Managing Perishable Inventory Systems as Non-perishable Ones

Hailun Zhang, Jiheng Zhang, Rachel Q. Zhang
Department of Industrial Engineering and Decision Analytics,
The Hong Kong University of Science & Technology,
hzhangaq@connect.ust.hk, jiheng@ust.hk, rzhang@ust.hk

Perishable inventory problems have a long history and involve two fundamental decisions, how much to order and how much old inventory to clear before expiration, that are known to be difficult to optimize due to the curse of dimensionality. Most early work ignores the clearance decision and focuses solely on the ordering decision until recently where heuristic clearance policies have been developed. In this paper, we approach the problem from a different angle by exploring its asymptotic behavior, i.e., perishability can be ignored in many cases and hence clearance of inventory is not necessary except at the beginning of the planning horizon when a system is large enough. Inspired by such asymptotic behavior, we examine simple policies that ignore clearance under minor conditions and establish theoretical bounds for them. The bounds not only vanish asymptotically, but also indicate a system size required to guarantee any given optimality gap. Numerical studies suggest that such policies can work very well for systems with reasonable sizes and practical management of complex perishable inventory systems is not so much harder than that of non-perishable ones.

Key words: perishable inventory, newsvendor problem, asymptotic analysis

1. Introduction

Consider the classical, periodic review perishable inventory system serving random customer demand. At the beginning of each period after the age distribution of the inventory on hand is observed, the amount of new inventory needed is determined and ordered which arrives immediately. Demand is then realized and met with inventory from the oldest to the newest. Unmet demand is lost. At the end of each period, inventory that has expired is discarded at a cost and a decision is then made on the amount of inventory to be cleared, referred to as the clearance decision. In the absence of a fixed ordering cost, Li and Yu (2014) have successfully characterized the structure of the optimal clearance decision as one governed by multiple thresholds, and established the monotonicity and bounded sensitivity of the optimal ordering and clearance policy using the concept of multi-modularity. In the presence of a fixed ordering cost, successful characterization of the optimal decisions is almost impossible for problems with general lifetimes. Regardless of whether an optimal structure is known or not, it is still difficult to compute an optimal policy due to the curse of dimensionality as both decisions depend on the age distribution of the inventory on hand. The longer the lifetime, the higher the dimension of the state space. Thus, although a large number of studies have been conducted on perishable inventory systems over the past decades, most focus on structural properties and/or offer heuristic policies without theoretical guarantee.

We will approach the problem from a different angle by exploring its asymptotic behavior as the demand and hence the inventory needed to meet it become large. Consider large supermarkets and online retail stores such as SF Best, JD fresh, and Amazon Fresh, each serving a large number of customers. According to a recent report (`https://www.investors.com/news/amazon-fresh-grocery-threatens-wal-mart-kroger/`), "They see the company's (Amazonfresh) grocery volume shooting up to $23.2 billion five years from now from an estimated $8.7 billion this year (2016)". In such applications, although different retailers may have customer bases of different sizes, each customer only buys for his or her own household and is likely to be similar across retailers selling the same product. Thus, demand in such retailers can be approximated as a compound Poisson process, i.e., the number of arrivals is Poisson and each arrival requires a random amount of inventory independent of the arrival rate. In the asymptotic regime, the arrival rate becomes large while the magnitude of the request from each arrival remains the same, i.e., the individual request does not scale up. Thus, we refer to the arrival rate of customers as the size of the system. We note that all of our results hold as long as the arrival process is a renewal process.

In the asymptotic limit, the system behaves as if demand is deterministic and hence inventory outdating can be completely avoided if it starts with a reasonably amount of inventory. Thus, clearance of inventory, if any, will only occur at the beginning of the planning horizon. This insight is very encouraging because, for large enough systems, the clearance decision guided by complex state dependent thresholds may be unnecessary if the initial inventory is low, and hence managing a perishable inventory system may be as easy as managing a non-perishable one. Inspired by this insight, we consider two very simple, asymptotically optimal heuristic policies to guide the ordering and clearance of initial inventory.

However, such simple heuristics inspired by the asymptotic behavior are meaningful only if they work well for systems of a realistic size, not just in the limit. To this end, we establish a theoretical bound on the performance gap, measuring how far from the true optimal a heuristic can achieve, for each of the two heuristic policies. These bounds vanish asymptotically implying asymptotic optimality of the heuristics and provide easy-to-calculate lower bounds of the system size that guarantee any given desired performance. Numerical experiments show that the lower bounds of the system sizes to achieve 2%, 3% and 5% performance gap are of reasonable sizes in various scenarios (see Tables 1 and 3). Since experiments on systems with sizes substantially smaller than the lower bounds of the required sizes also indicate performance within a few percentage from the optimal (see Tables 2 and 4), we also develop an efficient algorithm to compute the true optimal performance so that we can directly benchmark and more accurately evaluate the performance of the heuristic policies. Numerical comparisons against the optimal performance reveal that our heuristic policies work well even for small-to-medium-sized systems.

The paper is organized as follows. After a literature review in Section 2, we formulate the problem and define the asymptotic regime in Section 3. In Section 4, we show that both a fluid-based and newsvendor type of policy are asymptotically optimal to the original problem with the latter outperforms the former. In Section 5, we study a few extensions by examining the performance of the heuristic policies under backlogging, general demand distributions, the LIFO issuing rule (inventory is consumed from the newest to the oldest), general demand distributions, and when the distribution of the batch size is unknown and there is a fixed ordering cost. We present our efficient parallel algorithm to compute the optimal policy and objective in Section 6 and conclude the paper in Section 7. All the proofs can be found in the Appendix.

2. Literature review

2.1. Perishable inventory systems

The body of work on the management of periodic review perishable inventory systems is voluminous. We refer readers to Nahmias (1982) for a review of the early work, and Nahmias (2011) and Karaesmen et al. (2011) for reviews of more recent papers. Nahmias and Pierskalla (1973) characterize the structure of the optimal ordering policy when the lifetime of the product is two periods. For general lifetimes, Fries (1975) and Nahmias (1975) prove the monotonicity of the order quantity in on-hand inventory of different ages and the sensitivity of the order quantity to the age of on-hand inventory for two systems that differ in how the outdating costs are accounted for and are shown to be equivalent by Nahmias (1977b).

When clearance of inventory either at a cost or with a salvage value before the expiration date is allowed, Chen et al. (2014) study a joint pricing and inventory control problem for a perishable product. Employing the concept of $L^\natural$-convexity, they shed new light on perishable inventory management by obtaining comprehensive structural properties of the optimal policy in addition to providing a significantly simpler proof of a classical structural result. Li and Yu (2014) characterize the structure of the optimal control policy without pricing by using the concept of multi-modularity. Li et al. (2016) further study the structure of the optimal solution when demand is fulfilled with inventory from the newest to the oldest.

The state-dependent nature of the optimal policy for perishable inventory systems makes it extremely difficult to find such a policy, especially for products with long lifetimes. Therefore, much effort has been devoted to developing efficient heuristic ordering policies when the clearance of inventory before the expiration date is not allowed. Nahmias (1976, 1977a) propose myopic ordering policies by pretending that the lifetime of the product is shorter than it actually is to reduce the dimension of the state for the problems described by Fries (1975) and Nahmias (1975). Nandakumar and Morton (1993) propose a policy based on some analysis on the upper and lower bounds of the value function. Their policy is shown to perform near optimal in extensive numerical experiments when the lifetime is short. However, there is no theoretical analysis that can guarantee the performance, while the policy are "expected to work well when the demand distribution exhibits low variance". Limiting the problem to a single critical number policy, Cohen (1976) provides a closed-form expression of the optimal order-up-to level and Cooper (2001) derives bounds on the stationary distribution of outdated inventory. Chao et al. (2015) extend the framework of two-approximation analysis, developed by Levi et al. (2007), to prove a theoretical bound of the worst-case performance for periodic review perishable inventory systems with general demand distribution. Zhang et al. (2016) further extends the two-approximation method to the case with positive setup cost.

A few publications have tried to develop heuristics for systems that allow clearing inventory before the expiration date. Xue et al. (2012) derive piece-wise linear bounds on the value function, which leads to an order-up-to and deplete-down-to policy. Chen et al. (2014) develop a myopic heuristic policy for a joint inventory control and pricing problem. Our heuristic policies, which ignores the clearance decision except at the beginning of the planning horizon and whose order up to level in each period is either the mean demand or the newsvendor solution of the non-perishable system, work well for systems with moderate or high demand.

There is another stream of research on continuous review perishable systems where perishability is caused by disasters and hence the lifetime of the product considered is random, e.g., Gürler and Özkaya (2008), Baron et al. (2010) and David and Mehrez (1995). With random lifetimes, the methodology used to analyze the system under the (s, S) policy is very different. Most of the studies require the lifetime to be exponentially distributed for its memoryless property and some require that the models can be converted into single period ones.

Among all the work described above, only Chao et al. (2015) and Zhang et al. (2016) provide analytical bounds for their heuristics. Since the aim was to bound the performance regardless of the system size, these bounds can only guarantee the performance of their policies to be within 200% or 300% from the optimal and are of limited practical use because the actual performance gaps are likely to be much smaller as confirmed by extensive numerical studies. We provide much tighter analytical bounds for our heuristics that link their performance with the system size explicitly and are of practical significance in guiding the implementation of the heuristics.

2.2. Asymptotic analysis

Asymptotic analysis aims to explore the asymptotic properties of a complex system and has wide applications across sciences. Since the behavior of a system is usually more predictable in an asymptotic regime, asymptotic properties enhance the understanding of and help develop heuristic policies that work well for the original system. In the literature of Operations Research and Operations Management, asymptotic analysis, mostly as the customer arrival rate becomes large, is often used to understand complex stochastic systems. To name a few representative papers, Allon and Van Mieghem (2010) show that a Tailored Base-Surge policy, under which the same order is placed with the slower supplier in each period while orders are placed with the faster supplier following a simple order-up-to rule, is asymptotical optimal for a high-volume dual sourcing production system. Armony et al. (2009) use a fluid model to approximate a queuing system with delayed announcement. Bassamboo and Randhawa (2010) and Bassamboo et al. (2012) obtain some nice insights into the capacity decision of parallel queueing systems by analyzing some fluid models of the systems. More studies using fluid models to obtain practical insights include Bassamboo and Randhawa (2016), Perry and Whitt (2011) and Wu et al. (2017).

There has been a few asymptotic analyses of inventory systems. Consider the standard periodic review, lost sales inventory problem with a positive ordering lead time and zero fixed ordering cost. The state of the system includes the on-hand inventory as well as all the outstanding orders, and hence is multi-dimensional. Zipkin (2008) shows that the optimal ordering decision is monotone in the transformed state variables and provides bounds on the optimal ordering decision. Since the optimal ordering decision is state dependent and complex, optimization of this dynamic program problem suffers from the curse of dimensionality. However, Goldberg et al. (2016) show that a simple constant-order policy is asymptotically optimal when lead times become large and derive a lower bound on the lead time to achieve any desired performance gap. Xin and Goldberg (2016) further prove that, for the infinite-horizon variant of the same lost sales problem, the optimality gap of the same constant-order policy converges exponentially rapidly to zero. Xin and Goldberg (2017) also consider a dual sourcing inventory problem where one supplier is quicker but charges more, while the other is slower but charges less. They show that the Tailored Base-Surge heuristic policy is asymptotically optimal as the difference in the lead times between the two sources grows large, which is exactly the setting where the curse of dimensionality renders the problem intractable. In our paper, we will develop bounds for our heuristics that are asymptotically optimal as the demand becomes large and provide a lower bound of the size of the demand above which the performance gap of the heuristics are below any given level.

3. Problem description and asymptotic regime

Consider a periodic-review, single-product inventory system over a finite horizon of T periods where the product has a fixed lifetime of $L > 1$ periods. Thus, at the beginning of each period, there may be inventory with a remaining lifetime ℓ for $0 < \ell \leq L-1$.

The sequence of events and decisions

At the beginning of period t, an order of q_t units of new inventory is place at a unit cost of c and is received instantaneously. A random demand D_t then occurs and is met immediately as much as possible with inventory from the oldest to the newest at a price p. Unmet demand is lost. At the end of the period, expired inventory is discarded at a unit cost of θ. If there is inventory left, the firm decides on the amount z_t to be cleared, from the oldest to the newest, at a salvage value of s. A negative salvage value implies that there is a cost to clear a unit of inventory even before it expires. Any inventory carried over to the next period incurs a holding cost of h per unit. At the end of the planning horizon, all inventory will be salvaged at a value of c, which is commonly used in the literature (for example Chao et al. (2015)). Thus, there are two decisions to be made in each period, the order quantity at the beginning and the amount of inventory to be cleared at the end.

When $s = -\infty$, our problem reduces to the classical perishable inventory problem where clearance is not allowed. Since clearance of inventory before expiration is a common practice, e.g., retailers regularly sell products approaching their expiration dates at a discount to mitigate loss in practice, it should be considered as in recent work by Chen et al. (2014) and Li et al. (2016).

The state of the system and its evolution

Since there is no information update between the clearance decision at the end of a period and the ordering decision in the next period, we define $I_t(\ell)$, where $0 \le t \le T$ and $1 \le \ell \le L-1$, as the amount of inventory with remaining life up to ℓ at the end of period t, after the expired inventory is discarded but before the clearance decision is made. Then, $I_t(\ell)$ increases in ℓ, $I_t(L-1)$ is the total inventory, and the inventory profile $\mathbf{I}_t = \{I_t(\ell), 0 < \ell \le L-1\}$ describes the state of the system at the end of period t. Let $\mathcal{S} = \{(x_1, ..., x_{L-1}) \in \mathbb{R}^{L-1} : 0 \le x_1 \le ... \le x_{L-1}\}$ denote the set of all possible states.

For a given $\mathbf{I}_t$, the total available inventory with remaining life up to $\ell + 1$ after the clearance and ordering decisions are made is $[I_t(\ell+1) - z_t]^+ + q_{t+1}\mathbf{1}_{\{\ell=L-1\}}$ as clearance starts with the oldest. For a given demand D_{t+1}, the amount of inventory that will expire and hence be discarded at the end of period $t+1$ is $[I_t(1) - z_t - D_{t+1}]^+$. Thus, for $0 < \ell \le L-1$,

$$I_{t+1}(\ell) = \left[[I_t(\ell+1) - z_t]^+ + q_{t+1}\mathbf{1}_{\{\ell=L-1\}} - D_{t+1} - [I_t(1) - z_t - D_{t+1}]^+\right]^+. \qquad (1)$$

For notational purposes, let Γ be the mapping from $(\mathbf{I}_t, z_t, q_{t+1}, D_{t+1})$ to the new profile $\mathbf{I}_{t+1} = \Gamma(\mathbf{I}_t, z_t, q_{t+1}, D_{t+1})$.

Problem formulation

Let $V_t(\mathbf{I})$ be the optimal expected discounted profit at the end of period t for a given inventory profile $\mathbf{I} \in \mathcal{S}$, $\Omega(\mathbf{I}) = \{(z,q) : 0 \le z \le I(L-1), q \ge 0\}$ be the set of all feasible decisions for a given $\mathbf{I}$, and α be the discount factor. Then,

$$V_t(\mathbf{I}) = \max_{(z,q)\in\Omega(\mathbf{I})} \{\mathbb{E}[G(\mathbf{I},z,q,D) + \alpha V_{t+1}(\Gamma(\mathbf{I},z,q,D))]\}$$

for $0 \le t \le T-1$ where

$$G(\mathbf{I},z,q,D) = sz - h\,[I(L-1)-z]^+ - \alpha c q - \alpha\theta\,[I(1)-D-z]^+ + \alpha p \min\{D, I(L-1)-z+q\}$$

is the single-period profit function (revenue from clearance and sales less the purchasing and holding costs) for a given clearance and ordering decision $(z,q) \in \Omega(\mathbf{I})$ and demand D, and

$$V_T(\mathbf{I}) = \max_{0\le z\le I(L-1)} \left\{sz - h\,[I(L-1)-z]^+ + \alpha c\,[I(L-1)-z]^+\right\} = (\alpha c - h)I(L-1)$$

as no order will be placed after period T and any inventory remaining after period T will be salvaged at c.

Since one can recover the salvage value s as well as save the holding cost h by clearing one unit of inventory, it is optimal to simply deplete all the inventory remaining at the end of each period if $s + h \ge \alpha c$ where α is the discount factor. Thus, we assume $s + h - \alpha c < 0$ throughout the paper, which include the case where $s = -\infty$ where clearance will never occur. Lemma 1 allows us to take derivatives in the proofs when needed where $v_t(\mathbf{I},z,q,D) = G(\mathbf{I},z,q,D) + \alpha V_{t+1}(\Gamma(\mathbf{I},z,q,D))$ is the expected discounted profit at the end of period t for a given clearance and ordering decision $(z,q) \in \Omega(\mathbf{I})$, demand D and inventory profile $\mathbf{I}$.

Lemma 1. *$V_t(\mathbf{I})$ and $v_t(\mathbf{I},z,q,D)$ are continuous and differentiable component-wise except at finitely many points.*

3.1. Compound Poisson demand and the asymptotic regime

Although our analysis and results hold as long as the arrival process is a renewal process, we will assume that customers arrive according to a Poisson process for the ease of presentation. It is well understood that, in a given time interval or period, the number of customer arrivals is approximately Poisson with a rate n which measures the size of the customer base. With a general distribution on the batch size of each customer demand, a compound Poisson is a reasonably general distribution to represent total customer demand in a period. While different retailers may have customer bases of different sizes, i.e., different n, the requirements from each customer are likely to be similar. Thus, in this paper, we will focus on inventory systems with different n but independently and identically distributed batch sizes as summarized in the next assumption. Throughout the paper, we will use f_D and F_D to represent the probability density function and cumulative distribution function of a random variable D.

Assumption 1. *Consider a sequence of systems indexed by $n = 1, 2, \cdots$. In the nth system, the number of arrivals in each period is Poisson with rate n, each requiring a batch with mean λ and second moment σ^2.*

Under assumption 1, demand in each period is compound Poisson with mean $n\lambda$ and variance $n\sigma^2$ in the nth system. Throughout the paper, the superscript n denotes the nth system. For instance, demand in period t in the nth system is denoted by D_t^n. Such a demand form not only adequately represents real demand, it has several nice properties. As the system size n increases, the variability in demand D_t^n decreases, i.e., the coefficient of variation $\frac{\sigma}{\sqrt{n}\lambda}$ decreases. As $n \to \infty$, the variability vanishes and the distribution of the scaled demand

$$\tilde{D}_t^n := \frac{D_t^n - n\lambda}{\sqrt{n}\sigma} \tag{2}$$

converges to the standard normal $\Phi(\cdot)$. That is, once the system becomes large enough, demand becomes more predictable which eases the management of a perishable inventory system.

Below we present a nice property on the probability density function of the scaled demand that will be used later. For example, if the batch size required by each customer is constant, the probability mass function of the scaled demand is uniformly bounded by 1/2.

Lemma 2. *The probability density function of the scaled demand $\tilde{D}_t^n$ is uniformly bounded by a constant.*

4. Performance guarantee of two heuristic policies

In this section, we consider two very simple heuristic policies that ignore clearance except at the beginning, a fluid based and a newsvendor type, and establish performance bound as a function of the system size for each of them in Sections 4.1 and 4.2, respectively. We perform numerical experiments to demonstrate their performance against the bounds as well as the optimal in Section 4.3. To evaluate the performance of a policy, we define $\frac{V_0^n(\mathbf{I})-V_0^{\pi,n}(\mathbf{I})}{V_0^n(\mathbf{I})}$ as the (relative) performance gap for any policy π and need to obtain a lower bound of $V_0^n(\mathbf{I})$. By the monotonicity of the optimal value function, $V_0^n(\mathbf{I})$ is bounded from below for any given initial inventory as

$$\begin{aligned} V_0^n(\mathbf{I}) \geq V_0^n(\mathbf{0}) \geq & \sum_{t=0}^{T} \alpha^t \left[-(\alpha p - \alpha c + h)\mathbb{E}\left[D_1^n - n\lambda\right]^+ + \alpha(p-c)\lambda\right] \\ & - \frac{\alpha^L - \alpha^{T+1}}{1-\alpha}(\alpha c + \alpha\theta - h)\mathbb{E}\left[n\lambda - D_{[L]}^n\right]^+ \\ = & \frac{\alpha - \alpha^{T+2}}{1-\alpha}(p-c)n\lambda - \frac{1-\alpha^{T+1}}{1-\alpha}(\alpha p - \alpha c + h)\mathbb{E}\left[D_1^n - n\lambda\right]^+ \\ & - \frac{\alpha^L - \alpha^{T+1}}{1-\alpha}(\alpha c + \alpha\theta - h)\mathbb{E}\left[n\lambda - D_{[L]}^n\right]^+ \end{aligned} \tag{3}$$

where $D_{[t]} := D_1 + D_2 + \cdots + D_t$.

Throughout the paper, double superscripts "π, n" represent values and functions associated with the nth system under policy π. For instance, $V_t^{\pi,n}(\mathbf{I}_t^n)$ represents the expected profit of the nth system at time t for a given heuristic policy. It is notable that although the subsequent analysis is focused on $T < \infty$, all the results can be readily extended to infinite-horizon settings ($T = \infty$) as long as $\alpha < 1$.

4.1. A fluid based heuristic

By (2), the scaled demand $\frac{D_t^n}{n} \to \lambda$, a constant, as $n \to \infty$. Thus, we first study a system where $D_t = D$, a constant, referred to as the fluid model. With deterministic demand, one is able to match supply with demand by ordering up to D in each period and clearance of inventory, if any, will only occur at the beginning of the planning horizon. The next proposition provides a closed-form solution to the optimal clearance decision for the fluid model.

Proposition 1. *Under the fluid model, it is optimal to order up to D in each period and only clear*

$$z_0^* = \max\left\{I_0(L-1) - \ell_0 D, \max_{1\leq \ell \leq L-1}\left[I_0(\ell) - \ell D\right]^+\right\} \tag{4}$$

at the beginning of the planning horizon where

$$\ell_0 = \begin{cases} L-1, & \text{if } s + \frac{1-\alpha^{L-1}}{1-\alpha}h < \alpha^{L-1}c, \\ \min_{1\leq\ell\leq L-1}\left\{\ell : s + \frac{1-\alpha^\ell}{1-\alpha}h \geq \alpha^\ell c\right\}, & \text{otherwise.} \end{cases} \tag{5}$$

Here we provide some intuitive explanations for the optimal solution, and delay the formal proof to the appendix. If $I_0(\ell) > \ell D$ for some $1 \le \ell \le L-1$, then $\max_{1\le\ell\le L-1}[I_0(\ell) - \ell D]^+$ amount of inventory will expire and be discarded eventually if kept in the system. Thus, it is optimal to clear them as soon as possible. Of the remaining $I_0(L-1) - \max_{1\le\ell\le L-1}[I_0(\ell) - \ell D]^+$, it may be more profitable to clear some now, which brings in s and lowers the inventory cost, and order later which happens if $s + \frac{1-\alpha^\ell}{1-\alpha}h \ge \alpha^\ell c$ for some $1 \le \ell \le L-1$. Thus, one should hold inventory to cover at most ℓ_0 periods of demand $\ell_0 D$. If $s + \frac{1-\alpha^{L-1}}{1-\alpha}h < \alpha^{L-1}c$, $\ell_0 = L-1$, in which case, $I_0(L-1) - \max_{1\le\ell\le L-1}[I_0(\ell) - \ell D]^+ < \ell_0 D$ and one only clears those that will expire. Otherwise, clear down to $\ell_0 D$. Thus, ℓ_0 is the maximum number of periods for which a unit of inventory should be kept in the system.

Inspired by the optimal policy of the fluid model, we propose a simple fluid policy f for the nth system which mimics that of the fluid model: One orders up to $n\lambda$, the mean demand, in each period and clearance if any only occurs at the beginning of the planning horizon with
$z_0^{f,n} = \max\left\{ I_0^n(L-1) - n\lambda\ell_0, \max_{1\le\ell\le L-1}[I_0^n(\ell) - n\lambda\ell]^+ \right\}$, where ℓ_0 is given in (5). Propositions 2 and 3 provide bounds of performance gap of the fluid policies when $\mathbf{I}_0^n = \mathbf{0}$ and for a general initial state, respectively. These bounds will be used to determine the minimum system size that guarantees any given desired performance for the policy f. Note that the above bounds are independent of the salvage value s since the clearance decision can be ignored for a relatively large system except at the very beginning.

For convenience add a bar to the symbols under the fluid model, e.g., $\bar{\mathbf{I}}_t$ represents the inventory profile at the end of period t under the fluid model. Besides, add a bar to the symbols under the nth system to represent the scaled quantity, e.g., $\bar{\mathbf{I}}_t^n = \mathbf{I}_t^n/n$.

Proposition 2. *$V_0^n(\mathbf{0}) - V_0^{f,n}(\mathbf{0})$ is in the order of $O(\sqrt{n})$ and*

$$\begin{aligned} V_0^n(\mathbf{0}) - V_0^{f,n}(\mathbf{0}) \le & \sum_{t=1}^{T} \alpha^t \left\{ -cy^n + p\mathbb{E}\min(D_t^n, y^n) + (\alpha c - h)\mathbb{E}\left[y^n - D_t^n\right]^+ \right\} + (\alpha c - h + \theta)\mathbb{E}\left[\lambda n - D_{[L]}^n\right]^+ \\ & - \sum_{t=1}^{T} \alpha^t \left[-cn\lambda + p\mathbb{E}\min(D_t^n, n\lambda) + (\alpha c - h)\mathbb{E}\left[n\lambda - D_t^n\right]^+ \right]. \end{aligned} \tag{6}$$

Proof. Consider a nonperishable inventory system for which the event time line is the same as that for the perishable system except that the product is non-perishable. It is well-known in the literature, for example in Scarf (1959), that the optimal policy is to order up to $y^n = F_{D_t^n}^{-1}\left(\frac{p-c}{p-\alpha c+h}\right)$ in each period, where $F^{-1}(\cdot)$ stands for the reverse function of $F(\cdot)$. Observe that the optimal value function of the nonperishable system serves as an upper bound of the optimal value function of the nth system. Thus,

$$V_0^n(\mathbf{0}) \le \sum_{t=1}^{T} \left\{ \alpha^t - cy^n + p\mathbb{E}\min(D_t^n, y^n) + (\alpha c - h)\mathbb{E}\left[y^n - D_t^n\right]^+ \right\}.$$

Let $O_t^n = \left[I_{t-1}^n(1) - z_{t-1}^n - D_t^n\right]^+$ be the amount of inventory that is expired in period t. Then

$$\begin{aligned} V_0^{f,n}(\mathbf{0}) &= \sum_{t=1}^{T} \alpha^t \left[-c(\mathbb{E}D_{t-1}^n + \mathbb{E}O_{t-1}^n) + p\mathbb{E}\min(D_t^n, n\lambda) - h\mathbb{E}\left[n\lambda - D_t^n - O_t^n\right]^+ - \theta\mathbb{E}O_t^n\right] \\ &\ge \sum_{t=1}^{T} \alpha^t \left[-cn\lambda + p\mathbb{E}\min(D_t^n, n\lambda) + (\alpha c - h)\mathbb{E}\left[n\lambda - D_t^n\right]^+\right] - (\alpha c - h + \theta)\mathbb{E}O_t^n. \end{aligned}$$

Combining the above two inequalities leads to (6). Next we show (6) is of order $O(\sqrt{n})$. Note that, there exists a value ξ which lies between y^n and $n\lambda$ such that

$$V_0^n(\mathbf{0}) - V_0^{f,n}(\mathbf{0}) \le \sum_{t=1}^{T} \alpha^t \left[\frac{1}{2}(p - \alpha c + h) f_{D_t^n}(\xi)(y^n - n\lambda)^2 + (\alpha c - h + \theta)\mathbb{E}O_t^n\right].$$

Since

$$\frac{1}{2}(p - \alpha c + h) f_{D_t^n}(\xi)(y^n - n\lambda)^2 \le \frac{C_1}{2\sqrt{n}\sigma}(p - \alpha c + h)(y^n - n\lambda)^2$$

for some constant C_1 by Assumption 2 and $\mathbb{E}O_t^n \le \mathbb{E}\left[n\lambda - D_{[L]}^n\right]^+$ (see the similar result $O_t(x, y, d) \le B_t(y, d)$ in Section 6.2 of Chen et al. (2014)), we obtain

$$\begin{aligned} V_0^n(\mathbf{0}) - V_0^{f,n}(\mathbf{0}) \le & \sum_{t=1}^{T} \alpha^t \left\{-cy^n + p\mathbb{E}\min(D_t^n, y^n) + (\alpha c - h)\mathbb{E}\left[y^n - D_t^n\right]^+\right\} \\ & - \sum_{t=1}^{T} \alpha^t \left[-cn\lambda + p\mathbb{E}\min(D_t^n, n\lambda) + (\alpha c - h)\mathbb{E}\left[n\lambda - D_t^n\right]^+\right] \\ & + (\alpha c - h + \theta)\mathbb{E}\left[\lambda n - D_{[L]}^n\right]^+ \\ \le & \frac{\alpha - \alpha^{T+1}}{1 - \alpha} \left[\frac{C_1}{2\sqrt{n}\sigma}(p - \alpha c + h)(y^n - n\lambda)^2 + (\alpha c - h + \theta)\mathbb{E}\left[\lambda n - D_{[L]}^n\right]^+\right., \end{aligned}$$

which is of order $\sqrt{n}$ as $(y^n - n\lambda)/\sqrt{n} = F_{\tilde{D}_t^n}^{-1}\left(\frac{p-c}{p-\alpha c+h}\right) \to \Phi^{-1}\left(\frac{p-c}{p-\alpha c+h}\right)$.

By modifying the above proof slightly, we can show that $V_0^n(\mathbf{0}) - V_0^{f,n}(\mathbf{0}) \ge \Omega(\sqrt{n})$, where the notation $\Omega(g(n))$ means there exists a positive constant C which is greater than $Cg(n)$ when n is large enough. If we adopt the fluid policy in all periods except in period

T in which we order $q_T = \left[F^{-1}_{D^n_t}\left(\frac{p-c}{p-s}\right) - I^n_T(L-1)\right]^+$, then the value function is below the optimal. When n is large enough,

$$\begin{aligned}
V^n_0(\mathbf{0}) - V^{f,n}_0(\mathbf{0}) &\geq \alpha^T\{\mathbb{E}v^n_T(\mathbf{I}^n_T, 0, q_T, D^n_T) - \mathbb{E}v^n_T(\mathbf{I}^n_T, 0, n\lambda, D^n_T)\} \\
&\geq -\alpha^{T+1}c\mathbb{E}\left[I^n_T(L-1) - F^{-1}_{D^n_t}\left(\frac{p-c}{p-s}\right)\right]^+ + \alpha^T[\mathbb{E}v^n_T(\mathbf{0}, 0, q_T, D^n_T) - \mathbb{E}v^n_T(\mathbf{0}, 0, n\lambda, D^n_T)] \\
&\geq \alpha^T[\mathbb{E}v^n_T(\mathbf{0}, 0, q_T, D^n_T) - \mathbb{E}v^n_T(\mathbf{0}, 0, n\lambda, D^n_T)] \\
&= -\frac{\alpha^T}{2}\left(\frac{\partial^2 \mathbb{E}v^n_T(\mathbf{0}, 0, q, D^n_T)}{\partial q^2}\right)|_{q=\xi}\left[F^{-1}_{D^n_t}\left(\frac{p-c}{p-s}\right) - n\lambda\right]^2 \\
&= \frac{1}{2}\alpha^{T+1}(p-s)f_{D^n_t}(\xi)\left[F^{-1}_{D^n_t}\left(\frac{p-c}{p-s}\right) - n\lambda)\right]^2
\end{aligned}$$

where ξ is a value between $F^{-1}_{D^n_t}\left(\frac{p-c}{p-s}\right)$ and $n\lambda$. The third inequality follows as $\mathbb{E}\left[I^n_T(L-1) - F^{-1}_{D^n_t}\left(\frac{p-c}{p-s}\right)\right]^+ \leq \mathbb{E}\left[n\bar{D} - F^{-1}_{D^n_t}\left(\frac{p-c}{p-s}\right) - D^n_T\right]^+ = o(1)$. According to Assumption 1, $\frac{D^n_t - n\lambda}{\sqrt{n}} \to N(0, \sigma^2)$ as $n \to \infty$. Then, it is easy to see that $f_{D^n_t}(\xi) \sim \Omega(1/\sqrt{n})$ and $\left[F^{-1}_{D^n_t}\left(\frac{p-c}{p-s}\right) - n\lambda\right]^2 \sim \Omega(n)$. □

The following lemma is crucial for deriving the performance bound in Proposition 3.

Lemma 3. *For any given initial inventory profile $\mathbf{I}^n_0$, the optimal profit of the nth system $V^n_0(\mathbf{I}^n_0) \leq n\bar{V}_0(\bar{\mathbf{I}}_0)$ where $\bar{V}_0(\cdot)$ is the optimal profit of the fluid model with $D_t = \lambda$.*

Proof. Recall that $v^n_t(\mathbf{I}, z, q, D) = G(\mathbf{I}, z, q, D) + \alpha V^n_{t+1}(\Gamma(\mathbf{I}, z, q, D))$. Thus, if $v^n_t(\mathbf{I}, z, q, D)$ is concave in D, which we will show is true, then

$$V^n_t(\mathbf{I}) = \max_{(z,q)\in\Omega(\mathbf{I})} \mathbb{E}\{v^n_t(\mathbf{I}, z, q, D^n_{t+1})\} \leq \max_{(z,q)\in\Omega(\mathbf{I})}\{v^n_t(\mathbf{I}, z, q, \mathbb{E}D^n_{t+1})\}$$

for all t by Jensen's inequality. Thus, $V^n_0(\mathbf{I}^n_0) \leq n\bar{V}_0(\bar{\mathbf{I}}_0)$ by definition.

It suffices to show that $\frac{\partial v^n_t(\mathbf{I}, z, q, D)}{\partial D}$ is decreasing in D. First note that when $D < I(1) - z$ or $D > I(L-1) - z + q$, $\Gamma(\mathbf{I}, z, q, D)$ remains constant, thus $\frac{\partial V^n_{t+1}(\Gamma(\mathbf{I}, z, q, D))}{\partial D} = 0$. Thus,

$$\begin{aligned}
\frac{\partial v^n_t(\mathbf{I}, z, q, D)}{\partial D} &= \frac{\partial G(\mathbf{I}, z, q, D)}{\partial D} + \alpha\frac{\partial V^n_{t+1}(\Gamma(\mathbf{I}, z, q, D))}{\partial D} \\
&= \begin{cases} \alpha p + \alpha\theta, & \text{if } D < I(1) - z, \\ \alpha p + \alpha\frac{\partial V^n_{t+1}(\Gamma(\mathbf{I}, z, q, D))}{\partial D}, & \text{if } I(1) - z < D < I(L-1) - z + q, \\ 0 & \text{if } D > I(L-1) - z + q. \end{cases}
\end{aligned} \tag{7}$$

Consider the case of $I(1) - z < D < I(L-1) - z + q$. Note that, in this case, an increase (decrease) in D that is small enough results in a decrease (increase) of the same amount at the end of the period. By Chen et al. (2014), $V^n_{t+1}(\cdot)$ is $L^\natural$-concave (here we consider the profit maximization problem, while they consider the cost minimization), from which we know that $\frac{\partial V_{t+1}^n(\Gamma(\mathbf{I},z,q,D))}{\partial D}$ is decreasing in D. To prove that $\frac{\partial v_t^n(\mathbf{I},z,q,D)}{\partial D}$ is decreasing in D, in view of (7), it remains for us to show that

$$0 \leq \alpha p + \frac{\partial V_{t+1}^n(\Gamma(\mathbf{I},z,q,D))}{\partial D} \leq \alpha p + \alpha\theta. \tag{8}$$

The two profiles $\Gamma(\mathbf{I},z,q,D)$ and $\Gamma(\mathbf{I},z,q,D+\delta)$, where δ is a small enough positive constant, differ by a total of δ units of inventory. By simple difference analysis, we can show that $s \leq \frac{\partial V_{t+1}^n(\mathbf{I})}{\partial I(\ell)} \leq \alpha c - h$ for any $0 < \ell \leq L-1$. Thus, $-(\alpha c - h)\delta \leq V_{t+1}^n(\Gamma(\mathbf{I},z,q,D+\delta)) - V_{t+1}^n(\Gamma(\mathbf{I},z,q,D)) \leq 0$. It follows that $-\alpha(\alpha c - h) < \alpha \frac{\partial V_{t+1}^n(\Gamma(\mathbf{I},z,q,D))}{\partial D} < 0$, which leads to (8) as $p > c$. $\square$

PROPOSITION 3. *For any given initial inventory profile* $\mathbf{I}_0^n \in \mathcal{S}$,

$$V_0^n(\mathbf{I}_0^n) - V_0^{f,n}(\mathbf{I}_0^n) \leq K_1 \mathbb{E}\left[D_1^n - \lambda n\right]^+ + K_2 \mathbb{E}\left[\lambda n - D_{[L]}^n\right]^+ + H(n) = O(\sqrt{n})$$

where

$$\begin{aligned}
K_1 &= (L-1)c + \sum_{t=1}^{L-1}(t-1)\alpha^t h + \frac{\alpha - \alpha^{L+1}}{1-\alpha}p + \frac{\alpha^L - \alpha^{T+1}}{1-\alpha}(\alpha p - \alpha c + h), \\
K_2 &= \frac{\alpha^L - \alpha^{T+1}}{1-\alpha}(\alpha c + \alpha\theta - h), \\
H(n) &= (c + \alpha\theta)\mathbb{E}\max\left\{0, \lambda n - D_1^n, \cdots, (L-1)\lambda n - D_{[L-1]}^n\right\}.
\end{aligned}$$

Thus, we have a theoretical bound on the performance gap $\frac{V_0^n(\mathbf{I}) - V_0^{f,n}(\mathbf{I})}{V_0^n(\mathbf{I})}$ which is in the order of $O\left(\frac{1}{\sqrt{n}}\right)$ and vanishes asymptotically. Although being crude, the fluid analysis has both theoretical and practical significance. Theoretically, the above two propositions establish not only asymptotic optimality of policy f but also the speed at which its performance gap from the optimal vanishes as n becomes large. This is new in the literature of perishable inventory. A similar insight has been observed without a proof by Nandakumar and Morton (1993) that their heuristics are expected to work well when "the lifetimes are long or the demand distribution exhibits low variance".

Practically, it reveals the following important insights. (1) A family of simple and implementable policies are guaranteed to perform close enough to optimality. (2) As mentioned in the literature review, Chao et al. (2015) and Zhang et al. (2016) are the only papers in the literature on perishable inventory that provide theoretically guaranteed performance of 200% or 300% from the optimal for their proposed heuristics. Our bounds for the fluid heuristic are the first of its kind in that they are much tighter and can be arbitrary close to the optimal as the system size increases. Thus, they can be used to determine a system size that guarantees, say less than 5% performance gap. (3) Being simple is an important criterion for a policy to be implemented in practice. The fluid policy as well as the newsvendor policy we will examine in the next section, are both simple and easy to implement. (4) Our analysis delivers the message that managing perishable inventory systems is much simpler than it is understood for relatively large systems.

4.2. A newsvendor type of policy

While being simple, the fluid policy is rather crude and ignores demand uncertainty and some cost parameters. It compensates for the need to clear inventory throughout the planning horizon by a lower order-up-to level and leads to higher lost sales in most cases. In this section, we consider an enhanced policy, referred to as policy H, under which one orders up to $y^n = F_{D^n}^{-1}\left(\frac{p-c}{p-\alpha c+h}\right)$, the optimal order up to level for the corresponding non-perishable inventory system, and clears $z_0^{H,n} = \left[I_0^n(1) - F_{D^n}^{-1}\left(\frac{\alpha c-s-h}{\alpha c+\theta-h}\right)\right]^+$ amount of inventory at the beginning which is the solution of

$$\min_{0\le z\le I_0^n(L-1)} \mathbb{E}\left\{G(\mathbf{I}_0^n, z, y^n - I_0^n(L-1)+z, D_{t+1}^n) + \alpha(\alpha c - h)\left[y^n - D_1^n - [I_0^n(1) - D_1^n - z]^+\right]^+\right\},$$

the optimal cost if the lifetime of the inventory is two periods. Proposition 4 provides bounds for the absolute performance gap under the policy H for zero initial inventory and $I_0^n(L-1)\le y^n$, respectively. Note that both bounds are independent of the salvage value s. Combined with (3), we have the following bounds on the performance gap of the policy H, which vanish at an exponential speed.

PROPOSITION 4. *For any $\mathbf{I}_0^n \in \mathcal{S}$ and $I_0^n(L-1)\le y^n$,*

$$V_0^n(\mathbf{I}_0^n) - V_0^{H,n}(\mathbf{I}_0^n) \le \alpha(\alpha c - h+\theta)\mathbb{E}\left[y^n - D_1^n - D_2^n\right]^+ + \frac{\alpha^L - \alpha^T}{1-\alpha}(\alpha c - h + \theta)\mathbb{E}\left[y^n - D_{[L]}^n\right]^+ = o(1).$$

Proof. It is obvious that the optimal profit function under the nonperishable system, denoted by $\mathcal{V}_t^n(\mathbf{I}_t^n)$ with the understanding that all the inventory has infinite life time, is an upper bound of the optimal profit function of the perishable system. Let $O_t^{H,n}$ be the amount of inventory expiring at the end period of t under policy H, the following holds

$$V_0^n(\mathbf{I}_0^n) \le \max_{z,q}\left\{\mathbb{E}G(\mathbf{I}_0^n, z, q, D_1^n) + \alpha\mathbb{E}\mathcal{V}_1^n(\Gamma(\mathbf{I}_0^n, z, q, D_1^n))\right\}$$

$$
\begin{aligned}
&= \mathbb{E}G(\mathbf{I}_0^n, z_0^{H,n}, y^n - I_0^n(L-1) + z_0^{H,n}, D_1^n) + \alpha \mathbb{E}\mathcal{V}_1^n(\Gamma(\mathbf{I}_0^n, z_0^{H,n}, y^n - I_0^n(L-1) + z_0^{H,n}, D_1^n)) \\
&= V_0^{H,n}(\mathbf{I}_0^n) - \alpha \mathbb{E} V_1^{H,n}(\Gamma(\mathbf{I}_0^n, z_0^{H,n}, y^n - I_0^n(L-1) + z_0^{H,n}, D_1^n)) \\
&\quad + \alpha \mathbb{E}\mathcal{V}_1^n(\Gamma(\mathbf{I}_0^n, z_0^{H,n}, y^n - I_0^n(L-1) + z_0^{H,n}, D_1^n)) \\
&\leq V_0^{H,n}(\mathbf{I}_0^n) + (\alpha c - h + \theta) \sum_{t=2}^{T} \alpha^t \mathbb{E} O_t^{H,n}.
\end{aligned}
$$

The first inequality follows because the one period profit function of the original problem is bounded by that of the corresponding nonperishable system. The equalities follow due to the definitions of $z_0^{H,n}$, y^n, and $V_0^{H,n}(\mathbf{I}_0^n)$. The last inequality follows as the difference between $\mathcal{V}_1^n$ and $V_1^{H,n}$ under the same policy (that is, order up to y^n) is the expiration associated costs. Note that, each expired item incurs at most $\alpha c - h + \theta$ (disposal cost plus the cost to purchase a new unit later minus the inventory holding cost saved).

By (1), the total amount of inventory expiring during periods $2, \cdots, L$ satisfies

$$
\begin{aligned}
\sum_{t=2}^{L} \alpha^t \mathbb{E} O_t^{H,n} \leq \sum_{t=2}^{L} \mathbb{E} O_t^{H,n} &= \mathbb{E} \max \left\{0, I_0^n(1) - z_0^{H,n} - D_1^n, \cdots, I_0^n(L-1) - z_0^{H,n} - D_{[L-1]}^n, y^n - D_{[L]}^n \right\} \\
&\quad - \mathbb{E} \max \left(0, I_0^n(1) - z_0^{H,n} - D_1^n\right) \\
&\leq \mathbb{E} \max \{0, y^n - D_1^n - D_2^n\} = \mathbb{E} \left[y^n - D_1^n - D_2^n\right]^+
\end{aligned}
$$

Since $O_t^{H,n} \leq \left[y^n - \sum_{k=t-L+1}^{t} D_k^n\right]^+$ for $L < t \leq T$, so $\sum_{t=L}^{T} \alpha^t \mathbb{E} O_t^{H,n} \leq \frac{\alpha^L - \alpha^T}{1-\alpha} \mathbb{E} \left[y^n - D_{[L]}^n\right]^+$ (see the similar result $O_t(x, y, d) \leq B_t(y, d)$ in Section 6.2 of Chen et al. (2014)) as demand is independently and identically distributed. Thus, we have the desired bound.

Under Assumption 1, it is easy to see that $\frac{y^n}{n\lambda} \to 1$ as $n \to \infty$. Thus, $\mathbb{P}(D_1^n + D_2^n \leq y^n) = \mathbb{P}\left(\frac{D_1^n + D_2^n - 2n\lambda}{\sqrt{n}\sigma} \leq \frac{y^n - 2n\lambda}{\sqrt{n}\sigma}\right)$ decays exponentially as $n \to \infty$. So $\mathbb{E}\left[y^n - D_1^n - D_2^n\right]^+ \leq y^n \mathbb{P}(D_1^n + D_2^n \leq y^n) \to 0$ as $n \to \infty$. The same applies to $\mathbb{E}\left[y^n - D_{[L]}^n\right]^+$ and hence the bound is $o(1)$. □

For the special case of zero initial inventory, the bound in Proposition 4 reduces to

$$
V_0^n(\mathbf{0}) - V_0^{H,n}(\mathbf{0}) \leq \frac{\alpha^L - \alpha^T}{1-\alpha} (\alpha c - h + \theta) \mathbb{E}\left[y^n - D_{[L]}^n\right]^+ = o(1).
$$

Note that the performance bounds for both the fluid and newsvendor policies are independent of the salvage value s because the clearance decision can be ignored for a relatively large system except at the very beginning.

4.3. Numerical experiments

Since the bounds are functions of n, they provide lower bounds on the system size that guarantee any desired performance gap. These bounds help evaluate the effectiveness of the policies and provide practical reference.

Since the bounds overestimate the system sizes required to achieve a given performance of a heuristic and hence underestimate the effectiveness of the heuristics, we develop an efficient algorithm described in detail in Section 6 to search for the optimal decisions intelligently, which we also view as an important contribution of the paper. Coupled with parallel computing, we are able to solve relatively large problems optimally and provide more accurate evaluations of the heuristics.

To see how policies f and H fare, we consider various combinations of system parameters and batch size distributions.

1. System parameters: We fix $T=16$, $p=10$, $\theta=1$, $h=0.1c$, $\alpha=0.95$, and vary $c=2,4,6$, $s=0.5c,0.75c$ and $L=2,4$.

2. Batch size: We consider a batch size that follows constant, uniform on $[0,20]$, exponential and two-phase hyper-exponential distributions, all with a fixed mean of $\lambda=10$. Thus, the second moment of the deterministic, uniform and exponential distributions are $\sigma^2=100$, $400/3$, and 200, respectively. The two-phase hyper-exponential distribution is characterized by the parameter r. That is, with probability r the distribution is exponential with rate $r/5$, and with probability $1-r$ it is exponential with rate $(1-r)/5$. In our experiment, we set $r=1/8$ in which case the second moment $\sigma^2=450$.

Policy f: Since the results are similar, we only present the system sizes predicted by the bounds to achieve 2%, 3%, and 5% performance gaps in Table 1 and the percentage gaps in Table 2 for $\mathbf{I}_0=\mathbf{0}$, $c=4, L=4$ and $s=2$. While the performance improves as n increases, the crude fluid policy which sets the order-up-to level to the mean demand works well for relatively large systems, i.e., at $n=40$ when c is relatively high (see the Appendix for more numerical results). We also observe that zero initial inventory seems to have the worst performance. This can be explained as follows: recall the definition of relative performance gap $\frac{V_0^n(\mathbf{I})-V_0^{\pi,n}(\mathbf{I})}{V_0^n(\mathbf{I})}$. While the absolute gap $V_0^n(\mathbf{I})-V_0^{\pi,n}(\mathbf{I})$ is less affected by the initial inventory profile $\mathbf{I}$, the denominator $V_0^n(\mathbf{I})$ increases as $\mathbf{I}$ becomes large in terms of the total amount of inventory.

Performance gap	Constant	Uniform	Exponential	Hyper-exp
5%	32	43	64	104
3%	72	112	143	253
2%	149	191	253	508

Table 1 **Lower bounds on n to guarantee the desired performance gap under policy f for $L=4$ and $c=4$ when $\mathbf{I}_0^n=\mathbf{0}$**

n	Constant	Uniform	Exponential	Hyper-exp
10	8.17	9.73	11.55	14.67
20	5.80	6.82	8.18	11.27
30	4.73	5.54	6.67	9.39
40	4.10	4.79	5.78	8.19

Table 2 **Percentage gap of policy f for $L=4$ and $c=4$ when $\mathbf{I}_0^n=\mathbf{0}$**

Policy H: Table 3 presents the lower bounds on the system sizes that guarantee the desired performance gap for the same set of parameters. Again, the bounds are insensitive to the lifetime L and we only report the results for $L=4$. Note that the minimum required system sizes for policy H is much smaller than those for policy f and the bounds under heuristic H are very similar to those with $\mathbf{I}_0^n=\mathbf{0}$ as long as the total initial inventory is below y^n (and hence is omitted), indicating that policy H likely works very well even with non-zero initial inventory.

We also compare the performance of the heuristics against the true optimal for the same examples and Tables 4 provides a demo with more results in the Appendix. While the fluid heuristic works reasonably well for relatively large systems, policy H works extremely well even for very small systems. This is very encouraging in that a perishable inventory system may be managed just like a non-perishable one as long as the system is not too small and the initial inventory is not too "high".

Lastly, we compare policy H with the heuristic policies DB and PB in Chao et al. (2015) who proposed these policies for systems under backlogging and claimed that they can be extended to systems under lost sales. Since the authors did not provide any numerical study for the lost sales case, we will use the same parameters in this section (except that $s=1$, and vary $\theta=1,4$, to emphasize the effect of the disposal cost on the performance of different polices) but consider the demand distributions as in Nandakumar and Morton

(1993). As one can see Table 5, all heuristics work well for distributions with smaller c.v.s, while the results are mixed when the c.v. is relatively large (e.g., when demand is exponential). Given the extreme simplicity of policy H, it is not surprising that our heuristic does not outperform theirs in every case. However, if simplicity is a requirement for implementation, policy H can be much more appealing.

Performance gap	Constant	Uniform	Exponential	Hyper-exp
3%	6	9	14	35
2%	9	10	17	40
1%	10	14	20	49

Table 3 **Lower bounds on n to guarantee the desired performance gap under policy H for $c=4$ when $\mathbf{I}_0^n=\mathbf{0}$**

	$L=2$				$L=4$			
n	Const	Uniform	Exp	Hyp-exp	Const	Uniform	Exp	Hyp-exp
10	0.11	0.61	1.42	7.17	0.10	0.18	0.23	0.51
20	0.08	0.13	0.30	2.19	0.08	0.11	0.15	0.28
30	0.09	0.09	0.13	0.90	0.09	0.09	0.12	0.21

Table 4 **Percentage gap of policy H for $c=4$ when $\mathbf{I}_0^n=\mathbf{0}$**

5. Extensions

In this section, we extend our analysis beyond the base model by considering systems where unmet demand is backlogged, the distribution of batch sizes is unknown, a fixed ordering cost exists, and inventory is consumed last-in-first-out. We also examine numerically the performance of the heuristics for systems where the demand in each period is not compound Poisson.

5.1. The case of backlogging

Our results continue to hold when unmet demand is backlogged. In this case, we let $I_t(L-1)$ represent the net inventory level that can be positive or negative when there is backlogging and rewrite the state transition function in (1) for $\ell = L-1$

$$I_{t+1}(L-1) = [I_t(L-1) - z_t]^+ + q_{t+1} - D_{t+1} - [I_t(1) - z_t - D_{t+1}]^+ . \tag{9}$$

			Norm$(50, 20^2)$		Erlang-2		Exponential	
policy	c	θ	mean	max	mean	max	mean	max
H	2	1	0.06	0.06	0.28	0.29	1.84	1.88
	4	1	0.10	0.10	0.15	0.16	1.62	1.71
	6	1	0.17	0.18	0.05	0.05	0.93	1.04
	2	4	0.14	0.15	2.41	2.49	7.43	7.63
	4	4	0.14	0.14	0.30	0.31	3.65	3.86
	6	4	0.23	0.25	0.07	0.08	1.77	1.99
DB	2	1	1.79	1.83	2.07	2.10	2.22	2.27
	4	1	2.11	2.18	1.52	1.58	0.98	1.02
	6	1	1.08	1.15	0.15	0.16	0.14	0.16
	2	4	1.78	1.82	1.85	1.87	1.97	2.03
	4	4	2.10	2.18	1.42	1.48	0.64	0.71
	6	4	1.08	1.15	0.11	0.12	0.31	0.35
PB	2	1	1.64	1.70	1.69	1.77	1.78	1.90
	4	1	1.62	1.70	1.13	1.24	0.75	0.86
	6	1	0.71	0.76	0.20	0.29	1.11	1.27
	2	4	1.75	1.83	1.61	1.69	1.78	1.89
	4	4	1.64	1.72	1.11	1.22	0.62	0.71
	6	4	0.72	0.78	0.21	0.30	1.31	1.48

Table 5 **Percentage gap under polices H, DB and PB when $L=4$.“mean” and “max” are with respect to different initial inventory levels.**

If we view p as the unit cost for backlogging, our profit maximization formulation is the same as the cost minimization one in Chao et al. (2015). Following similar proofs, we can establish all the bounds in Sections 4.1 and 4.2 for the fluid policy f and newsvendor-type policy H in which case the order-up-to level is modified to $F_{D^n}^{-1}\left(\frac{p-c+\alpha c}{p+h}\right)$, the optimal order up to level for the corresponding non-perishable inventory system with backlogging.

A comparison of our policy H with policies DB and PB for the same examples with general demand distributions in Chao et al. (2015) and initial state $\mathbf{0}$ is summarized in Table 6. While all the policies perform well, our policy H outperforms when demand is Erlang-2 with a c.v. of 0.71 and underperforms when demand is exponential in which case the c.v. is higher at 1. We also add a comparison of the policies when demand is Normal (50, 20^2) which has an even smaller c.v.$= 0.4$ in the same table and all three policies perform equally well with our policy slightly better.

	Normal (50, 20^2)			Erlang-2			Exponential		
c, p, θ	H	DB	PB	H	DB	PB	H	DB	PB
0, 5, 5	0.03	3.30	5.56	0.10	0.56	0.50	1.08	1.21	0.13
0, 10, 5	0.00	7.64	11.88	0.29	1.37	2.30	2.49	0.88	0.29
0, 5, 10	0.03	3.51	5.54	0.28	0.56	0.29	3.31	1.86	0.58
5, 10, 5	0.12	0.48	0.58	0.11	0.26	0.36	1.37	0.29	0.02
5, 5, 10	0.29	0.04	0.05	0.02	0.08	0.01	0.86	0.39	0.42
5, 5, 5	0.29	0.03	0.05	0.01	0.08	0.02	0.44	0.22	0.26
5, 10, 0	0.12	0.34	0.58	0.03	0.07	0.09	0.41	0.18	0.06
10, 10, 5	0.17	0.08	0.12	0.06	0.11	0.12	1.03	0.37	0.07
10, 10, 10	0.17	0.10	0.12	0.10	0.11	0.10	0.53	0.34	0.38
10, 5, 5	0.57	0.21	0.12	0.01	0.02	0.00	0.26	0.44	0.32
10, 10, 0	0.17	0.06	0.12	0.03	0.10	0.16	0.52	0.28	0.02

Table 6 **Percentage gap of polices H, DB and PB under backlogging when $L = 4$.**

5.2. Unknown distribution of the batch sizes

So far, we have assumed that the distribution of batch sizes is known. Since the demand distribution of D_t^n approaches the normal distribution with mean $n\lambda$ and variance $n\sigma^2$ as $n \to \infty$, a modified H policy with the distribution of the demand replaced by the normal distribution, referred to as policy $\pi = \star$, should work well for large enough systems regardless of the batch size distribution or the total demand distribution. That is, under policy $\star$, we clear $z_0^{\star,n} = \left[I_0^n(1) - n\lambda - \sqrt{n}\sigma\Phi^{-1}\left(\frac{\alpha c - s - h}{\alpha c + \theta - h}\right)\right]^+$ amount of initial inventory and order up to $y^{\star,n} = n\lambda + \sqrt{n}\sigma\Phi^{-1}\left(\frac{p-c}{p-\alpha c+h}\right)$ at the beginning of each period. This policy only requires information about the first two moments of the batch size distribution under compound Poisson or the total demand. In Proposition 5, we present a theoretical performance gap bound for policy $\star$ when customer arrivals follow a Poisson distribution. Although the bound is a rough one due to technical difficulties, it vanishes asymptotically, implying asymptotic optimality of policy $\star$.

PROPOSITION 5. *Let m_3 be the third moment of the distribution of the batch size. There exists a constant $C(\lambda,\sigma,m_3)$ in a closed form and $\hat{n}$ such that, for $n>\hat{n}$,*

$$V_0^n(\mathbf{I}_0^n) - V_0^{\star,n}(\mathbf{I}_0^n) \leq \frac{C(\lambda,\sigma,m_3)}{\sqrt{n}}.$$

For the same examples in the previous sections, for $(\sigma^2, m_3) \in \{(1,1),(1.5,2),(2,3)\}$, Table 7 illustrates the system sizes that guarantee the desired relative performance gap for $L=4$ (the results for $L=2$ are similar and hence omitted). Numerical results (see the Appendix for details) show that although the required system sizes provided by the bounds are very large as expected, the heuristic works very well for systems of a reasonable size.

			(σ^2, m_3)		
Performance gap	c	s	(1,1)	(1.5,2)	(2,3)
3%	2	1	1107	651	469
	4	2	249	157	120
	6	4	175	110	85
2%	2	1	1194	714	521
	4	2	281	181	140
	6	4	192	124	98
1%	2	1	1398	862	643
	4	2	357	238	187
	6	4	236	162	131

Table 7 Lower bounds of n for desired performance gap under policy $\star$

5.3. There is a fixed ordering cost

Suppose that there is a fixed cost K for each order placed regardless of the order quantity, in which case it is not clear what the optimal structure looks like for systems with general lifetimes and few heuristic policies have been developed in the literature. For simplicity, we will focus on the infinite horizon case with $\mathbf{I}_0 = \mathbf{0}$ and leave out the subscript t in the cost functions. The optimal ordering strategy of the fluid model in the presence of a fixed ordering cost K is to order just enough to cover demand for t^* periods where t^* is either L or minimizes the total cost

$$\left[ctD + K + \frac{1-\alpha^t - t\alpha^t + t\alpha^{t+1}}{(1-\alpha)^2} Dh\right] / (1-\alpha^t), \tag{10}$$

whichever is smaller, and no clearance is needed. It is obvious that $t^* = 1$ in our asymptotic regime as the impact of the fixed ordering cost vanishes when $n \to \infty$ and $D = n\lambda \to \infty$. Thus, the corresponding fluid policy f is the same as that without a fixed ordering cost, i.e., order up to $n\lambda$ in every period in the nth system, and is asymptotically optimal in the presence of a fixed ordering cost.

However, in reality, even a "fixed" ordering cost will change as the order quantity exceeds a certain level. For instance, another truck would need to be dispatched if the order quantity exceeds the initial truck capacity and the fixed cost will double or triple as the order quantity increases. In this case, the fixed ordering cost needs to be scaled up in the asymptotic regime to preserve its impact in the original problem.

Suppose that the fixed cost is piece-wise linear as aK if the order quantity is in $[(a-1)m\lambda + 1, am\lambda]$, $a = 1, 2, ...$, for some positive integer m. Then, t^* is either L or minimizes the total cost

$$\left[cnDt + \left\lceil \frac{nDt}{m} \right\rceil K + \frac{1 - \alpha^t - t\alpha^t + t\alpha^{t+1}}{(1-\alpha)^2} nDh \right] / (1 - \alpha^t), \tag{11}$$

whichever is smaller, and no clearance is needed. Then, t^* decreases in m and the order-up-to level of the corresponding fluid policy f is $n\lambda t^*$ for the nth system. The next proposition states that the fluid policy f is not only asymptotically optimal but also has computable performance guarantee.

Proposition 6. *The optimality gap of the fluid policy f satisfies*

$$V^n(\mathbf{0}) - V^{f,n}(\mathbf{0}) \le \frac{\alpha}{1-\alpha}(p + \theta + h + K)\mathbb{E}|D_1^n - n\lambda| + \left[\frac{\alpha}{1-\alpha}(p+\theta) + \frac{1-\alpha^{t^*}}{(1-\alpha)^2}h\right] \mathbb{E}\,[n\lambda - D_t^n]^+ = O(\sqrt{n}).$$

5.4. The LIFO issuing rule

In practice, consumers may have the option to pick the freshest items, i.e., inventory is depleted on a last-in-first-out (LIFO) basis. Li et al. (2016) is one of the few that explores such a scenario and the structure of the optimal policy is even more complex under the LIFO rule.

From the modeling perspective, it suffices to modify the inventory dynamics in (1) to

$$I_{t+1}(\ell) = \left[[I_t(\ell+1) - z_t]^+ + q_{t+1}\mathbf{1}_{\{\ell = L-1\}} - [D_{t+1} - q_{t+1} - I_t(L-1) + I_t(\ell+1)]^+ - O_{t+1}\right]^+,$$

for $0 < \ell \le L-1$, where $O_{t+1} = \min\{I_t(1), [I_t(L-1) + q_{t+1} - D_{t+1}]^+\}$ is the amount of inventory that will be expired by the end of period $t+1$.

It is obvious that under the LIFO issuing rule, more inventory will be expired and hence, clearance of inventory may be unavoidable. Following a similar proof of Proposition 4 which provides a performance bound in the order of $o(1)$ under the FIFO issuing rule, we can establish a performance bound of heuristic H in the order of $\sqrt{n}$ under the LIFO issuing rule. Although policy H is likely not to perform as well under the LIFO rule, it is still asymptotically optimal.

COROLLARY 1. *For any* $\mathbf{I}_0^n \in \mathcal{S}$, $V_0^n(\mathbf{I}_0^n) - V_0^{H,n}(\mathbf{I}_0^n)$ *is of order* $\sqrt{n}$. *For* $\mathbf{I}_0^n = \mathbf{0}$,

$$V_0^n(\mathbf{0}) - V_0^{H,n}(\mathbf{0}) \leq \frac{\alpha^L - \alpha^T}{1-\alpha}(\alpha c - h + \theta)\mathbb{E}\min\left\{[y^n - D_1^n]^+, \cdots, [y^n - D_L^n]^+\right\} = O(\sqrt{n}).$$

Next, we modify the order-up-to level of policy H and devise a simple clearance policy in each period that takes into account the total amount of inventory and the amount of inventory that has remaining lifetime of one period. We refer to this heuristic policy as the RH heuristic policy whose asymptotic optimality can also be established with a similar effort.

The ordering policy of RH sets the order-up-to level at $y^n - \gamma I^n(1)$ where $\gamma = \frac{\alpha c - h + \theta}{p - \alpha c + h}$ by the following reason: The marginal benefit of ordering-up-to $x+1$ rather than x in period t is approximately

$$-c + p\mathbb{P}(D_t^n \geq x) - \theta\mathbb{P}(x - I_t^n(1) < D_t^n \leq x) + (\alpha c - h)\mathbb{P}(D_t^n < x - I_t^n(1)) \tag{12}$$

as it costs c to order the extra unit, enables one more unit of sales at p when $D_t^n \geq x$, potentially causes the disposal of one unit of old inventory at θ when $x - I^n(1) < D^n \leq x$, and gains up to $\alpha c - h$ due to extra inventory in the future when $D^n < x - I^n(1)$. Setting (12) to be 0 yields the order-up-to level $x = y^n - \gamma I^n(1)$.

The clearance policy of RH is inspired by an optimal property characterized in Li et al. (2016) that, when the oldest inventory is below a certain level, its chance of expiration is increased by the subsequent large order quantity and thus should be cleared.

- If $I_t^n(L-1) < y^n - \gamma I_t^n(1)$, clear $I_t^n(1)$ when $I_t^n(1) \leq y^n - F_{D_t^n}^{-1}\left(\frac{\alpha c - s - h}{\alpha\theta + \alpha s}\right)$ and no clearance otherwise.
- If $I_t^n(L-1) \geq y^n - \gamma I_t^n(1)$, clear $\min\left\{I_t^n(1), \left[I_t^n(L-1) - F_{D_t^n}^{-1}\left(\frac{\alpha p - s - h - \alpha\theta\mathbb{P}(D^n \leq I^n(L-1) - I^n(1))}{\alpha p}\right)\right]^+\right\}$.

When $I_t^n(L-1) < y^n - \gamma I_t^n(1)$, an order will be placed in period $t+1$. Since it saves θ and gains at least s if $D^n \leq y^n - I^n(1)$, the marginal benefit of clearing one unit of the oldest inventory $I_t^n(1)$ at the end of period t and then ordering one fresh unit at the beginning of period $t+1$ is at least $s+h-\alpha c+\alpha\mathbb{P}(D_t^n \leq y^n - I_t^n(1))(\theta+s) \geq s+h-\alpha c+\alpha\mathbb{P}\left(D_t^n \leq F_{D_t^n}^{-1}\left(\frac{\alpha c-s-h}{\alpha\theta+\alpha s}\right)\right)(\theta+s)=0$ when $I^n(1) \leq y^n - F_{D^n}^{-1}\left(\frac{\alpha c-s-h}{\alpha\theta+\alpha s}\right)$.

When $I_t^n(L-1) \geq y^n - \gamma I_t^n(1)$, the ordering decision in period $t+1$ will depend on the amount of inventory cleared at the end of period t. Since it is most likely that no order will be placed in period $t+1$, the marginal benefit of clearing one unit of inventory with remaining lifetime of 1 period is at least $s+h-\alpha p\mathbb{P}(D_t^n > I_t^n(L-1)) + \alpha\theta\mathbb{P}(D_t^n \leq I_t^n(L-1) - I_t^n(1)) \geq 0$ as you sell one less unit at p when $D_t^n > I_t^n(L-1)$ and dispose one less unit at the cost θ when $D_t^n \leq I_t^n(L-1) - I_t^n(1)$ in period $t+1$, resulting in the heuristic clearance decision.

We conduct a numerical experiment on the same system parameters and demand distributions as in Section 4.3 for $L=4$ under the LIFO rule. To emphasize the effect of the disposal cost on the performance of the polices, we fix the salvage value $s=1$ and vary the disposal cost $\theta=1,4$. As one can see in Table EC.8 in the Appendix, policy RH works very well for systems with reasonable sizes. Since policy H ignores clearance except at the beginning, policy RH should outperform H when the disposal cost θ is high, which is confirmed in Table 8 when $\theta=4$.

Lastly, we examine the performance of policy RH under the same settings in Li et al. (2016), i.e., demand in each period is uniform with different c.v.s. and the results are summarized in Table 9 (we omit the simple case with $\theta=0.2$). Compared with the performance of policies $MH2$ and $MH3$ proposed in Li et al. (2016), policy RH outperforms $MH2$ but underperforms $MH3$. Given that their heuristics require solving a non-linear optimization problem to find both the ordering and clearance decisions in each period, heuristic RH that provides closed forms for both decisions is of practical significance.

5.5. General demand distributions

So far, we have assumed that demand in each period follows a compound Poisson distribution. We now ask whether the fluid and newsvendor policies can continue to perform well under other demand distributions. We conduct a numerical experiment on the same system parameters but with five alternative demand distributions with the same mean for $L=4$. These distributions are Uniform$(25,35)$, Uniform$(20,40)$, Normal$(30,5^2)$, Normal$(30,10^2)$ and Exponential$(1/30)$. The performance gap under the two policies with zero initial inventory are summarized in Table EC.9 (see the appendix for details). As we can see, although

		Constant		Uniform		Exponential		Hyper-exp	
Policy	c	mean	max	mean	max	mean	max	mean	max
RH	2	0.39	2.00	0.50	2.55	0.50	3.28	0.61	5.81
	4	0.79	1.98	0.97	2.01	1.21	2.61	2.46	4.91
	6	1.07	2.12	1.38	2.41	1.76	3.12	3.53	5.91
H	2	4.61	4.95	5.83	6.30	7.61	8.07	11.07	12.64
	4	3.86	4.88	4.63	5.93	5.98	7.66	11.04	12.69
	6	3.87	5.28	4.60	6.32	5.63	7.90	9.67	13.58

Table 8 Mean and maximum percentage performance gap of policies H and RH under the LIFO issuing rule for $\theta = 4$ and $L = 4$.

	$p = 5$				$p = 10$			
L	3	4	5	6	3	4	5	6
c.v.=0.2								
θ= 0.5	0.92	0.77	0.68	0.62	0.62	0.54	0.52	0.51
θ= 1.0	0.98	0.69	0.58	0.54	0.59	0.52	0.51	0.51
c.v.=0.3								
θ= 0.5	1.28	1.01	0.86	0.74	0.76	0.62	0.58	0.57
θ= 1.0	1.37	0.89	0.70	0.62	0.72	0.59	0.57	0.56
c.v.=0.4								
θ= 0.5	2.06	1.65	1.41	1.23	1.20	1.02	0.96	0.94
θ= 1.0	2.08	1.40	1.14	1.04	1.13	0.98	0.95	0.94
c.v.=0.5								
θ= 0.6	2.12	1.72	1.49	1.32	1.31	1.11	1.06	1.04
θ= 1.0	2.20	1.47	1.21	1.11	1.22	1.06	1.04	1.04

Table 9 Percentage performance gap of RH under the LIFO rule.

the fluid policy may not work so well when demand is exponential or normal and highly uncertain, the mean performance gap of the policy H are fairly small for all the distributions.

6. An efficient parallel algorithm for finding an optimal solution

The bounds developed in sections 4.1 and 4.2 provide a rough idea about the system sizes required to guarantee a desired performance gap under the heuristics. The minimum system size may be much smaller and a comparison with the optimal performance can provide more accurate information. However, solving a dynamic program to find an optimal solution is time consuming for problems with reasonable large sizes. In this section, we will exploit some properties of the problems that will aid in finding optimal solutions using parallel computing for the discrete version of the problem. With discrete units, the state variable $\mathbf{I}$ is of dimension $L-1$ and for convenience we let the ℓth element be the inventory with remaining lifetime of ℓ in this section. Let J be the maximum order quantity in each period, which can be the order quantity from the newsvendor problem with non-perishable inventory. Then, the dynamic programming problem is one with a state space of size $(J+1)^{L-1}$ and the complexity of finding an optimal solution increases exponentially in the life time L. Furthermore, since $0 \le q \le J$ and $0 \le z \le J(L-1)$, an exhaustive search finding the optimal decisions for a given state will require roughly $O(J^2)$ effort. Thus the overall complexity of solving the dynamic program is $O(TJ^{L+1})$.

Next, we identify some structural properties of the optimal solution in Proposition 7 using a tree representation of the state space. We then develop an efficient algorithm to reduce the computational complexity to $O(TJ^{L-1})$. Using parallel computing, we can solve a problem where the system is of a reasonable size in minutes.

Figure 1 demonstrates the way the state is evolved when one unit of inventory is depleted for the case with $L = 4$ and $J = 2$. For instance, clearing one unit at a time, the state evolves according to the sequence $(2,2,2) \to (1,2,2) \to \cdots \to (0,0,2) \to (0,0,1) \to (0,0,0)$. So the diagram is a tree where a predecessor has a unique successor which may have up to $L-1$ predecessors. Next, we present some properties of the optimal decisions along a path from the bottom of the tree to the state $\mathbf{0}$, which will greatly reduce the computational complexity.

Proposition 7. *For any state, if the optimal clearance quantity (z) is positive, then it is optimal for its successor to clear one less and order the same amount of inventory; else, it is optimal for its successor to clear nothing and order the same or one more.*

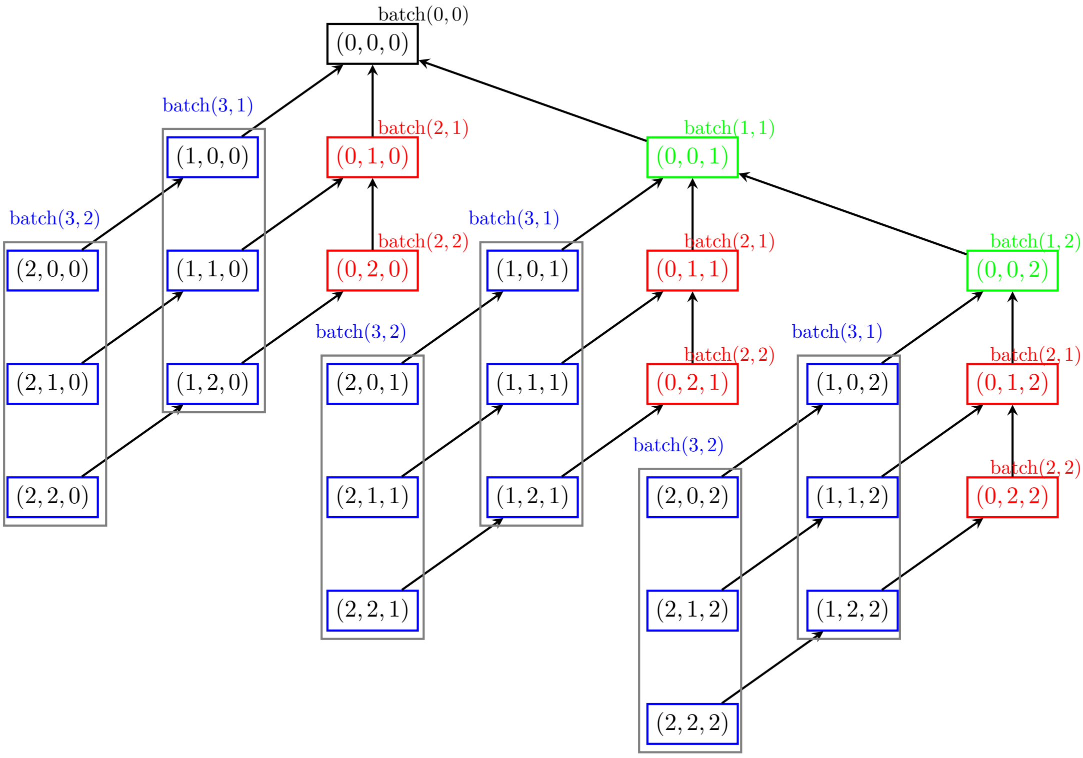

Figure 1 **A tree representation of the states when** $L = 4, J = 2$**.**

For instance, if the optimal action of state $(2, 2, 2)$ is $(z^*, q^*) = (3, 2)$, then we know from the above proposition that the optimal action of $(1, 2, 2)$ is $(2, 2)$; if $(z^*, q^*) = (0, 2)$, then the optimal action of $(1, 2, 2)$ is $(0, 2)$ or $(0, 3)$.

In each period from the last to the first, we need to evaluate all the possible options of (z, q) for each state, referred to as a job. Since there are a total of $(J+1)^{L-1}$ number of states, a number that is far bigger than the available number of computing units, we send the jobs for evaluating and finding the optimal decisions for all the states in parallel batch by batch. Several parallel computing application program interfaces (API) such as OpenCl and CUDA can take advantage of to utilize a large number of computing cores (about 2000) in a graphics processing unit. The computation of the optimal decisions in each period starts with the state $\mathbf{0}$ and followed by its predecessors.

To take advantage of the properties of the optimal decisions along the paths on a tree, we devise a batching and sequencing rule for sending the jobs. We batch the jobs with the same index (ℓ, k) where $L - \ell$ is the position of the first non-zero element in a state and k is its value. The jobs are sent to be processed according to the ascending ordering first of ℓ and then of k. We say that a batch (ℓ, k) has a smaller index than (ℓ', k') if $\ell < \ell'$ or

$\ell = \ell'$ but $k < k'$. The next proposition shows that such a batching and sequencing rule guarantees that, starting from state $\mathbf{0}$, all the predecessors will be processed after their successors as demonstrated in Figure 1. An algorithm based on the proposition can be easily devised.

PROPOSITION 8. *All the successors of batch (ℓ, k) must be in a batch with smaller indices.*

Algorithm 1 Parallel Computing of Dynamic Program

Initialization $V(\cdot) = V_T(\cdot)$;
for $t = T : 1$ **do**
  compute the optimal policy for state 0 (i.e., batch $(0,0)$) using an exhaustive search;
  for $\ell = 1 : L$ **do**
    for $k = 1 : (J-1)$ **do**
      compute the optimal policy for all states in batch (ℓ, k) using proposition 7;
    end for
  end for
  Update $V(\cdot) = V_{t-1}(\cdot)$;
end for
Output $V(\cdot) = V_0(\cdot)$.

7. Conclusions

It is well known that perishable inventory problems are difficult to optimize due to curse of dimensionality and even heuristic ordering and clearance policies involving solving complicated equations or optimization problems. Inspired by the asymptotic behavior of perishable inventory systems, we consider two extremely simple policies that ignore clearance except at the very beginning of the planning horizon and order up to the mean demand (called a fluid based policy) and a newsvendor type of solution (called a newsvendor type of policy), respectively. We provide analytical performance bounds for the policies which vanish as the size of the system increases, implying asymptotic optimality of the policies and providing lower bounds of the system size to achieve any desired performance. Numerical experiments against the bounds as well as the optimal solutions show that these policies either work reasonably well or extremely well for small-to-medium-sized systems, suggesting that managing perishable inventory systems of reasonable sizes may be just as simple as managing a non-perishable ones. In addition, it is observed that when lifetime is small ($L = 2$) or the demand distribution has high coefficient of variations (e.g., hyper-exponential), the performance of our policies deteriorates. It remains an open question for simple polices in such circumstances. Given that analytical and tight bounds on the performance of heuristic policies are rarely established in the perishable inventory literature, our research makes a significant contribution to both the existing literature and real practice.

Acknowledgments

References

Allon G, Van Mieghem JA (2010) Global dual sourcing: Tailored base-surge allocation to near-and offshore production. *Mgt. Sci.* 56(1):110–124.

Armony M, Shimkin N, Whitt W (2009) The impact of delay announcements in many-server queues with abandonment. *Oper. Res.* 57(1):66–81.

Baron O, Berman O, Perry D (2010) Continuous review inventory models for perishable items ordered in batches. *Mathematical Methods of Operations Research* 72(2):217–247.

Bassamboo A, Randhawa RS (2010) On the accuracy of fluid models for capacity sizing in queueing systems with impatient customers. *Oper. Res.* 58(5):1398–1413.

Bassamboo A, Randhawa RS (2016) Scheduling homogeneous impatient customers. *Mgt. Sci.* 62(7):2129–2147.

Bassamboo A, Randhawa RS, van Mieghem JA (2012) A little flexibility is all you need: On the asymptotic value of flexible capacity in parallel queuing systems. *Oper. Res.* 60(6):1423–1435.

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

Chen X, Pang Z, Pan L (2014) Coordinating inventory control and pricing strategies for perishable products. *Oper. Res.* 62(2):284–300.

Cohen MA (1976) Analysis of single critical number ordering policies for perishable inventories. *Oper. Res.* 24(4):726–741.

Cooper WL (2001) Pathwise properties and performance bounds for a perishable inventory system. *Oper. Res.* 49(3):455–466.

David I, Mehrez A (1995) An inventory model with exogenous failures. *Oper. Res.* 43(5):902–903.

Durrett R (2010) *Probability: Theory and Examples* (Cambridge University Press), 4th edition.

Fries BE (1975) Optimal ordering policy for a perishable commodity with fixed lifetime. *Oper. Res.* 23(1):46–61.

Goldberg DA, Katz-Rogozhnikov DA, Lu Y, Sharma M, Squillante MS (2016) Asymptotic optimality of constant-order policies for lost sales inventory models with large lead times. *Math. Oper. Res.* 41(3):898–913.

Gürler Ü, Özkaya BY (2008) Analysis of the (s, s) policy for perishables with a random shelf life. *IIe Transactions* 40(8):759–781.

Karaesmen IZ, Scheller-Wolf A, Deniz B (2011) Managing perishable and aging inventories: review and future research directions. *Planning production and inventories in the extended enterprise*, 393–436 (Springer).

Levi R, Pál M, Roundy RO, Shmoys DB (2007) Approximation algorithms for stochastic inventory control models. *Math. Oper. Res.* 32(2):284–302.

Li Q, Yu P (2014) Multimodularity and its applications in three stochastic dynamic inventory problems. *Manufacturing & Service Operations Management* 16(3):455–463.

Li Q, Yu P, Wu X (2016) Managing perishable inventories in retailing: Replenishment, clearance sales, and segregation. *Oper. Res.* 64(6):1270–1284.

Nahmias S (1975) Optimal ordering policies for perishable inventory—ii. *Oper. Res.* 23(4):735–749.

Nahmias S (1976) Myopic approximations for the perishable inventory problem. *Mgt. Sci.* 22(9):1002–1008.

Nahmias S (1977a) Higher-order approximations for the perishable-inventory problem. *Oper. Res.* 25(4):630–640.

Nahmias S (1977b) Technical note—comparison between two dynamic perishable inventory models. *Oper. Res.* 25(1):168–172.

Nahmias S (1982) Perishable inventory theory: A review. *Oper. Res.* 30(4):680–708.

Nahmias S (2011) *Perishable inventory systems*, volume 160 (Springer).

Nahmias S, Pierskalla WP (1973) Optimal ordering policies for a product that perishes in two periods subject to stochastic demand. *Naval Research Logistics Quarterly* 20(2):207–229.

Nandakumar P, Morton TE (1993) Near myopic heuristics for the fixed-life perishability problem. *Mgt. Sci.* 39(12):1490–1498.

Perry O, Whitt W (2011) A fluid approximation for service systems responding to unexpected overloads. *Oper. Res.* 59(5):1159–1170.

Scarf H (1959) Bayes solutions of the statistical inventory problem. *The Annals of Mathematical Statistics* 30(2):490–508.

Tyurin IS (2010) Refinement of the upper bounds of the constants in lyapunov's theorem. *Russian Mathematical Surveys* 65(3):586–588.

Wu CA, Bassamboo A, Perry O (2017) Service systems with dependent service and patience times .

Xin L, Goldberg DA (2016) Optimality gap of constant-order policies decays exponentially in the lead time for lost sales models. *Oper. Res.* 64(6):1556–1565.

Xin L, Goldberg DA (2017) Asymptotic optimality of tailored base-surge policies in dual-sourcing inventory systems. *Mgt. Sci.* .

Xue Z, Ettl M, Yao DD (2012) Managing freshness inventory – optimal policy, bounds and heuristics. Technical report, Columbia University.

Zhang H, Shi C, Chao X (2016) Approximation algorithms for perishable inventory systems with setup costs. *Operations Research* 64(2):432–440.

Zipkin P (2008) On the structure of lost-sales inventory models. *Oper. Res.* 56(4):937–944.

Appendix: Managing Perishable Inventory Systems as Non-perishable Ones

EC.1. Technical proofs

Proof of Lemma 1. It is easy to see that the lemma holds for $t = T$. Suppose it is true for $t+1 \leq T$. Since $G(\mathbf{I}, z, q, D)$ is piecewise linear and differentiable except at finitely many points, $v_t(\mathbf{I}, z, q, D)$ is differentiable in each element almost everywhere. Let $(z^*(\mathbf{I}), q^*(\mathbf{I}))$ be the maximizer of $\mathbb{E}[v_t(\mathbf{I}, z, q, D)]$. Since $\frac{\partial v_t(\mathbf{I},z,q,D)}{\partial z}\Big|_{\substack{z=z^*(\mathbf{I})\\ q=q^*(\mathbf{I})}} = \frac{\partial v_t(\mathbf{I},z,q,D)}{\partial q}\Big|_{\substack{z=z^*(\mathbf{I})\\ q=q^*(\mathbf{I})}} = 0$, using the chain rule, the following holds

$$\begin{aligned}\frac{\partial V_t(\mathbf{I})}{\partial I(\ell)} &= \mathbb{E}\left[\frac{\partial v_t(\mathbf{I}, z, q, D)}{\partial I(\ell)}\right]_{\substack{z=z^*(\mathbf{I})\\ q=q^*(\mathbf{I})}} + \mathbb{E}\left[\frac{\partial v_t(\mathbf{I}, z, q, D)}{\partial z}\Big|_{\substack{z=z^*(\mathbf{I})\\ q=q^*(\mathbf{I})}} \frac{\partial z^*(\mathbf{I})}{\partial I(\ell)}\right] \\ &\quad + \mathbb{E}\left[\frac{\partial v_t(\mathbf{I}, z, q, D)}{\partial q}\Big|_{\substack{z=z^*(\mathbf{I})\\ q=q^*(\mathbf{I})}} \frac{\partial q^*(\mathbf{I})}{\partial I(\ell)}\right] \\ &= \mathbb{E}\left[\frac{\partial v_t}{\partial I(\ell)}(\mathbf{I}, z^*(\mathbf{I}), q^*(\mathbf{I}), D)\right]\end{aligned}$$

exists almost everywhere, where in the first equality we used the fact that $\frac{\partial z^*(\mathbf{I})}{\partial I(\ell)}$ and $\frac{\partial q^*(\mathbf{I})}{\partial I(\ell)}$ are bounded by Corollary 1 of Chen et al. (2014). Hence, $V_t(\mathbf{I})$ is differentiable in each component almost everywhere. $\square$

Proof of Lemma 2. Let $\varphi(t)$ be the characteristic function of the batch size distribution, then the characteristic function of $\hat{D}_t^n$ can be expressed as (after some calculation) $\exp\left(-\frac{1}{2}t^2\frac{\varphi''(\xi)}{\sigma^2}\right)$. Now by the discussion on P.98 of Durrett (2010), $\varphi''(\xi)$ is bounded near 0, so by Theorem 3.3.5 of Durrett (2010), we reach the desired conclusion. $\square$

Proof of Proposition 1. We first note that $z_0^* \geq \max_{1\leq \ell \leq L-1} [I_0(\ell) - \ell D]^+$ as it only incurs holding costs to hold a unit that is due to expire. Likewise, $z_t^* = 0$ when $t \geq 1$ as any unit of inventory that will be cleared at t should either be depleted at the end of period

0 or not be ordered in the first place. It is obvious that, at the beginning of period t, an order is placed only if there is not enough inventory to meet the demand of that period. Thus, $q_t^* = \left[D_t - I_{t-1}(L-1) + z_{t-1}^*\right]^+$ and there is no inventory will expire after period 0.

Next, we derive the optimal clearance quantity z_0^*. Let $q_1^*(z_0)$ be the optimal order quantity given that we deplete z_0 at the end of period 0. We first claim that $z_0^* \le [I_0(L-1) - D]^+$ because, for any $z_0 > [I_0(L-1) - D]^+$, an order up to D_1 is placed in period 1 and $\frac{dq_1^*(z)}{dz} \mid_{z=z_0} = 1$. In this case, $\frac{dv_0(\mathbf{I}_0, z, q_1^*(z))}{dz} \mid_{z=z_0} = s + h - \alpha c < 0$, thus z_0 cannot be optimal. For any $\max_{1\le \ell \le L-1} [I_0(\ell) - \ell D]^+ \le z_0 < [I_0(L-1) - D]^+$, $q_1^*(z_0) = 0$ and

$$\begin{aligned}\frac{\partial v_0(\mathbf{I}_0, z, 0, D)}{\partial z} \mid_{z=z_0} &= \frac{\partial G(\mathbf{I}_0, z, 0, D)}{\partial z} \mid_{z=z_0} + \alpha \frac{\partial V_1(\Gamma(\mathbf{I}_0, z, 0, D))}{\partial z} \mid_{z=z_0} \\ &= s + h + \alpha \frac{\partial V_1(\Gamma(\mathbf{I}_0, z, 0, D))}{\partial z} \mid_{z=z_0}. \end{aligned} \tag{EC.1}$$

Fix $\delta > 0$. Note that $\Gamma(\mathbf{I}_0, z_0 + \delta, 0, D)$ and $\Gamma(\mathbf{I}_0, z_0, 0, D)$ differ by δ units. If the extra δ units are consumed in period $\ell = \ell(z_0) \in \{2, \cdots, L-1\}$, which depends on z_0, then

$$V_1(\Gamma(\mathbf{I}_0, z_0, 0, D)) - V_1(\Gamma(\mathbf{I}_0, z_0 + \delta, 0, D) = \left(\alpha^{\ell-1} c - \sum_{i=0}^{\ell-2} \alpha^i h\right) \delta = \left(\alpha^{\ell-1} c - \frac{1-\alpha^{\ell-1}}{1-\alpha} h\right) \delta.$$

That is, with the extra δ units, one needs to order δ units less in period ℓ but incurs the holding costs. Thus, by (EC.1), $\frac{\partial v_0(\mathbf{I}_0, z, 0, D)}{\partial z} \mid_{z=z_0} = s + h - \alpha \left(\alpha^{\ell-1} c - \frac{1-\alpha^{\ell-1}}{1-\alpha} h\right) = s + \frac{1-\alpha^{\ell}}{1-\alpha} h - \alpha^{\ell} c$, which is increasing in ℓ.

If $s + \frac{1-\alpha^{L-1}}{1-\alpha} h < \alpha^{L-1} c$, then $\frac{\partial v_0(\mathbf{I}_0, z, 0, D)}{\partial z} \mid_{z=z_0} < 0$, and one depletes as little inventory as possible and only those inventory items that are due to expire, i.e., $z_0^* = \max_{1\le \ell \le L-1} [I_0(\ell) - \ell D]^+$. Otherwise, by the definition of ℓ_0 in (5), z_0^* is the minimum quantity such that $\ell(z_0^*) \le \ell_0$, i.e., $I_0(L-1) - z_0^* \le \ell D$. Combing these two cases (4) holds. □

Proof of Proposition 3. By Lemma 3,

$$V_0^n(\mathbf{I}_0^n) - V_0^{f,n}(\mathbf{I}_0^n) \le n \bar{V}_0(\bar{\mathbf{I}}_0) - V_0^{f,n}(\mathbf{I}_0^n) = n \sum_{t=0}^{T} \alpha^t \mathbb{E} \Delta \bar{G}_t^n,$$

where

$$\Delta \bar{G}_t^n = G(\bar{\mathbf{I}}_t, \bar{z}_t^*, \bar{q}_{t+1}^*, \lambda) - G(\bar{\mathbf{I}}_t^n, \bar{z}_t^n, \bar{q}_{t+1}^n, \bar{D}_{t+1}^n)$$

is the difference between the single-period costs of the fluid model and that of the stochastic model under the fluid policy in period t. According to the definition of the fluid policy,

$$\Delta \bar{G}_t^n = h\left[\bar{I}_t^n(L-1) - \bar{I}_t(L-1)\right] + \alpha c \left(\left[\lambda - \bar{I}_t^n(L-1)\right]^+ - \left[\lambda - \bar{I}_t(L-1)\right]^+\right) + \alpha\theta \bar{O}_{t+1}^n$$

$$
\begin{aligned}
&+\alpha p\left(\lambda-\min\left\{\bar{D}^n_{t+1},\lambda+\left[\bar{I}^n_t(L-1)-\lambda\right]^+\right\}\right)\\
&\le h\left[\bar{I}^n_t(L-1)-\bar{I}_t(L-1)\right]+\alpha c\left(\left[\lambda-\bar{I}^n_t(L-1)\right]^+-\left[\lambda-\bar{I}_t(L-1)\right]^+\right)+\alpha\theta\bar{O}^n_{t+1}\\
&+\alpha p\left[\lambda-\bar{D}^n_{t+1}\right]^+. \qquad \text{(EC.2)}
\end{aligned}
$$

1. When $t\ge L$, using the order-up-to structure and given that the lifetime of the inventory is L, it is easy to show that $\bar{I}_t(L-1)=0$ and $\bar{I}^n_t(L-1)=\left[\lambda-\bar{D}^n_t-\bar{O}^n_t\right]^+\le\lambda$, so in this case

$$
\begin{aligned}
\Delta\bar{G}^n_t &\le -(\alpha c-h)\left[\lambda-\bar{D}^n_t-\bar{O}^n_t\right]^+ +\alpha\theta\bar{O}^n_{t+1}+\alpha p\left[\lambda-\bar{D}^n_{t+1}\right]^+\\
&\le -(\alpha c-h)\left[\lambda-\bar{D}^n_t\right]^+ +(\alpha c-h)\bar{O}^n_t+\alpha\theta\bar{O}^n_{t+1}+\alpha p\left[\lambda-\bar{D}^n_{t+1}\right]^+.
\end{aligned}
$$

Using $\mathbb{E}\bar{O}^n_t\le\mathbb{E}\left[\lambda-\bar{D}^n_{[L]}\right]^+$, we have

$$
\begin{aligned}
\sum_{t=L}^T\alpha^t\mathbb{E}\Delta\bar{G}^n_t &\le \sum_{t=L}^T\alpha^t\left((\alpha p-\alpha c+h)\mathbb{E}\left[\bar{D}^n_1-\lambda\right]^+ +(\alpha c+\alpha\theta-h)\mathbb{E}\left[\lambda-\bar{D}^n_{[L]}\right]^+\right)\\
&=\frac{\alpha^L-\alpha^{T+1}}{1-\alpha}(\alpha p-\alpha c+h)\mathbb{E}\left[\bar{D}^n_1-\lambda\right]^+ +\frac{\alpha^L-\alpha^{T+1}}{1-\alpha}(\alpha c+\alpha\theta-h)\mathbb{E}\left[\lambda-\bar{D}^n_{[L]}\right]^+\\
&=\frac{\alpha^L-\alpha^{T+1}}{1-\alpha}(\alpha p-\alpha c+h)\mathbb{E}\left[\bar{D}^n_1-\lambda\right]^+ +K_2\mathbb{E}\left[\lambda-\bar{D}^n_{[L]}\right]^+.
\end{aligned}
$$

2. When $0\le t<L$,

$$
\begin{aligned}
\sum_{t=0}^{L-1}\alpha^t\Delta\bar{G}_t &\le \alpha c\left\{\sum_{t=0}^{L-1}\alpha^t\left[\bar{I}_t(L-1)-\bar{I}^n_t(L-1)\right]^+\right\}+h\left\{\sum_{t=0}^{L-1}\alpha^t\left[\bar{I}^n_t(L-1)-\bar{I}_t(L-1)\right]^+\right\}\\
&\quad+\alpha\theta\sum_{t=0}^{L-1}\alpha^t\bar{O}^n_{t+1}+\alpha p\sum_{t=0}^{L-1}\alpha^t\left[\lambda-\bar{D}^n_{t+1}\right]^+.
\end{aligned}
$$

The last two terms are easy to bound and, in the second term,

$$
\begin{aligned}
\left[\bar{I}^n_t(L-1)-\bar{I}_t(L-1)\right]^+ &= \left[\bar{I}^n_{t-1}(L-1)+\bar{q}^n_t-\bar{O}^n_t-\bar{D}^n_t\right]^+-\left[\bar{I}_{t-1}(L-1)+\bar{q}_t-\lambda\right]^+\\
&\le \left[\bar{I}^n_{t-1}(L-1)-\bar{I}_{t-1}(L-1)\right]^+ +\left[\lambda-\bar{D}^n_t\right]^+\\
&=\sum_{k=1}^t\left[\lambda-\bar{D}^n_k\right]^+.
\end{aligned}
$$

It remains for us to find an upper bound of the first term. Let $\bar{\tau}\le L-1$ be the first time such that $\bar{I}_t(L-1)<\lambda$ and $\tau^n=\min\{t:\bar{I}^n_t(L-1)<\lambda\}$. We only need to consider those periods in which the stochastic model orders more than the fluid model, i.e., the periods from τ^n to $\bar{\tau}$ if $\tau^n < \bar{\tau}$. The first term is the extra ordering cost beyond that incurred in the fluid model, which is bounded by $c[(\bar{\tau}-\tau^n)\lambda - \bar{I}^n_{\tau^n}(L-1) + \bar{I}_{\bar{\tau}}(L-1)]\mathbf{1}_{\{\bar{\tau}>\tau^n\}} \leq c\left[\bar{I}_{\tau^n}(L-1) - \bar{I}^n_{\tau^n}(L-1)\right]^+$.

Note that

$$\left[\bar{I}_t(L-1) - \bar{I}^n_t(L-1)\right]^+ \leq \sum_{k=1}^{t}\left(\left[\bar{D}^n_k - \lambda\right]^+ + \bar{O}^n_k\right),$$

And the first term is bounded by

$$c\sum_{k=1}^{\tau^n}\left(\left[\bar{D}^n_k - \lambda\right]^+ + \bar{O}^n_k\right) \leq c\sum_{k=1}^{L-1}\left(\left[\bar{D}^n_k - \lambda\right]^+ + \bar{O}^n_k\right).$$

By taking expectation and using the fact that $\sum\limits_{k=1}^{t}\bar{O}^n_k = \max\left\{0, \lambda - \bar{D}^n_1, \cdots, \lambda t - \bar{D}^n_{[t]}\right\}$, we obtain the desired bound. Noting that both $\mathbb{E}\left[\bar{D}^n_1 - \lambda\right]^+$ and $H(n)$ are of order $\frac{1}{\sqrt{n}}$, and $\mathbb{E}\left[\lambda - \bar{D}^n_{[L]}\right]^+$ is of lower order of $\frac{1}{\sqrt{n}}$ (see the end of the proof of Proposition 4), the conclusion follows by multiplying both sides of (EC.2) by n. □

LEMMA EC.1. *$\left|F_{\tilde{D}^n_t}(x) - \Phi(x)\right| \leq \frac{C_0}{\sqrt{n}}$ for any $x>0$, where $C_0 = \frac{m_3+2\lambda^3}{2\sigma^3}$ and m_3 is the third moment of the batch size distribution.*

Proof of Lemma EC.1. Since a Poisson process is infinitely divisible, D^n_t can be written as $\sum\limits_{i=1}^{n} X_i$ where $X_i = \xi_{i,1} + \cdots + \xi_{i,N(1)} \geq 0$, $\xi_{i,j}$ are independently and identically distributed random variables representing the batch size, and $N(1) \sim \text{Poisson}(1)$. Then, the scaled demand can be written as the normalized sum as $\tilde{D}^n_t = \frac{\sum\limits_{i=1}^{n} X_i - \lambda n}{\sqrt{n}\sigma}$. Since X_i are i.i.d. with mean λ and second moment σ^2, the Berry-Esseen Theorem (see Tyurin (2010)) establishes the convergence

$$\left|F_{\tilde{D}^n_t}(x) - \Phi(x)\right| \leq \frac{\mathbb{E}|X_i-\lambda|^3}{2\sigma^3\sqrt{n}}.$$

Since $\mathbb{E}|X_i-\lambda|^3 = \mathbb{E}(X_i-\lambda)^3 + 2\mathbb{E}|X_i-\lambda|^3\mathbf{1}_{\{X_i\leq\lambda\}} \leq \mathbb{E}(X_i-\lambda)^3 + 2\lambda^3 = m_3 + 2\lambda^3$, we obtain the desired result. □

LEMMA EC.2. *Let $\beta = F_{D^n}(y^n) = \frac{p-c}{p-\alpha c+h}$, $\gamma = F_{D^n}(w^n) = \frac{\alpha c-s-h}{\alpha c+\theta-h}$ and*

$$C_2(y,n) = \min\left\{\phi(\Phi^{-1}\left(y - C_0/\sqrt{n}\right)), \phi(\Phi^{-1}\left(y + C_0/\sqrt{n}\right))\right\} \tag{EC.3}$$

where $\phi(\cdot)$ is the probability density function of the standard normal distribution. Then,

$$|y^n - y^{\star,n}| \leq \frac{\sigma C_0}{C_2(\beta,n)}, \text{ when } n > \frac{C_0^2}{\min\{\beta^2,(1-\beta)^2\}}, \tag{EC.4}$$

$$|w^n - w^{\star,n}| \le \frac{\sigma C_0}{C_2(\gamma,n)}, \ \textit{when } n > \frac{C_0^2}{\min\{\gamma^2,(1-\gamma)^2\}}, \tag{EC.5}$$

$$\mathbb{E}\left[y^{\star,n} - \sum_{t=1}^{L} D_t^n\right]^+ \le y^{\star,n}(1-\Theta^\star)^{Ln}, \ \textit{when } n > \left[\frac{\sigma\Phi^{-1}\left(\frac{p-c}{p-\alpha c+h}\right)}{(L-1)\lambda}\right]^2, \tag{EC.6}$$

$$\mathbb{E}\left[y^{n} - \sum_{t=1}^{L} D_t^n\right]^+ \le \left[y^{\star,n} + \frac{\sigma C_0}{C_2(\beta,n)}\right](1-\Theta)^{Ln}, \ \textit{when } y^{\star,n} + \frac{\sigma C_0}{C_2(\beta,n)} < L\lambda n, \tag{EC.7}$$

where $\Theta^\star = \frac{(L\lambda n - y^{\star,n})^2}{4L^2n^2(2\lambda^2+\sigma^2)}, \Theta = \frac{\left(L\lambda n - y^{\star,n} - \frac{\sigma C_0}{C_2(\beta,n)}\right)^2}{4L^2n^2(2\lambda^2+\sigma^2)}$.

Proof of Lemma EC.2. Since the proofs of (EC.5) and (EC.7) are similar to those of (EC.4) and (EC.6), respectively, we will only prove (EC.4) and (EC.6). Since $F_{D_t^n}(y^n) = F_{\tilde{D}_t^n}\left(\frac{y^n - n\lambda}{\sqrt{n}\sigma}\right) = \beta = \Phi(\frac{y^{\star,n}-n\lambda}{\sqrt{n}\sigma})$, by Lemma EC.1,

$$\left|\Phi\left(\frac{y^{\star,n}-n\lambda}{\sqrt{n}\sigma}\right) - \Phi\left(\frac{y^{n}-n\lambda}{\sqrt{n}\sigma}\right)\right| \le \frac{C_0}{\sqrt{n}}. \tag{EC.8}$$

Using the Taylor expansion, there exists some ξ between $\frac{y^{\star,n}-n\lambda}{\sqrt{n}\sigma}$ and $\frac{y^{n}-n\lambda}{\sqrt{n}\sigma}$,

$$\frac{|y^n - y^{\star,n}|}{\sqrt{n}\sigma} \le \frac{C_0}{\phi(\xi)\sqrt{n}} \le \frac{C_0}{\min\left\{\phi\left(\frac{y^{\star,n}-n\lambda}{\sqrt{n}\sigma}\right), \phi\left(\frac{y^{n}-n\lambda}{\sqrt{n}\sigma}\right)\right\}\sqrt{n}} = \frac{C_0}{C_2(\beta,n)\sqrt{n}},$$

where the last inequality follows from (EC.3) and $\Phi^{-1}(\beta - C_0/\sqrt{n}) \le \frac{y^n-n\lambda}{\sqrt{n}\sigma} \le \Phi^{-1}(\beta + C_0/\sqrt{n})$, by (EC.8). When $n > \frac{C_0^2}{\min\{\beta^2,(1-\beta)^2\}}$, we have $C_2(\beta,n) > 0$ and the above upper bound is finite.

Same as in the proof of Lemma EC.1, we can write the demand as $D_t^n = \sum_{i=1}^{n} X_{ti}$ and have

$$\mathbb{E}\left[y^{\star,n} - \sum_{t=1}^{L} D_t^n\right]^+ \le y^{\star,n}\mathbb{P}\left(\sum_{t=1}^{L} D_t^n \le y^{\star,n}\right) = y^{\star,n}\mathbb{P}\left(\sum_{t=1}^{L}\sum_{i=1}^{n}\left(\frac{y^{\star,n}}{Ln} - X_{ti}\right) \ge 0\right).$$

Similarly, we have a bound of the same form for $\mathbb{E}\left[y^n - \sum_{t=1}^{L} D_t^n\right]^+$. Now using Lemma 10 of Goldberg et al. (2016), the desired result follows. $\square$

Proof of Proposition 5. First let n be an integer such that all the results in Lemma EC.2 hold when $n > \mathrm{n}$. Throughout the proof we assume $n > \mathrm{n}$. Since the derivative of $V_t^n(\cdot)$ with respect to any inventory level is bounded by $\alpha c - h$,

$$\begin{aligned} V_0^n(\mathbf{I}_0^n) - V_0^{\star,n}(\mathbf{I}_0^n) &= \max_{(z,q)\in\Omega(\mathbf{I})} \{\mathbb{E}G(\mathbf{I}_0^n, z, q, D_1^n) + \alpha\mathbb{E}V_1^n(\mathbf{I}_1^n)\} - \{\mathbb{E}G(\mathbf{I}_0^n, z_0^{\star,n}, q_1^{\star,n}, D_1^n) + \alpha\mathbb{E}V_1^{\star,n}(\mathbf{I}_1^{\star,n})\} \\ &\le \max_{(z,q)\in\Omega(\mathbf{I})} \left\{\mathbb{E}G(\mathbf{I}_0^n, z, q, D_1^n) + \alpha(\alpha c - h)\mathbb{E}\left[I_0^n(L-1) - z + q - D_1^n - O_1^n\right]^+\right\} + \alpha V_1^n(\mathbf{0}) \end{aligned}$$

$$
\begin{aligned}
&- \{\mathbb{E}G(\mathbf{I}_0^n, z^{\star,n}, q^{\star,n}, D_1^n) + \alpha(\alpha c - h)\mathbb{E}I_1^{\star,n}(L-1)\} \\
&+ \{\alpha(\alpha c - h)\mathbb{E}I_1^{\star,n}(L-1) - \alpha\mathbb{E}V_1^{\star,n}(\mathbf{I}_1^{\star,n})\} \\
=& \left\{M(\mathbf{I}_0^n, z_0^{H,n}, q_1^{H,n}) - M(\mathbf{I}_0^n, z_0^{\star,n}, q_1^{\star,n})\right\} \\
&+ \{\alpha(\alpha c - h)\mathbb{E}I_1^{\star,n}(L-1) - \alpha\mathbb{E}V_1^{\star,n}(\mathbf{I}_1^{\star,n}) + \alpha V_1^{\star,n}(\mathbf{0})\} \\
&+ [\alpha V_1^n(\mathbf{0}) - \alpha V_1^{\star,n}(\mathbf{0})] \\
=& \text{(I)} + \text{(II)} + \text{(III)}
\end{aligned}
$$

where $M(\mathbf{I}_0^n, z, q) \triangleq \mathbb{E}G(\mathbf{I}_0^n, z, q, D_1^n) + \alpha(\alpha c - h)\mathbb{E}\left[I_0^n(L-1) - z + q - D_1^n - O_1^n\right]^+$. Next we bound each of the three terms on the right hand side.

Using the second order Taylor expansion, and by the fact that $(z_0^{H,n}, q_1^{H,n})$ is the maximizer of $M(\mathbf{I}_0^n, z, q)$, we have

$$
\begin{aligned}
\text{(I)} =& -\frac{\partial^2 M(\mathbf{I}_0^n, z, q)}{2\partial z^2}\,|_{(\mathrm{z},\mathrm{q})} \left(z_0^{\star,n} - z_0^{H,n}\right)^2 - \frac{\partial^2 M(\mathbf{I}_0^n, z, q)}{2\partial q^2}\,|_{(\mathrm{z},\mathrm{q})} \left(q_1^{\star,n} - q_1^{H,n}\right)^2 \\
&- \frac{\partial^2 M(\mathbf{I}_0^n, z, q)}{\partial z \partial q}\,|_{(\mathrm{z},\mathrm{q})} \left(q_1^{\star,n} - q_1^{H,n}\right)\left(z_1^{\star,n} - z_1^{H,n}\right)
\end{aligned}
$$

where z is between $z_0^{\star,n}$ and $z_0^{H,n}$, and q is between $q_1^{\star,n}$ and $q_1^{H,n}$. Let C_1 be a uniform bound of the probability density function of the scaled demand $\tilde{D}^n$ in Assumption 2. Note that

$$
\begin{aligned}
-\frac{\partial^2 M(\mathbf{I}_0^n, z, q)}{2\partial z^2}\,|_{(\mathrm{z},\mathrm{q})} &= \alpha(p - \alpha c + h) f_{D^n}(I_0^n(L-1) - \mathrm{z} + \mathrm{q}) + \alpha(\alpha c - h + \theta) f_{D^n}(I_0^n(L-1) - \mathrm{z}) \\
&\leq \alpha(p + \theta) C_1 n^{-1/2}, \\
-\frac{\partial^2 M(\mathbf{I}_0^n, z, q)}{2\partial q^2}\,|_{(\mathrm{z},\mathrm{q})} &= \alpha(p - \alpha c + h) f_{D^n}(I_0^n(L-1) - \mathrm{z} + \mathrm{q}) \leq \alpha(p - \alpha c + h) C_1 n^{-1/2}, \\
-\frac{\partial^2 M(\mathbf{I}_0^n, z, q)}{\partial z \partial q}\,|_{(\mathrm{z},\mathrm{q})} &= \alpha(p - \alpha c + h) f_{D^n}(I_0^n(L-1) - \mathrm{z} + \mathrm{q}) \leq \alpha(p - \alpha c + h) C_1 n^{-1/2}.
\end{aligned}
$$

Combined the above with Lemma EC.2, we obtain

$$
\text{(I)} \leq \frac{\alpha C_1 (\sigma C_0)^2}{\sqrt{n}} \left[\frac{p + \theta}{C_2^2(\beta, n)} + \frac{p - \alpha c + h}{C_2(\gamma, n)} \left(\frac{1}{C_2(\gamma, n)} + \frac{1}{C_2(\beta, n)}\right)\right].
$$

(II) is bounded by the loss incurred due to the expected expiry of inventory items in $\mathbf{I}_1^{\star,n}$ during periods $2, \cdots, L$, i.e.,

$$
\sum_{t=2}^{L} O_t^{\star,n} = \max\left\{0, I_0^n(1) - z_0^{\star,n} - D_1^n, \cdots, I_0^n(L-1) + q_1^{\star,n} - z_0^{\star,n} - D_{[L]}^n\right\}
$$

$$
\begin{aligned}
&- \max\{0, I_0^n(1) - z_0^{\star,n} - D_1^n\} \\
&\le \max\{0, I_0^n(L-1) + q_1^{\star,n} - D_1^n - D_2^n\} \\
&\le \max\{0, y^n \vee y^{\star,n} - D_1^n - D_2^n\}.
\end{aligned}
$$

Since each unit of expired inventory can bring at most a loss of $\alpha c - h + \theta$, we have

$$
\begin{aligned}
(\mathrm{II}) &\le (\alpha c - h + \theta)\mathbb{E}\sum_{t=2}^{L} O_t^{\star,n} \le (\alpha c - h + \theta)\mathbb{E}\max\{0, y^n \vee y^{\star,n} - D_1^n - D_2^n\} \\
&\le (\alpha c - h + \theta)\left[y^{\star,n} + \frac{\sigma C_0}{C_2(\beta,n)}\right](1-\Theta)^{2n}
\end{aligned}
$$

where the last inequality follows from Lemma EC.2. Following the same procedure as in the proof of Proposition 2, we can obtain

$$
\begin{aligned}
(\mathrm{III}) &\le \frac{\alpha - \alpha^{T+1}}{1-\alpha}\left\{\frac{C_1}{2\sigma\sqrt{n}}(p - \alpha c + h)(y^n - y^{\star,n})^2 + (\alpha c - h + \theta)\mathbb{E}\left[y^{\star,n} - D_{[L]}^n\right]^+\right\} \\
&\le \frac{\alpha - \alpha^{T+1}}{1-\alpha}\left\{\frac{C_1}{2\sigma\sqrt{n}}(p - \alpha c + h)\left(\frac{\sigma C_0}{C_2(\beta,n)}\right)^2 + (\alpha c - h + \theta)y^{\star,n}(1-\Theta^\star)^{Ln}\right\}
\end{aligned}
$$

where the last inequality follows from Lemma EC.2. The proof is complete by observing that the sum of the three bounds is of the order $O(\sqrt{n})$ (note that $C_2(\cdot,n)$ is bounded away from 0 when $n \to \infty$). $\square$

Proof of Proposition 6. Since

$$
V^n(\mathbf{0}) - V^{f,n}(\mathbf{0}) = \left[V^n(\mathbf{0}) - n\bar{V}(\mathbf{0})\right] + \left[n\bar{V}(\mathbf{0}) - V^{f,n}(\mathbf{0})\right],
$$

we only need to show that the two terms in [ ] are of the order $O(\sqrt{n})$. We first note that $\frac{\partial V^n(\mathbf{I})}{\partial I(\ell)} \le \alpha p + K$ as any extra unit of inventory can avoid at most an order at the cost K now and delayed sales at αp. Since $V^n(\mathbf{I})$ and $n\bar{V}(\mathbf{I})$ follow a recursive equation of the same form, we obtain

$$
\begin{aligned}
V^n(\mathbf{0}) &= \max_{z,q\in\Omega(\mathbf{I})}\{\mathbb{E}G(\mathbf{0},z,q,D_1^n) + \alpha\mathbb{E}V^n(\Gamma(\mathbf{0},z,q,D_1^n))\} \\
&\le \max_{z,q\in\Omega(\mathbf{I})}\{\mathbb{E}G(\mathbf{0},z,q,n\lambda) + \alpha\mathbb{E}V^n(\Gamma(\mathbf{0},z,q,n\lambda)) + \alpha(p+\theta+h+K)\mathbb{E}|D_1^n - n\lambda|\} \\
&\le \dots \le \frac{\alpha}{1-\alpha}(p+\theta+h+K)\mathbb{E}|D_1^n - n\lambda| + n\bar{V}(\mathbf{0}). \qquad \text{(EC.9)}
\end{aligned}
$$

Thus, $[V^n(\mathbf{0}) - n\bar{V}(\mathbf{0})] = O(\sqrt{n})$. Next we identify an upper bound of the second term by a sample-path argument. Note that the difference between $n\bar{V}(\mathbf{0})$ and $V^{f,n}(\mathbf{0})$ is caused by the deviation of demand from its mean $n\lambda$. Furthermore, any demand above the mean $[D_t^n - n\lambda]^+$ is either lost, which has no effect on the difference, or satisfied, which will bring more profit even though it may reduce sales in the future due to discounting. Any demand below its mean $[n\lambda - D_t^n]^+$ will incur a revenue loss of αp and a holding cost bounded by $h + \alpha h + ... + \alpha^{t^*-1} h = \frac{1-\alpha^{t^*}}{1-\alpha} h$, plus a potential disposal cost $\alpha\theta$, compared with $n\bar{V}(\mathbf{0})$. Thus,

$$\begin{aligned} n\bar{V}(\mathbf{0}) - V^{f,n}(\mathbf{0}) &\leq \sum_{t=0}^{\infty} \alpha^t \left[\alpha(p+\theta) + \frac{1-\alpha^{t^*}}{1-\alpha} h \right] \mathbb{E}\,[n\lambda - D_t^n]^+ \\ &= \left[\frac{\alpha}{1-\alpha}(p+\theta) + \frac{1-\alpha^{t^*}}{(1-\alpha)^2} h \right] \mathbb{E}\,[n\lambda - D_t^n]^+ , \end{aligned}$$

and $[n\bar{V}(\mathbf{0}) - V^{f,n}(\mathbf{0})] = O(\sqrt{n})$. □

Proof of Proposition 7. By a similar argument to Theorem 1 of Chen et al. (2014), we can show that, for any given $\mathbf{I}$,

$$0 \leq \frac{\partial z^*(\mathbf{I})}{\partial I_\ell} \leq 1, \quad -1 \leq \frac{\partial q^*(\mathbf{I})}{\partial I_\ell} \leq 0, \quad 1 \leq \ell \leq L-1. \tag{EC.10}$$

That is, with one fewer unit of inventory as we move up the path, at most one fewer unit is cleared and at most one more unit is ordered. Suppose the path starts from state $\mathbf{I} > \mathbf{0}$ at the bottom of the tree.

If $z^*(\mathbf{I}) = 0$, then along the path to $\mathbf{0}$, the optimal clearance quantity is 0 and the optimal order quantity increases at most by one unit. Now suppose that $z^*(\mathbf{I}) > 1$. We want to show that if $\mathbf{I}'$ is the successor of $\mathbf{I}$, then $z^*(\mathbf{I}') = z_0 - 1$, and $q^*(\mathbf{I}') = q^*(\mathbf{I})$.

Let $q^*(\mathbf{I}, z) = \arg\max_{q \geq 0} \{v(\mathbf{I}, z, q)\}$ (if multiple maximizers exist, pick the largest one). Then,

$$\max_{z,q} \{v(\mathbf{I}, z, q)\} = \max_{z} \{v(\mathbf{I}, z, q^*(\mathbf{I}, z))\}.$$

If $z^*(\mathbf{I}) = I(L-1)$, then it is optimal to clear the entire inventory and order the same quantity at any of the successors. Otherwise, at its immediate predecessor $\mathbf{I}'$,

$$\begin{aligned} &v(\mathbf{I}', z^*(\mathbf{I}) - 1, q^*(\mathbf{I}', z^*(\mathbf{I}) - 1)) = v(\mathbf{I}, z^*(\mathbf{I}), q^*(\mathbf{I}, z^*(\mathbf{I}))) - s \\ &\qquad > v(\mathbf{I}, z^*(\mathbf{I}) + 1, q^*(\mathbf{I}, z^*(\mathbf{I}) + 1)) - s = v(\mathbf{I}', z^*(\mathbf{I}), q^*(\mathbf{I}', z^*(\mathbf{I}))), \end{aligned}$$

which implies $z^*(\mathbf{I}') = z^*(\mathbf{I}) - 1$ and hence $q^*(\mathbf{I}') = q^*(\mathbf{I})$ as the inventory profiles after clearance are the same under the two states. □

EC.2. Extended numerical experiments

Tables EC.1 and EC.2 summarise the lower bounds of n to guarantee certain levels of performance gap for polices f and H. Tables EC.3-EC.9 provide numerical results for various polices discussed in this paper under various scenarios and initial inventory levels.

Performance gap	c	Constant	Uniform	Exponential	Hyper-exp
5%	2	59	79	91	184
	4	32	43	64	104
	6	25	33	35	50
3%	2	133	177	251	471
	4	72	112	143	253
	6	44	52	68	74
2%	2	256	334	491	1001
	4	149	191	253	508
	6	69	101	88	137

Table EC.1 Lower bounds on n to guarantee the desired performance gap under policy f for $L = 4$ when $\mathbf{I}_0^n = \mathbf{0}$.

Performance gap	c	L	Constant	Uniform	Exponential	Hyper-exp
3%	2	2	6	10	15	38
	4	2	6	9	14	35
	6	2	6	8	14	33
	2	4	3	4	5	15
	4	4	2	2	5	12
	6	4	2	2	4	11
2%	2	2	8	11	18	42
	4	2	9	10	17	40
	6	2	7	11	16	37
	2	4	3	4	7	19
	4	4	2	4	6	15
	6	4	2	2	5	13
1%	2	2	10	14	23	52
	4	2	10	14	20	49
	6	2	8	13	19	46
	2	4	5	6	9	23
	4	4	3	5	9	21
	6	4	2	5	8	19

Table EC.2 Lower bounds on n to guarantee the desired performance gap under policy H when $\mathbf{I}_0^n = \mathbf{0}$.

			Constant		Uniform		Exponential		Hyper-exp	
n	c	s	mean	max	mean	max	mean	max	mean	max
10	2	1.0	9.71	10.43	11.34	12.30	13.46	14.77	17.55	19.95
	4	2.0	7.41	8.17	8.76	9.73	10.28	11.55	12.66	14.67
	6	3.0	4.54	5.27	5.37	6.30	6.14	7.29	6.53	8.05
	2	1.5	9.79	10.51	11.42	12.40	13.60	14.89	18.27	20.20
	4	3.0	7.55	8.29	8.91	9.87	10.50	11.72	13.24	15.07
	6	4.5	4.66	5.47	5.51	6.51	6.32	7.65	6.94	9.14
20	2	1.0	6.81	7.38	8.12	8.66	9.35	10.48	14.92	14.97
	4	2.0	5.35	5.80	6.24	6.82	7.55	8.18	9.66	11.27
	6	3.0	3.28	3.72	3.85	4.39	4.45	5.18	5.66	6.58
	2	1.5	6.86	7.45	8.14	8.73	9.44	10.56	15.00	15.10
	4	3.0	5.41	5.89	6.31	6.92	7.77	8.31	10.13	11.45
	6	4.5	3.34	3.85	3.91	4.53	4.57	5.39	5.71	7.15
30	2	1.0	4.98	6.03	7.32	7.05	8.01	8.55	9.42	12.38
	4	2.0	4.19	4.73	4.76	5.54	6.09	6.67	8.60	9.39
	6	3.0	2.63	3.03	3.12	3.56	3.88	4.23	4.45	5.59
	2	1.5	5.13	6.08	7.46	7.11	8.02	8.62	9.42	12.49
	4	3.0	4.21	4.80	4.81	5.63	6.36	6.78	8.78	9.53
	6	4.5	2.69	3.13	3.23	3.67	3.92	4.40	4.53	6.00
40	2	1.0	5.63	5.23	5.03	6.10	4.15	7.40	4.92	10.77
	4	2.0	3.36	4.10	5.08	4.79	5.41	5.78	3.73	8.19
	6	3.0	2.17	2.63	2.65	3.07	3.49	3.66	4.27	4.92
	2	1.5	5.63	5.27	5.03	6.15	4.15	7.47	4.92	10.86
	4	3.0	3.42	4.16	5.09	4.86	5.41	5.87	4.27	8.32
	6	4.5	2.25	2.71	2.72	3.17	3.49	3.80	4.85	5.25

Table EC.3 **Mean and maximum percentage gap under policy f for $L=4$. “mean” and “max” are with respect to different initial inventory levels, the same below.**

n	c	s	Constant	Uniform	Exponential	Hyper-exp
10	2	1.0	10.43	12.30	14.77	19.95
	4	2.0	8.17	9.73	11.55	14.67
	6	3.0	5.27	6.30	7.29	8.05
	2	1.5	10.51	12.40	14.89	20.20
	4	3.0	8.29	9.87	11.72	15.00
	6	4.5	5.38	6.41	7.40	8.16
20	2	1.0	7.38	8.66	10.48	14.97
	4	2.0	5.80	6.82	8.18	11.27
	6	3.0	3.72	4.39	5.18	6.58
	2	1.5	7.45	8.73	10.56	15.10
	4	3.0	5.89	6.92	8.31	11.45
	6	4.5	3.79	4.47	5.27	6.67
30	2	1.0	6.03	7.05	8.55	12.38
	4	2.0	4.73	5.54	6.67	9.39
	6	3.0	3.03	3.56	4.23	5.59
	2	1.5	6.08	7.11	8.62	12.49
	4	3.0	4.80	5.63	6.78	9.53
	6	4.5	3.09	3.63	4.31	5.67
40	2	1.0	5.23	6.10	7.40	10.77
	4	2.0	4.10	4.79	5.78	8.19
	6	3.0	2.63	3.07	3.66	4.92
	2	1.5	5.27	6.15	7.47	10.86
	4	3.0	4.16	4.86	5.87	8.32
	6	4.5	2.68	3.13	3.73	5.00

Table EC.4 Percentage gap of policy f for $L = 4$ when $\mathbf{I}_0^n = \mathbf{0}$.

			$L=2$				$L=4$			
n	c	s	Const	Uniform	Exp	Hyp-exp	Const	Uniform	Exp	Hyp-exp
10	2	1.0	0.13	0.37	0.89	3.79	0.05	0.06	0.09	0.31
	4	2.0	0.11	0.61	1.42	7.17	0.10	0.18	0.23	0.51
	6	3.0	0.33	0.82	1.81	9.34	0.32	0.38	0.48	0.93
	2	1.5	0.18	0.57	1.49	6.45	0.01	0.01	0.01	0.24
	4	3.0	0.14	0.73	1.97	9.74	0.00	0.02	0.03	0.22
	6	4.5	0.18	0.73	2.04	10.74	0.04	0.05	0.06	0.17
20	2	1.0	0.07	0.07	0.20	1.46	0.04	0.04	0.05	0.11
	4	2.0	0.08	0.13	0.30	2.19	0.08	0.11	0.15	0.28
	6	3.0	0.19	0.26	0.42	2.41	0.19	0.25	0.32	0.57
	2	1.5	0.03	0.06	0.27	2.33	0.01	0.00	0.01	0.00
	4	3.0	0.01	0.05	0.27	2.87	0.01	0.01	0.02	0.03
	6	4.5	0.02	0.05	0.21	2.54	0.02	0.03	0.04	0.07
30	2	1.0	0.02	0.03	0.07	0.68	0.02	0.03	0.04	0.07
	4	2.0	0.09	0.09	0.13	0.90	0.09	0.09	0.12	0.21
	6	3.0	0.16	0.20	0.25	0.95	0.16	0.20	0.25	0.43
	2	1.5	0.00	0.01	0.05	1.01	0.00	0.00	0.00	0.00
	4	3.0	0.01	0.01	0.04	1.05	0.01	0.01	0.01	0.02
	6	4.5	0.02	0.03	0.04	0.77	0.02	0.03	0.03	0.05

Table EC.5 Percentage gap of policy H when $\mathbf{I}_0^n = \mathbf{0}$.

				Constant		Uniform		Exponential		Hyper-exponential	
L	n	c	s	mean	max	mean	max	mean	max	mean	max
2	10	2	1.0	0.12	0.13	0.35	0.37	0.85	0.89	3.69	3.79
		4	2.0	0.10	0.11	0.56	0.61	1.32	1.42	6.74	7.17
		6	3.0	0.30	0.33	0.75	0.82	1.61	1.81	8.24	9.34
		2	1.5	0.17	0.18	0.55	0.57	1.44	1.49	6.28	6.45
		4	3.0	0.13	0.14	0.68	0.73	1.85	1.97	9.18	9.74
		6	4.5	0.17	0.18	0.65	0.73	1.82	2.04	9.55	10.74
	20	2	1.0	0.06	0.07	0.07	0.07	0.19	0.20	1.41	1.46
		4	2.0	0.08	0.08	0.12	0.13	0.28	0.30	2.05	2.19
		6	3.0	0.18	0.19	0.24	0.26	0.38	0.42	2.15	2.41
		2	1.5	0.03	0.03	0.05	0.06	0.25	0.27	2.27	2.33
		4	3.0	0.01	0.01	0.04	0.05	0.25	0.27	2.72	2.87
		6	4.5	0.02	0.02	0.05	0.05	0.19	0.21	2.28	2.54
	30	2	1.0	0.02	0.02	0.03	0.03	0.06	0.07	0.65	0.68
		4	2.0	0.09	0.09	0.09	0.09	0.13	0.13	0.84	0.90
		6	3.0	0.15	0.16	0.19	0.20	0.24	0.25	0.86	0.95
		2	1.5	0.00	0.00	0.01	0.01	0.05	0.05	0.98	1.01
		4	3.0	0.01	0.01	0.01	0.01	0.04	0.04	0.99	1.05
		6	4.5	0.02	0.02	0.03	0.03	0.04	0.04	0.69	0.77
4	10	2	1.0	0.05	0.05	0.06	0.08	0.09	0.15	0.32	0.74
		4	2.0	0.09	0.10	0.17	0.18	0.21	0.26	0.48	0.92
		6	3.0	0.28	0.32	0.34	0.38	0.42	0.48	0.80	1.08
		2	1.5	0.01	0.04	0.01	0.09	0.02	0.25	0.35	1.29
		4	3.0	0.00	0.03	0.02	0.11	0.03	0.29	0.28	1.48
		6	4.5	0.04	0.06	0.05	0.11	0.06	0.27	0.20	1.46
	20	2	1.0	0.04	0.04	0.04	0.04	0.05	0.06	0.12	0.22
		4	2.0	0.08	0.08	0.10	0.11	0.14	0.15	0.25	0.34
		6	3.0	0.17	0.19	0.22	0.25	0.28	0.32	0.50	0.57
		2	1.5	0.01	0.01	0.00	0.01	0.01	0.04	0.00	0.01
		4	3.0	0.01	0.01	0.01	0.02	0.02	0.04	0.04	0.40
		6	4.5	0.02	0.02	0.03	0.03	0.04	0.05	0.07	0.33
	30	2	1.0	0.02	0.02	0.03	0.03	0.04	0.04	0.07	0.08
		4	2.0	0.08	0.09	0.08	0.09	0.12	0.12	0.23	0.21
		6	3.0	0.13	0.16	0.17	0.20	0.24	0.25	0.39	0.43
		2	1.5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01
		4	3.0	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.02
		6	4.5	0.02	0.02	0.02	0.03	0.03	0.03	0.04	0.11

Table EC.6 **Mean and maximum percentage gap under policy H.**

				Constant		Uniform		Exponential		Hyper-exp	
L	n	c	s	mean	max	mean	max	mean	max	mean	max
		2	1.0	0.09	0.10	0.13	0.14	0.28	0.31	0.92	1.01
		4	2.0	0.26	0.27	0.36	0.39	0.89	0.97	4.42	4.82
	10	6	3.0	0.29	0.31	0.64	0.71	1.79	2.00	10.96	12.53
		2	1.5	0.15	0.16	0.34	0.35	0.91	0.93	3.74	3.82
		4	3.0	0.26	0.28	0.49	0.53	1.46	1.55	7.18	7.54
		6	4.5	0.17	0.18	0.56	0.63	2.00	2.22	12.41	13.73
		2	1.0	0.03	0.03	0.03	0.03	0.06	0.06	0.32	0.35
		4	2.0	0.09	0.10	0.10	0.10	0.18	0.18	1.22	1.33
2	20	6	3.0	0.22	0.24	0.23	0.24	0.38	0.42	2.67	2.99
		2	1.5	0.01	0.01	0.03	0.03	0.13	0.14	1.25	1.28
		4	3.0	0.02	0.02	0.03	0.03	0.16	0.17	1.97	2.06
		6	4.5	0.05	0.06	0.04	0.04	0.19	0.21	2.81	3.08
		2	1.0	0.02	0.02	0.03	0.03	0.03	0.03	0.13	0.15
		4	2.0	0.08	0.08	0.08	0.08	0.10	0.11	0.49	0.53
	30	6	3.0	0.16	0.17	0.18	0.19	0.24	0.25	1.02	1.13
		2	1.5	0.00	0.00	0.01	0.01	0.02	0.03	0.49	0.50
		4	3.0	0.02	0.02	0.01	0.01	0.03	0.03	0.67	0.71
		6	4.5	0.03	0.03	0.02	0.02	0.04	0.04	0.85	0.93
		2	1.0	0.08	0.08	0.08	0.09	0.14	0.20	0.38	0.78
		4	2.0	0.13	0.14	0.15	0.16	0.20	0.25	0.37	0.80
	10	6	3.0	0.31	0.34	0.32	0.36	0.44	0.50	1.10	1.35
		2	1.5	0.05	0.07	0.04	0.12	0.10	0.32	0.49	1.44
		4	3.0	0.02	0.05	0.03	0.11	0.05	0.30	0.24	1.45
		6	4.5	0.08	0.09	0.04	0.10	0.07	0.28	0.42	1.62
		2	1.0	0.03	0.03	0.04	0.04	0.06	0.07	0.25	0.34
		4	2.0	0.09	0.09	0.10	0.11	0.13	0.14	0.25	0.32
4	20	6	3.0	0.21	0.23	0.21	0.24	0.28	0.32	0.55	0.65
		2	1.5	0.00	0.01	0.02	0.02	0.04	0.07	0.22	0.51
		4	3.0	0.01	0.01	0.02	0.02	0.02	0.05	0.07	0.43
		6	4.5	0.05	0.05	0.03	0.03	0.04	0.05	0.11	0.36
		2	1.0	0.02	0.02	0.03	0.03	0.04	0.05	0.15	0.16
		4	2.0	0.07	0.08	0.08	0.08	0.11	0.11	0.19	0.20
	30	6	3.0	0.15	0.17	0.17	0.19	0.24	0.25	0.46	0.47
		2	1.5	0.00	0.00	0.01	0.01	0.02	0.02	0.08	0.22
		4	3.0	0.02	0.02	0.01	0.01	0.02	0.02	0.04	0.12
		6	4.5	0.03	0.03	0.02	0.02	0.03	0.03	0.06	0.12

Table EC.7 **The mean and maximum percentage gap for policy $\star$.**

			Constant		Uniform		Exponential		Hyper-exp	
n	c	θ	mean	max	mean	max	mean	max	mean	max
5	2	1	0.33	1.45	0.67	2.19	0.80	2.97	1.48	5.89
	4	1	3.06	4.14	4.19	5.44	5.79	7.27	12.15	14.08
	6	1	1.08	1.58	1.30	1.82	2.05	2.71	5.39	7.52
	2	4	0.36	4.10	0.50	5.52	0.60	7.41	1.25	13.73
	4	4	1.75	3.70	2.54	4.88	3.86	7.15	8.73	16.10
	6	4	3.12	5.28	3.74	5.82	5.74	8.75	13.30	19.70
10	2	1	0.47	1.20	0.52	1.45	0.62	1.94	0.96	3.74
	4	1	1.83	2.47	2.79	3.66	3.64	4.78	7.45	9.14
	6	1	0.84	1.45	0.79	1.12	1.03	1.46	2.66	3.90
	2	4	0.50	3.11	0.47	3.69	0.50	4.88	0.81	9.08
	4	4	0.90	2.13	1.44	3.01	1.98	4.06	4.92	8.92
	6	4	2.05	3.42	2.18	3.64	2.99	4.97	7.04	10.83
20	2	1	0.28	0.73	0.44	1.00	0.50	1.30	0.68	2.36
	4	1	1.59	1.99	1.93	2.44	2.46	3.19	4.48	5.86
	6	1	0.44	2.00	0.53	1.71	0.64	1.34	1.23	1.83
	2	4	0.39	2.00	0.50	2.55	0.50	3.28	0.61	5.81
	4	4	0.79	1.98	0.97	2.01	1.21	2.61	2.46	4.91
	6	4	1.07	2.12	1.38	2.41	1.76	3.12	3.53	5.91
30	2	1	0.31	0.64	0.39	0.81	0.45	1.02	0.47	1.58
	4	1	1.14	1.41	1.60	1.98	1.97	2.52	3.78	4.59
	6	1	0.41	2.41	0.42	2.12	0.55	1.72	0.96	1.31
	2	4	0.45	1.68	0.48	2.05	0.52	2.59	0.57	5.32
	4	4	0.56	2.35	0.79	2.10	1.01	2.05	1.89	3.65
	6	4	0.92	2.52	1.04	2.25	1.35	2.45	2.58	4.36

Table EC.8 Mean and maximum percentage gap for policy RH under the LIFO issuing rule for $L = 4$.

Policy	L	c	s	Unif(25,35)	Unif(20,40)	Norm(30,5^2)	Norm(30,10^2)	Exp(1/30)
f	2	2	1.0	3.02	6.84	4.87	10.12	20.72
		4	2.0	2.41	5.65	3.75	7.64	8.55
		6	3.0	1.54	3.81	2.30	4.52	0.86
		2	1.5	3.04	6.88	4.92	10.34	26.75
		4	3.0	2.45	5.74	3.81	7.94	18.17
		6	4.5	1.57	3.88	2.35	4.81	11.90
	4	2	1.0	3.02	6.84	4.87	10.69	27.48
		4	2.0	2.41	5.65	3.75	8.52	19.06
		6	3.0	1.54	3.81	2.30	5.57	8.98
		2	1.5	3.04	6.88	4.92	10.78	28.89
		4	3.0	2.45	5.74	3.82	8.65	21.03
		6	4.5	1.57	3.88	2.35	5.67	10.77
H	2	2	1.0	0.00	0.01	0.03	0.22	8.57
		4	2.0	0.04	0.06	0.06	0.32	19.47
		6	3.0	0.11	0.18	0.13	0.61	35.05
		2	1.5	0.00	0.00	0.00	0.33	15.10
		4	3.0	0.01	0.01	0.01	0.39	27.05
		6	4.5	0.02	0.02	0.01	0.54	40.69
	4	2	1.0	0.00	0.01	0.03	0.05	1.73
		4	2.0	0.04	0.06	0.06	0.13	2.33
		6	3.0	0.11	0.18	0.13	0.32	2.81
		2	1.5	0.00	0.00	0.00	0.01	3.07
		4	3.0	0.01	0.01	0.01	0.02	3.67
		6	4.5	0.02	0.02	0.01	0.04	3.32

Table EC.9 Percentage gap for policies f and H (with general demand distribution).