$\mathscr{T}_{2\Pi} = \left\{t : p_t < 1 \text{ and } x_t^{RPB} \geq y_t^{OPT}\right\}; \tag{18}$$

$$\mathscr{T}_{2H} = \left\{t : p_t < 1 \text{ and } x_t^{RPB} < y_t^{OPT} \text{ and } \exists s \in [t, (t+m-1)\wedge T], y_{t,s}^{OPT} > \overline{y_{t,s}^{RPB}}\right\}; \tag{19}$$

$$\mathscr{T}_{2M} = \left\{t : p_t < 1 \text{ and } x_t^{RPB} < y_t^{OPT} \text{ and } \forall s \in [t, (t+m-1)\wedge T], y_{t,s}^{OPT} \leq \overline{y_{t,s}^{RPB}}\right\}. \tag{20}$$

Lemma 1 shows that the backlogging cost of OPT can cover that of RPB within set $\mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi}$, which is formally stated below. We relegate its detailed proof to the electronic companion.

Lemma 1. *For each sample path f_T, we have*

$$\sum_{t=1}^{T} \Pi_t^{OPT} \geq \sum_{t=1}^{T} \left[\Pi_t^{RPB} \cdot \mathbb{1}\left(t \in \mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi}\right)\right]. \tag{21}$$

Lemma 2 shows how to amortize the sum of marginal holding and outdating costs against that of OPT, which follows identical arguments used in Chao et al. (2015a); hence its proof is omitted.

Lemma 2. *For each sample path f_T, we have*

$$\sum_{t=1}^{T}\left(H_t^{OPT} + \Theta_t^{OPT}\right) \geq \frac{mh+\theta}{2(m-1)h+\theta} \sum_{t=1}^{T}\left[(H_t^{RPB} + \Theta_t^{RPB}) \cdot \mathbb{1}\left(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H}\right)\right]. \tag{22}$$

The next result is a key lemma for this paper. It suggests that the total setup cost incurred by OPT can cover the total setup cost incurred by RPB for periods within the set $\mathscr{T}_{2M}$. Note that $\mathscr{T}_{2M}$ is the most "problematic" set of periods because we are unsure whether the trimmed ending inventory levels of RPB are higher or lower than those of OPT due to randomized decisions.

Lemma 3. *For each sample path f_T, the following inequality holds for the OPT and RPB policies,*

$$\sum_{t=1}^{T} \mathbb{1}(q_t^{OPT} > 0) \geq \sum_{t=1}^{T} \mathbb{1}(q_t^{RPB} > 0 \text{ and } t \in \mathscr{T}_{2M}). \tag{23}$$

Proof. For a fixed sample path f_T, each period t deterministically belongs to one of the sets in $\mathscr{T}_{1H}$, $\mathscr{T}_{1\Pi}$, $\mathscr{T}_{2\Pi}$, $\mathscr{T}_{2H}$ and $\mathscr{T}_{2M}$. Denote all the periods in $\mathscr{T}_{2M}$ with positive ordering quantities by

$$\left\{t \in \mathscr{T}_{2M} : q_t^{RPB} > 0\right\} = \{t_1, t_2, \ldots, t_n\}.$$

See Figure 1 for an illustration of all such periods.

If this set is empty, i.e., $n = 0$, then (23) trivially holds. Otherwise, it follows from the definition of $\{t_1, t_2, \ldots, t_n\}$ that the right hand side of (23) is equal to n. Thus, it suffices to prove that OPT places at least n orders. We first show that,

$$\sum_{t=1}^{t_1} \mathbb{1}(q_t^{OPT} > 0) \geq 1. \tag{24}$$

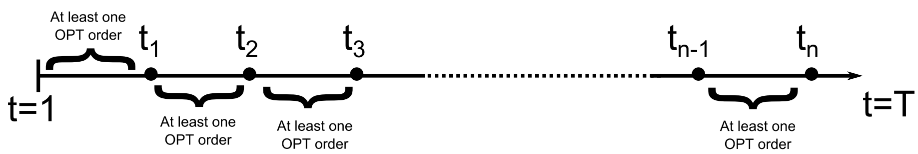

Figure 1 Illustration of both Inequality (24) and Inequality (25) for sample path f_T.

This is true because both OPT and RPB have the same starting inventory and are facing the same demand, if OPT has not yet placed any order before or in period t_1, then we cannot have $y_{t_1}^{OPT} > x_{t_1}^{RPB}$. Hence, OPT places at least one order as shown in Figure 1.

If $n = 1$, then (23) is the same as (24) which has already been proved. Now suppose $n \geq 2$. We want to show that

$$\sum_{t=t_i+1}^{t_{i+1}} \mathbb{1}(q_t^{OPT} > 0) \geq 1, \qquad \text{for all } i = 2, 3, \ldots, n. \tag{25}$$

That is, OPT places at least one order in each and every interval $\{t_i + 1, \ldots, t_{i+1}\}$. (Figure 1 provides an illustration of (25).)

Fix an $i \in \{2, 3, \ldots, n\}$. Since $t_i \in \mathscr{T}_{2M}$ and the RPB policy orders in t_i, it follows from the definition of $\mathscr{T}_{2M}$ that for all $s \in \{t_i, t_i + 1, \ldots, t_i + m - 1\}$, it holds that $y_{t_i,s}^{OPT} \leq y_{t_i,s}^{RPB}$. This implies $y_{t_i,s}^{OPT} \leq y_{t_i,s}^{RPB}$ for all $s \in \{t_i, t_i + 1, \ldots, T\}$ since by (15), the difference of $y_{t_i,s}^{OPT}$ and $y_{t_i,s}^{RPB}$ is fixed when $s > t_i + m - 1$. This shows

$$y_{t_{i+1},t_i}^{OPT} \leq y_{t_{i+1},t_i}^{RPB}. \tag{26}$$

We prove (25) by contradiction. Assume that, on the contrary, OPT does not place any order from $t_i + 1$ to t_{i+1}. By the definition of trimmed on-hand inventory, we can see that, for any policy P, y_{t_{i+1},t_i}^{P} is always a lower bound of $x_{t_{i+1}}^{P}$ and $y_{t_{i+1}}^{P}$. Thus, since we have assumed that OPT makes no order from t_i to t_{i+1}, we have $y_{t_{i+1}}^{OPT} = y_{t_{i+1},t_i}^{OPT}$. Similarly, we have $y_{t_{i+1},t_i}^{RPB} \leq x_{t_{i+1}}^{RPB}$. Together with (26), we have

$$y_{t_{i+1}}^{OPT} = y_{t_{i+1},t_i}^{OPT} \leq y_{t_{i+1},t_i}^{RPB} \leq x_{t_{i+1}}^{RPB}, \tag{27}$$

which contradicts to the fact that $t_{i+1} \in \mathscr{T}_{2M}$ with $x_{t_{i+1}}^{RPB} < y_{t_{i+1}}^{OPT}$. This proves that (25) holds for all $i = 2, 3, \ldots, n$.

Finally, summing up (24) and (25) for all $i = 2, 3, \ldots, n$ proves (23). **Q.E.D.**

Remark 1. A key step in the proof of Lemma 3 is the construction of set $\mathscr{T}_{2M}$ defined in (20). This definition is very different from the original set $\mathscr{T}_{2M}^{\mathbf{LS}}$ defined in Levi and Shi (2013) for non-perishable inventory systems, which only requires the relationship between the aggregate inventory positions $y_t^{OPT} \leq y_t^{\overline{RPB}}$. This is because for non-perishable inventory systems, the condition on aggregate inventory levels is sufficient to establish a similar result as (23). But for perishable

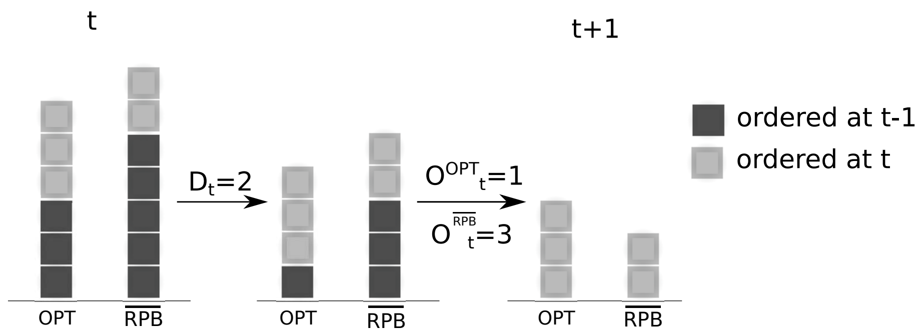

Figure 2 A counterexample with $m=2$ to show the importance of age distribution in set decomposition

inventory systems where the age information becomes critical, $y_t^{OPT} \leq y_t^{\overline{RPB}}$ is not sufficient to guarantee the same type of result. For example, consider the following simplest nontrivial setting with $m=2$, the set $\mathscr{T}_{2M}$ defined in (20) can be written as

$$\mathscr{T}_{2M} = \left\{ t : p_t < 1 \text{ and } x_t^{RPB} < y_t^{OPT}, y_t^{OPT} \leq y_t^{\overline{RPB}}, \text{ and } y_{t,t+1}^{OPT} \leq y_{t,t+1}^{\overline{RPB}} \right\}.$$

However, the set $\mathscr{T}_{2M}^{\mathbf{LS}}$ used in Levi and Shi (2013) is given by

$$\mathscr{T}_{2M}^{\mathbf{LS}} = \left\{ t : p_t < 1 \text{ and } x_t^{RPB} < y_t^{OPT}, y_t^{OPT} \leq y_t^{\overline{RPB}} \right\}.$$

Now consider the following scenario as depicted in Figure 2. In period t, OPT has 3 inventory units of age 1 and 3 new inventory units of age 0; RPB has 5 inventory units of age 1 and 2 new inventory units of age 0. Suppose the demand is 2 in period t and is satisfied using FIFO issuance. At the end of period t, OPT has 1 outdating unit and $\overline{\text{RPB}}$ has 3 outdating units. In period $t+1$, OPT (having 3 units on-hand) overtakes $\overline{\text{RPB}}$ (having only 2 units on-hand).

Note that in this case, period t satisfies the conditions in $\mathscr{T}_{2M}^{\mathbf{LS}}$ but does not satisfy the last condition in $\mathscr{T}_{2M}$. Thus, though the inventory level of $\overline{\text{RPB}}$ is higher in period t, after one period we would have $y_{t,t+1}^{OPT} > y_{t,t+1}^{\overline{RPB}}$ (because $\overline{\text{RPB}}$ has more outdating inventory units in period t). The difficulty arises because OPT needs not to make any order in period $t+1$ to overtake the inventory level of $\overline{\text{RPB}}$. Therefore, it cannot be argued that OPT has placed more orders and the result in Lemma 3 cannot be established for the set $\mathscr{T}_{2M}^{\mathbf{LS}}$. □

By the definition of $\mathscr{T}_{2H}$, the indicator function $\mathbb{1}(t \in \mathscr{T}_{2H})$ (i.e., whether a period t belongs to the set $\mathscr{T}_{2H}$) cannot be determined only by $\mathscr{F}_t$ (the information up to time t). Hence in order to compare the expected costs between RPB and OPT, we require the demand process to be associated, a concept introduced in Chao et al. (2015a) that embodies a notion of positive correlation between future demands. Lemma 4 establishes that under that condition, in each period t, the expected marginal holding and outdating costs with known information $\mathbb{1}(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H})$ is in fact greater than or equal to those without knowing the information.

LEMMA 4. *If the demand process $D_1, \ldots, D_T$ is associated, then for each period $t = 1, \ldots, T$,*

$$\begin{aligned} &\mathbb{E}\Big[\mathbb{1}\Big(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H}\Big) \mid \mathscr{F}_t\Big] \cdot \mathbb{E}\Big[H_t^{RPB} + \Theta_t^{RPB} \mid \mathscr{F}_t\Big] \\ \le\; &\mathbb{E}\Big[\mathbb{1}\Big(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H}) \cdot (H_t^{RPB} + \Theta_t^{RPB}\Big) \mid \mathscr{F}_t\Big]. \end{aligned} \tag{28}$$

Proof. Conditioning on $\mathscr{F}_t$, we have the following four cases (a) $p_t = 1$, (b) $p_t < 1$ and $x_t^{RPB} \ge y_t^{OPT}$, (c) $p_t < 1$, $x_t^{RPB} < y_t^{OPT}$ and $y_t^{\overline{RPB}} < y_t^{OPT}$, and (d) $p_t < 1$, $x_t^{RPB} < y_t^{OPT} \le y_t^{\overline{RPB}}$.

First, we notice that $\mathbb{P}(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H} \mid \mathscr{F}_t)$ takes value 0 or 1 if and only if either case (a) or (b) or (c) happens; in such cases, it is straightforward to verify that (28) holds.

Thus, in the remainder of this proof, we only focus our attention on the non-trivial case (d), which implies that $0 < \mathbb{P}(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H} \mid \mathscr{F}_t) = \mathbb{P}(t \in \mathscr{T}_{2H} \mid \mathscr{F}_t) < 1$.

According to the definition of marginal holding and outdating costs, it is clear that conditioning on $\mathscr{F}_t$, $H_t^{RPB} + \Theta_t^{RPB}$ is decreasing in future demands $(D_t, \ldots, D_{t+m})$. If we can show that $\mathbb{1}\big(t \in \mathscr{T}_{2H} \mid \mathscr{F}_t\big)$ is also decreasing in future demands $(D_t, \ldots, D_{t+m})$, then (28) follows since the demand process is associated (see Essary et al. (1967)).

To show that $\mathbb{1}\big(t \in \mathscr{T}_{2H} \mid \mathscr{F}_t\big)$ is decreasing in future demands $(D_t, \ldots, D_{t+m})$, we first define what we call *switching* events, for $t \le s < (t + m - 1) \wedge T$,

$$A_{t,s} \triangleq \Big\{[Y_{t,s}^{\overline{RPB}} \ge Y_{t,s}^{OPT}] \cap [Y_{t,s+1}^{\overline{RPB}} < Y_{t,s+1}^{OPT}]\Big\}.$$

Given case (d), by the definition of $\mathscr{T}_{2H}$, the event $\{t \in \mathscr{T}_{2H}\}$ happens if and only if the switching event $A_{t,s}$ occurs for some $s \in [t, (t + m - 1) \wedge T)$.

We claim that the occurrence of the switching event $A_{t,s}$ implies that $E_s^{\overline{RPB}} > 0$, i.e., $\overline{\text{RPB}}$ must have some outdating units in period s. (To see this, $A_{t,s}$ happens if OPT overtakes the trimmed inventory level of $\overline{\text{RPB}}$ in period $s + 1$, which is impossible if no inventory units in $\overline{\text{RPB}}$ expires in period t as both systems face the same demands.) The fact that $E_s^{\overline{RPB}} > 0$ implies that no inventory units ordered after period $s - m + 1$ are consumed by demand at the beginning of period $s + 1$. Hence, we must have $Y_{t,s+1}^{\overline{RPB}} = q_{[s-m+2,t]}^{\overline{RPB}}$ with probability 1. By the same argument, we can see that if $E_s^{\overline{RPB}} > 0$ and $Y_{t,s+1}^{\overline{RPB}} = q_{[s-m+2,t]}^{\overline{RPB}}$ with probability 1, then $Y_{t,s+1}^{\overline{RPB}} < Y_{t,s+1}^{OPT}$. Thus, the switching event $A_{t,s}$ can be rewritten as

$$A_{t,s} = \Big\{[Y_{t,s}^{\overline{RPB}} \ge Y_{t,s}^{OPT}] \cap [E_s^{\overline{RPB}} > 0] \cap [q_{[s-m+2,t]}^{\overline{RPB}} < Y_{t,s+1}^{OPT}]\Big\}. \tag{29}$$

Thus, we can rewrite the event as

$$\{t \in \mathscr{T}_{2H}\} = \left\{ \bigcup_{s=t}^{(t+m-1)\wedge T-1} [E_s^{\overline{RPB}} > 0] \bigcap \left[q_{[s-m+2,t]}^{\overline{RPB}} < Y_{t,s+1}^{OPT} \right] \right\}. \tag{30}$$

Since $E_s^{\overline{RPB}}$ are the outdating units in period s (which were ordered in period $s - m + 1 \le t$), it is clear that given $\mathscr{F}_t$, $E_s^{\overline{RPB}}$ is decreasing in the demand process after t, i.e., $\mathbb{1}(E_s^{\overline{RPB}} > 0 \mid \mathscr{F}_t)$ is decreasing in $(D_t, D_{t+1}, \ldots, D_T)$. Because the trimmed inventory level $Y_{t,s+1}^{OPT}$ is decreasing in $(D_t, D_{t+1}, \ldots, D_T)$, and $q_{[s-m+2,t]}^{\overline{RPB}}$ is already known deterministically at time t, it follows that $\mathbb{1}\left(q_{[s-m+2,t]}^{\overline{RPB}} < Y_{t,s+1}^{OPT} \mid \mathscr{F}_t\right)$ is also decreasing in $(D_t, D_{t+1}, \ldots D_T)$. Given case (d), by (30), we conclude that $\mathbb{1}\big(t \in \mathscr{T}_{2H} \mid \mathscr{F}_t\big)$ is indeed decreasing in future demands $(D_t, \ldots, D_{t+m})$. **Q.E.D.**

Remark 2. The main idea behind the proof of Lemma 4 is to apply the concept of associated random variables. It can be shown that $\mathbb{1}(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H})$ is a decreasing function of future demands. Together with the fact that $H_t^{RPB} + \Theta_t^{RPB}$ is also a decreasing function of future demands, the result then follows from the properties of associated stochastic processes. To (intuitively) see why $\mathbb{1}(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H})$ is a decreasing function of future demands, we consider the same example as shown in Figure 2. If d_t is equal to 4, then we have $y_{t,t+1}^{OPT} = y_{t,t+1}^{RPB} = 2$ and in this case $t \in \mathscr{T}_{2M}$. We can also see that if $d_t < 4$, then t falls in the set $\mathscr{T}_{2H}$ while if $d_t \ge 4$, then it falls in the set $\mathscr{T}_{2M}$. This implies that it is more likely to be in the set $\mathscr{T}_{1H} \bigcup \mathscr{T}_{2H}$ when the demand is lower. □

Let Z_t^{RPB} be a random variable defined by

$$Z_t^{RPB} \triangleq \begin{cases} \Lambda_t, \text{ if } p_t = 1; \\ p_t K, \text{ otherwise}, \end{cases} \tag{31}$$

where $\Lambda_t \triangleq \frac{mh+\theta}{2(m-1)h+\theta} \mathbb{E}[H_t^{RPB} + \Theta_t^{RPB} | F_t] = \mathbb{E}[\Pi_t^{RPB} | F_t]$, and p_t is the ordering probability.

Lemma 5 below follows from the standard arguments of conditional expectation, the construction of RPB, and (31). We relegate its detailed proof to the electronic companion.

Lemma 5. *Let $\mathscr{C}(RPB)$ be the expected total cost incurred by the RPB policy. Then we have,*

$$\mathscr{C}(RPB) \le \left[3 + \frac{(m-2)h}{mh+\theta} \right] \sum_{t=1}^{T} \mathbb{E}[Z_t^{RPB}].$$

Combining Lemmas 2, 3, 4, and 5, we are able to prove Theorem 1, and we again relegate its detailed proof to the electronic companion.

5. Computational Experiments

An important question is how well RPB performs empirically. In this section we report on numerical results. All computations were done using Matlab R2014a on a desktop computer with an Intel(R) Xeon(R) CPU E31230 @ 3.20 Ghz.

The purpose of our computational study is two-fold. **(a)** Under i.i.d. demands, we would like to test the empirical performance of RPB against the heuristic policy proposed in Nahmias (1978) (denoted by NA). To the best of our knowledge, NA is the only benchmark heuristic algorithm (without any worst-case performance guarantees) available for perishable inventory control problems with setup costs. This approximate (s, S)-type policy is reported to perform very close to optimality under i.i.d. demands, with average error under 1% and maximum error of 3.7%. The question is whether our RPB policy can achieve, if not better, at least a comparable performance to NA under i.i.d. demands. **(b)** Under correlated demands, since there are no benchmark heuristic algorithms reported in the existing literature, we can only compare the performance of RPB against an optimal policy (denoted by OPT). Thus in our experimental setting under correlated demands we generate a sufficiently rich set of small problem samples, where the optimal policy can be evaluated (using brute-force dynamic programming) with reasonable computational effort to provide a baseline for comparison.

Following Levi and Shi (2013), the proposed RPB policies can be parametrized to obtain a general class of policies, and the worst-case analysis discussed above can be viewed as choosing one parameter value that achieves a worst-case performance guarantee for any problem instance (see the discussion at the end of §3.2). Alternatively, one can try different parameter values for a given problem instance, and identify a parameter that empirically performs the best (in terms of average and/or worst-case performance) for that particular instance. This gives rise to policies that have at least the same performance guarantees, but are likely to perform better empirically, as we refined the parameters according to the specific instance being solved.

5.1. Design of Experiments and Performance Metrics

We have conducted our experiments under both an i.i.d. demand setting and a correlated demand setting, using similar examples as in Chao et al. (2015b). **(a)** Under the i.i.d. demand setting, following Nahmias (1976, 1977), we test two demand distributions, i.e., exponential distribution and Erlang-2 distribution, both with mean 10. **(b)** Under the correlated demand setting, we adopt the Markov modulated demand process (MMDP) with three states of the economy. The MMDP is governed by the state of the economy: poor (1), fair (2), and good (3). If the state of the economy in period t is i $(i = 1, 2, 3)$, then the demand in period t is iD_t, where D_t has mean 10 and follows one of the following distributions: exponential, and Erlang-2. The state of the economy follows a Markov chain with transition probabilities

$$p_{11} = 0.6, p_{12} = 0.3, p_{13} = 0.1, p_{21} = 0.4, p_{22} = 0.2, p_{23} = 0.4, p_{31} = 0.1, p_{32} = 0.3, \text{ and } p_{33} = 0.6.$$

This Markov chain is stochastically monotone (and hence the demand process is associated).

The parameters of our computational model are chosen as follows. The holding cost for all testing instances is normalized to $h = 1$, and the discount factor is $\alpha = 0.95$. We set the planning horizon $T = 20$, the per-unit backlogging cost $b \in \{5, 10, 15\}$, the per-unit outdating cost $\theta \in \{5, 10, 15\}$ and the setup cost $K \in \{5, 10, 15\}$. We also set the product lifetime $m = 2$. (We note that Nahmias (1978) focused on $m = 2$ only and commented on the expensive computational overhead for $m = 3$.) The system is initially empty. We define the *performance ratio* of our RPB policy against a benchmark policy P as $r = \left(\frac{\mathscr{C}(RPB)}{\mathscr{C}(P)} - 1\right) \times 100\%$, where $\mathscr{C}(RPB)$ is the cost of RPB and $\mathscr{C}(P)$ is the cost of policy P. The benchmark policy P is NA (proposed by Nahmias (1978)) under i.i.d. demands whereas P is OPT (solved using brute-force dynamic program) under correlated demands.

5.2. Numerical Results

Table EC.1 from the electronic companion summarizes our numerical results for the i.i.d. demand setting. The performance ratio is very small for all test instances, which suggests that the performance of RPB is comparable (in fact slightly better but not statistically significant) to that of NA under i.i.d. demands. This also suggests both RPB and NA perform extremely close to optimality (in that NA's error is on average under 1% of the optimal solution). We also notice that both algorithms perform better when the setup cost is relatively small. Table EC.2 provides the total cost breakdown for the i.i.d. demand setting (varying the values of K while fixing $b = \theta = 5$). We observe that the setup cost component accounts for a significant portion of the total costs, and the percentage of setup costs increases in K. The ordering frequency is around 70% to 80% with $m = 2$, and decreases in K.

Table EC.3 summarizes our numerical results for the correlated demand setting. The performance ratio is small for all test instances, with average error under 7% and maximum error of 11.28%. This indicates that the RPB policy performs well, and it is significantly better than the theoretical worst-case performance guarantee.

Under both the i.i.d. and correlated demand settings, our numerical results suggest that the performance of RPB is rather insensitive to b or θ, but improves as K gets smaller. We also test the stability of RPB – the coefficient of variation of the total costs due to randomization. We consider the case with exponential demand with mean 10, and evaluate the costs of RPB (with randomized decisions) for the same cost parameters and demand realizations. We find that the coefficient of variation of costs (standard deviation divided by mean cost) is at most 1.57%, suggesting that the adverse effect due to randomization is rather small and RPB is quite stable.

Supplemental Material

An electronic companion to this paper is available at http://or.journal.informs.org/.

Acknowledgments

The authors thank the area editor Chung-Piaw Teo, the anonymous associate editor, and two referees for their constructive comments and suggestions, which helped improve both the content and the exposition of this paper significantly. The research of Huanan Zhang and Cong Shi is partially supported in part by NSF grants CMMI-1362619 and CMMI-1451078. The research of Xiuli Chao is partially supported by NSF grants CMMI-1131249 and CMMI-1362619.

References

Box, G. E. P., G. M. Jenkins, G. C. Reinsel. 2008. *Time Series Analysis: Forecasting and Control*. Wiley, Fourth Edition.

Chao, X., X. Gong, C. Shi, C. Yang, H. Zhang, S. X. Zhou. 2015a. Approximation algorithms for capacitated perishable inventory systems with positive lead times. Working paper, University of Michigan, Ann Arbor. Available at http://www.umich.edu/~shicong/papers/perishable2.pdf.

Chao, X., X. Gong, C. Shi, H. Zhang. 2015b. Approximation algorithms for perishable inventory systems. *Operations Research* **63**(3) 585–601.

Essary, J. D., F. Proschan, D. W. Walkup. 1967. Association of random variables, with applications. *The Annals of Mathematical Statistics* **38**(5) 1466–1474.

Fries, B. 1975. Optimal ordering policy for a perishable commodity with fixed lifetime. *Operational Research* **23**(1) 46–61.

Gallego, G., Ö. Özer. 2001. Integrating replenishment decisions with advance demand information. *Management Science* **47**(10) 1344–1360.

Heath, D. C., P. L. Jackson. 1994. Modeling the evolution of demand forecasts with application to safety stock analysis in production/distribution system. *IIE Transactions* **26**(3) 17–30.

Karaesmen, I. Z., A. Scheller-Wolf, B. Deniz. 2011. Managing perishable and aging inventories: Review and future research directions. G. Karl Kempf, Pınar Keskinocak, Reha Uzsoy, eds., *Planning Production and Inventories in the Extended Enterprise: A State of the Art Handbook, Volume 1*. International Series in Operations Research and Management Science, Vol. 151 (Springer, New York), 393–436.

Levi, R., G. Janakiraman, M. Nagarajan. 2008. A 2-approximation algorithm for stochastic inventory control models with lost-sales. *Mathematics of Operations Research* **33**(2) 351–374.

Levi, R., M. Pál, R. O. Roundy, D. B. Shmoys. 2007. Approximation algorithms for stochastic inventory control models. *Mathematics of Operations Research* **32**(4) 821–838.

Levi, R., C. Shi. 2013. Approximation algorithms for the stochastic lot-sizing problem with order lead times. *Operations Research* **61**(3) 593–602.

Lian, Z., L. Liu. 1999. A discrete-time model for perishable inventory systems. *Annals of Operations Research* **87**(0) 103–116.

Lian, Z., L. Liu, M. F. Neuts. 2005. A discrete-time model for common lifetime inventory systems. *Mathematics of Operations Research* **30**(3) 718–732.

Nahmias, S. 1975. Optimal ordering policies for perishable inventory-II. *Operational Research* **23**(4) 735–749.

Nahmias, S. 1976. Myopic approximations for the perishable inventory problem. *Management Science* **22**(9) 1002–1008.

Nahmias, S. 1977. Higher order approximations for the perishable inventory problem. *Operations Research* **25**(4) 630–640.

Nahmias, S. 1978. The fixed charge perishable inventory problem. *Operations Research* **26**(3) 464–481.

Prastacos, G. P. 1984. Blood inventory management: an overview of theory and practice. *Management Science* **30** 777–800.

Shi, C., H. Zhang, X. Chao, R. Levi. 2014. Approximation algorithms for capacitated stochastic inventory systems with setup costs. *Naval Research Logistics (NRL)* **61**(4) 304–319.

Brief Bio:

Huanan Zhang is a Ph.D. candidate in the Department of Industrial and Operations Engineering at the University of Michigan. His primary research interest lies in stochastic optimization and online learning problems with applications to inventory and supply chain management, and revenue management.

Cong Shi is an assistant professor in the Department of Industrial and Operations Engineering at the University of Michigan. His research lies in stochastic optimization and online learning theory with applications to inventory and supply chain management, and revenue management. He won first prize in the 2009 George Nicholson Student Paper Competition.

Xiuli Chao is a professor in the Department of Industrial and Operations Engineering at the University of Michigan. His recent research interests include stochastic modeling and analysis, inventory control, game applications in supply chains, and data-driven optimization.

This page is intentionally blank. Proper e-companion title page, with INFORMS branding and exact metadata of the main paper, will be produced by the INFORMS office when the issue is being assembled.

Electronic Companion to

"Approximation Algorithms for Perishable Inventory Systems with Setup Costs"

by Zhang, Shi and Chao

EC.1. Computational Complexity of RPB

We show that the RPB policy is computationally very efficient with computational complexity $\mathrm{O}(C^m T)$ for some positive constant C, where m is product lifetime and T is the length of the planning horizon. We also stress that the RPB policy can be efficiently implemented in an online manner, i.e., the decision at any time is computed based only on the current observed state of the system and does not depend on future decisions. This is a desired property if one wishes to avoid the prohibitive computational burden of solving large dynamic programs.

Next we prove that RPB has a computational complexity $\mathrm{O}(C^m T)$. For each period $t = 1, \ldots, T$, the complexity for evaluating the marginal holding and outdating costs is $\mathrm{O}(C^m)$ since there are m layers of integration involved in the exact computation. For practical purposes, evaluating the expected marginal costs using Monte Carlo simulations (e.g., generating 10000 sample paths according to the joint conditional demand distribution) can cut down the computational time dramatically. This computational overhead is unavoidable as the complexity is the same as computing the single-period outdating costs. This suggests that even *myopic policies* that only minimize the current-single-period cost have to incur this computational overhead. Since the complexity for carrying out bisection search is $\mathrm{O}(\log U)$ (where U is an upper bound on the balancing quantities), the algorithm runs in time $\mathrm{O}(C^m T \log U) \approx \mathrm{O}(C^m T)$. In contrast, computing the exact optimal policy using dynamic programming is exponential in the length of the planning horizon T.

EC.2. Omitted Proofs of Lemmas and Theorems

Proof of Lemma 1. By the definition of Π_t and the construction of sets $\mathscr{T}_{1\Pi}$ and $\mathscr{T}_{2\Pi}$, we have

$$\sum_{t\in\mathscr{T}_{1\Pi}\bigcup\mathscr{T}_{2\Pi}} \Pi_t^{RPB} = b \sum_{t\in\mathscr{T}_{1\Pi}\bigcup\mathscr{T}_{2\Pi}} (d_t - y_t^{RPB})^+ \le b \sum_{t\in\mathscr{T}_{1\Pi}\bigcup\mathscr{T}_{2\Pi}} (d_t - y_t^{OPT})^+ \le \sum_{t=1}^{T} \Pi_t^{OPT},$$

where the first inequality holds since $y_t^{OPT} \le y_t^{RPB}$ when $t \in \mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi}$. **Q.E.D.**

Proof of Lemma 5. Using the standard arguments of conditional expectations and by the construction of RPB, we have

$$\mathscr{C}(RPB) = \sum_{t=1}^{T} \mathbb{E}[H_t^{RPB} + \Theta_t^{RPB} + \Pi_t^{RPB} + K \cdot \mathbb{1}(q_t^{RPB} > 0)]$$

$$
\begin{aligned}
&= \sum_{t=1}^{T} \mathbb{E}\left[\mathbb{E}[H_t^{RPB} + \Theta_t^{RPB} + \Pi_t^{RPB} + K \cdot \mathbb{1}(q_t^{RPB} > 0) \mid \mathscr{F}_t]\right] \\
&= \sum_{t=1}^{T} \Big\{\mathbb{E}\Big[\mathbb{E}\Big[\mathbb{1}\Big(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{1\Pi}\Big) \cdot \left(\mathbb{E}[H_t^{RPB} + \Theta_t^{RPB} + \Pi_t^{RPB} + K \cdot \mathbb{1}(q_t^{RPB} > 0)\right) \mid \mathscr{F}_t]\Big]\Big] \\
&\quad + \mathbb{E}\Big[\mathbb{E}\Big[\mathbb{1}\Big(t \in \mathscr{T}_{2H} \bigcup \mathscr{T}_{2M} \bigcup \mathscr{T}_{2\Pi}\Big) \cdot \left(\mathbb{E}[H_t^{RPB} + \Theta_t^{RPB} + \Pi_t^{RPB} + K \cdot \mathbb{1}(q_t^{RPB} > 0)\right) \mid \mathscr{F}_t]\Big]\Big]\Big\} \\
&\le \left(3 + \frac{(m-2)h}{mh+\theta}\right) \sum_{t=1}^{T} \Big\{\mathbb{E}\Big[\mathbb{E}\Big[\mathbb{1}\Big(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{1\Pi}\Big) \cdot \Lambda_t \mid \mathscr{F}_t\Big]\Big] \\
&\quad + \mathbb{E}\Big[\mathbb{E}\Big[\mathbb{1}\Big(t \in \mathscr{T}_{2H} \bigcup \mathscr{T}_{2M} \bigcup \mathscr{T}_{2\Pi}\Big) \cdot p_t \cdot K \mid \mathscr{F}_t\Big]\Big]\Big\} \\
&= \left[3 + \frac{(m-2)h}{mh+\theta}\right] \sum_{t=1}^{T} \mathbb{E}[Z_t^{RPB}],
\end{aligned}
$$

where the second equality holds because conditioning on $\mathscr{F}_t$, $\mathbb{1}(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{1\Pi})$ and $\mathbb{1}(t \in \mathscr{T}_{2H} \bigcup \mathscr{T}_{2M} \bigcup \mathscr{T}_{2\Pi})$ are known; the inequality is valid due to the construction of RPB; and the last equality follows from the definition of Z_t^{RPB}. **Q.E.D.**

Proof of Theorem 1. Denote the expected total cost of OPT by $\mathscr{C}(OPT)$, then we have

$$
\begin{aligned}
\mathscr{C}(OPT) &= \mathbb{E}\left[H_t^{OPT} + \Theta_t^{OPT} + \Pi_t^{OPT} + \sum_{t=1}^{T} K \cdot \mathbb{1}\left(q_t^{OPT} > 0\right)\right] \\
&\ge \sum_{t=1}^{T} \mathbb{E}\Big[\frac{mh+\theta}{2(m-1)h+\theta} \cdot \Big[\mathbb{1}\Big(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H}\Big)(H_t^{RPB} + \Theta_t^{RPB})\Big] \\
&\quad + \mathbb{1}\Big(t \in \mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi}\Big)\Pi_t^{RPB} + K \cdot [\mathbb{1}(t \in \mathscr{T}_{2M}) \cdot \mathbb{1}(q_t^{RPB} > 0)]\Big] \\
&= \sum_{t=1}^{T} \mathbb{E}\Big[\mathbb{E}\Big[\frac{mh+\theta}{2(m-1)h+\theta} \cdot \Big[\mathbb{1}\Big(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H}\Big)(H_t^{RPB} + \Theta_t^{RPB})\Big] \\
&\quad + \mathbb{1}\Big(t \in \mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi}\Big)\Pi_t^{RPB} + K \cdot [\mathbb{1}(t \in \mathscr{T}_{2M}) \cdot \mathbb{1}(q_t^{RPB} > 0)] \mid \mathscr{F}_t\Big]\Big] \\
&\ge \sum_{t=1}^{T} \mathbb{E}\Big[\frac{mh+\theta}{2(m-1)h+\theta}\mathbb{P}(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H} \mid \mathscr{F}_t) \cdot \mathbb{E}\left[H_t^{RPB} + \Theta_t^{RPB} \mid \mathscr{F}_t\right] \\
&\quad + \mathbb{P}\Big(t \in \mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi} \mid \mathscr{F}_t\Big) \cdot \mathbb{E}\left[\Pi_t^{RPB} \mid \mathscr{F}_t\right] + \mathbb{P}\big(t \in \mathscr{T}_{2M} \mid \mathscr{F}_t\big) \cdot \mathbb{E}\left[K \cdot \mathbb{1}(q_t^{RPB} > 0) \mid \mathscr{F}_t\right]\Big] \\
&= \sum_{t=1}^{T} \mathbb{E}[Z_t^{RPB}] \ge \frac{1}{3 + \frac{(m-2)h}{mh+\theta}} \mathscr{C}(RPB).
\end{aligned}
$$

The first inequality follows from Lemmas 2 and 3; and the last inequality follows from Lemma 5.

Now it remains to check the validity of the second inequality. Lemma 4 implies that

$$
\mathbb{E}\Big[\mathbb{1}\Big(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H}\Big) \cdot (H_t^{RPB} + \Theta_t^{RPB}) \mid \mathscr{F}_t\Big] \ge \mathbb{P}\Big(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H} \mid \mathscr{F}_t\Big) \cdot \mathbb{E}\Big[H_t^{RPB} + \Theta_t^{RPB} \mid \mathscr{F}_t\Big],
$$

which gives us the desired inequality for the holding and outdating cost part.

By conditioning on the information set $\mathscr{F}_t$, the indicator function $\mathbb{1}(t \in \mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi})$ becomes known. Hence we can write

$$\mathbb{E}\Big[\mathbb{1}\Big(t \in \mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi}\Big) \cdot \Pi_t^{RPB} \mid \mathscr{F}_t\Big] = \mathbb{P}\Big(t \in \mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi} \mid \mathscr{F}_t\Big) \cdot \mathbb{E}\Big[\Pi_t^{RPB} \mid \mathscr{F}_t\Big].$$

Moreover, given the information set $\mathscr{F}_t$, the setup cost $K \cdot \mathbb{1}(q_t^{RPB} > 0)$ is fully determined by the RPB policy. Thus, we have

$$\mathbb{E}\left[K \cdot [\mathbb{1}(t \in \mathscr{T}_{2M}) \cdot \mathbb{1}(q_t^{RPB} > 0)] \mid \mathscr{F}_t\right] = \mathbb{P}\big(t \in \mathscr{T}_{2M} \mid \mathscr{F}_t\big) \cdot \mathbb{E}\left[K \cdot \mathbb{1}(q_t^{RPB} > 0) \mid \mathscr{F}_t\right].$$

Combining the above arguments, we have shown that the second inequality holds as well. **Q.E.D.**

EC.3. Numerical Results

		$\mathscr{C}$(**RPB**)	$\mathscr{C}$(**NA**)	r	$\mathscr{C}$(**RPB**)	$\mathscr{C}$(**NA**)	r	$\mathscr{C}$(**RPB**)	$\mathscr{C}$(**NA**)	r
Erlang-2 Demand Distribution										
	b \ θ	**5**			**10**			**15**		
K=5	**5**	297.25	302.29	-1.67%	333.15	338.44	-1.56%	348.80	357.93	-2.55%
	10	396.60	404.00	-1.83%	461.15	468.85	-1.64%	501.66	514.70	-2.53%
	15	462.90	467.29	-0.94%	543.17	550.72	-1.37%	611.38	621.04	-1.56%
K=10	**5**	373.85	373.01	0.23%	415.99	415.74	0.06%	433.30	432.60	0.16%
	10	485.15	483.24	0.40%	540.60	548.27	-1.40%	599.16	609.58	-1.71%
	15	546.26	550.10	-0.70%	635.19	639.33	-0.65%	711.02	720.90	-1.37%
K=15	**5**	436.60	431.18	1.26%	491.79	488.57	0.66%	513.15	507.83	1.05%
	10	558.40	555.65	0.49%	617.95	615.51	0.40%	675.26	673.32	0.29%
	15	638.09	632.08	0.95%	730.94	729.16	0.24%	798.46	798.31	0.02%
Exponential Demand Distribution										
K=5	**5**	405.08	405.41	-0.08%	459.94	456.56	0.74%	497.16	493.69	0.70%
	10	564.60	568.52	-0.69%	662.38	668.79	-0.96%	745.61	751.37	-0.77%
	15	723.69	725.65	-0.27%	833.76	835.11	-0.16%	932.33	946.55	-1.50%
K=10	**5**	475.35	471.06	0.91%	538.64	532.61	1.13%	570.12	561.88	1.47%
	10	644.25	642.84	0.22%	752.16	751.73	0.06%	823.58	820.74	0.35%
	15	784.91	778.17	0.87%	917.32	918.22	-0.10%	1013.45	1028.58	-1.47%
K=15	**5**	545.31	539.14	1.14%	591.76	585.37	1.09%	647.76	635.45	1.94%
	10	718.31	710.05	1.16%	823.89	813.72	1.25%	901.32	893.41	0.89%
	15	838.34	828.72	1.16%	988.31	985.02	0.33%	1097.46	1094.77	0.25%

Table EC.1 Computational Performance of RPB and NA under IID Demands ($h = 1$, $\alpha = 0.95$)

	K	$\sum \mathbb{E}(H_t)$	$\sum \mathbb{E}(\Pi_t)$	$\sum \mathbb{E}(\Theta_t)$	Setup Costs
Erlang-2 Demand Distribution	5	20.3%	45.7%	13.1%	20.9%
	10	15.8%	43.4%	9.9%	30.9%
	15	15.9%	41.9%	9.6%	32.6%
Exponential Demand Distribution	5	16.4%	51.8%	16.0%	15.7%
	10	15.2%	45.1%	14.4%	25.3%
	15	12.9%	43.0%	12.7%	31.4%

Table EC.2 Total Cost Breakdown under IID Demands ($h=1$, $\alpha=0.95$, $b=\theta=5$)

		$\mathscr{C}$(**RPB**)	$\mathscr{C}$(**OPT**)	r	$\mathscr{C}$(**RPB**)	$\mathscr{C}$(**OPT**)	r	$\mathscr{C}$(**RPB**)	$\mathscr{C}$(**OPT**)	r
Erlang-2 Demand Distribution										
	b \ θ	**5**			**10**			**15**		
K=5	**5**	295.23	276.19	6.89%	327.36	309.49	5.77%	362.55	334.12	8.51%
	10	398.12	375.35	6.07%	469.11	439.96	6.63%	506.90	486.04	4.29%
	15	471.22	441.81	6.66%	554.29	531.72	4.24%	618.99	598.76	3.38%
K=10	**5**	350.14	327.37	6.96%	396.91	363.01	9.34%	419.10	387.96	8.03%
	10	462.57	428.58	7.93%	532.03	495.59	7.35%	576.67	541.95	6.41%
	15	533.63	495.76	7.64%	617.82	588.01	5.07%	692.38	656.61	5.45%
K=15	**5**	410.63	374.66	9.60%	449.33	412.98	8.80%	486.51	437.21	11.28%
	10	523.01	478.81	9.23%	595.11	548.33	8.53%	642.63	594.73	8.05%
	15	592.84	546.61	8.46%	682.17	641.47	6.34%	762.60	711.16	7.23%
Exponential Demand Distribution										
K=5	**5**	389.61	369.13	5.55%	438.46	416.09	5.38%	467.46	447.23	4.52%
	10	559.15	524.17	6.67%	639.06	622.86	2.60%	718.02	687.53	4.43%
	15	648.68	629.39	3.06%	801.82	770.81	4.02%	889.96	868.51	2.47%
K=10	**5**	443.38	417.68	6.15%	498.09	466.14	6.85%	525.32	498.21	5.44%
	10	631.57	575.17	9.81%	719.86	675.37	6.59%	784.01	741.35	5.75%
	15	727.03	681.08	6.75%	878.62	824.75	6.53%	964.56	924.04	4.39%
K=15	**5**	507.38	461.33	9.98%	565.08	511.92	10.38%	589.62	542.68	8.65%
	10	688.26	623.02	10.47%	783.96	724.92	8.14%	838.85	792.19	5.89%
	15	808.49	729.79	10.78%	929.15	875.81	6.09%	1027.84	976.51	5.26%

Table EC.3 Computational Performance of RPB and OPT under MMDP Demands ($h=1$, $\alpha=0.95$)