Perishable Inventory Systems

Qing Li and Peiwen Yu

Abstract Many chemical, food and pharmaceutical products are perishable and effective management of their inventory is important to customer service, the bottom line, and the environment. This chapter provides a review of recent literature on perishable inventory systems. The chapter covers models with one class of demand and one location, multiple classes of demand and one location, as well as one class of demand and multiple locations. The focus of the review is on structural properties of the optimal policies and intuitive heuristic policies. The chapter concludes with a discussion on empirical studies and several ideas for future research.

1 Introduction

Perishable inventory management is one of the most researched areas in Operations Management/Operations Research. There have been at least four reviews of the literature: Prastacos (1982), Nahmias (1982), Karaesmen et al. (2011) and Nahmias (2011). The review by Prastacos (1982) focuses on blood products, and the other three are more general. Since the last two reviews in 2011, the area has been experiencing a resurgence of interest. In this review, we focus on recent results that are not discussed in the earlier reviews.

In our view, there are three reasons for the renewed interest in this research area. First, driven by market development, there has been much interest in e-commerce, health care operations and sustainable operations (in particular waste reduction), and many problems in these areas are related to perishable inventory systems. Second, although perishable inventory systems are known to be difficult to analyze, this research area has received a boost from the successful application of tools such

Qing Li
Hong Kong University of Science and Technology, e-mail: imqli@ust.hk

Peiwen Yu
Chongqing University, e-mail: ypw@cqu.edu.cn

as multimodularity and $L^\natural$-convexity in large-scale dynamic programs. Third, perishable inventory systems, which are typically computationally challenging, have proved to be great test beds for new ideas in algorithm design, another area that has been experiencing increased research recently.

In this chapter, we focus on research on the structure of optimal inventory decisions. Research on perishable inventory management with pricing decisions is reviewed in Chapter 14 (e.g., Chen et al. 2014, Hu et al. 2016). Research on approximation and learning algorithms for perishable inventory systems is reviewed in Chapters 12, 15 and 16 (e.g., Chao et al. 2015, Chao et al. 2018). The remainder of this review is organized as follows. Section 2 reviews the results for joint replenishment and clearance sales when there is only one class of demand and inventories are depleted either on a first-in-first-out (FIFO) basis or last-in-first-out (LIFO) basis. Section 3 discusses how the results under the FIFO rule can be generalized to allow multiple classes of demand. Section 4 reviews the results on perishable inventory systems that involve multiple locations. Section 5 reviews models with endogenous product lifetimes. Section 6 discusses the results in empirical research. The review concludes with a discussion of potential future research directions in Section 7.

2 Models with One Class of Demand and Location

Consider a firm that sells perishable products with an n-period lifetime. The firm purchases products at unit cost c. The products can be sold either at a regular price, r, or a clearance sale price, s. With a regular price, the demand in a period is random. Let D represent the regular demand. Unmet demand is lost. Demand under the clearance sales price scenario is abundant, and thus a firm can control how many items it sells for a given price. Without loss of generality, it is assumed that items have zero value once they expire. Items that expire incur outdating cost θ per unit and are removed from the shelf and disposed of. Items that are carried over to the next period incur holding cost h per unit. Profits received in future periods are discounted by discount factor α. Clearance sales are a common strategy to reduce a mismatch between supply and demand for perishable goods. To effectively use clearance sales to reduce a mismatch, it is important that retailers choose the right timing and sales depth and coordinate such sales with replenishment decisions.

Depending on the sequence by which items in different age groups are used to fulfill the regular demand, the analysis of these models requires different analytical tools and leads to different optimal replenishment and clearance sale policies.

Li and Yu (2014) study a problem in which a firm controls inventory issuance. This problem is particularly relevant to blood banks and e-commerce platforms that sell grocery products. The problem's sequence of events is as follows. At the beginning of each period, the firm's initial state is $\mathbf{y} = (y_1, y_2, ..., y_{n-1})$, which represents the inventory level after regular demand is fulfilled but before any clearance sales. Here, y_i represents inventory with a remaining lifetime of i periods. The firm must decide on an order quantity q of new items and the amount of inventory that will be carried over to the next period, denoted by $\mathbf{z} = (z_1, z_2, ..., z_{n-1})$. Then, after regular demand is realized, the firm decides on an issuing policy to meet the demand. At the end of each period, the items that expired in that period are removed and disposed of. Let d_i denote the amount of regular demand that is met by the inventories with a remaining lifetime of i periods. Let $\mathbf{d} = (d_1, d_2, ..., d_n)$ and $O(D) = \{\mathbf{d} : \sum_{i=1}^{n} d_i \le D, 0 \le d_n \le q, 0 \le d_i \le z_i \text{ for } 1 \le i \le n-1\}$.

The dynamic programming formulation is as follows:

$$\pi_t(\mathbf{y}) = s \sum_{i=1}^{n-1} y_i + \max_{\mathbf{0} \le \mathbf{z} \le \mathbf{y}, q \ge 0} u_t(\mathbf{z}, q),$$

where

$$u_t(\mathbf{z}, q) = -(s+h) \sum_{i=1}^{n-1} z_i - \alpha c q + \alpha \mathbb{E} \max_{\mathbf{d} \in O(D)} \Bigg\{ r \sum_{i=1}^{n} d_i - \theta(z_1 - d_1)$$
$$+ \pi_{t+1}(z_2 - d_2, ..., z_{n-1} - d_{n-1}, q - d_n) \Bigg\},$$

and $\pi_{T+1}(\mathbf{y}) = s \sum_{i=1}^{n-1} y_i$. In the formulation above, $s \sum_{i=1}^{n-1} (y_i - z_i)$ represents the revenue from clearance sales. Since $\sum_{i=1}^{n-1} z_i$ units of inventory are carried over to the next period, they incur total holding cost of $h \sum_{i=1}^{n-1} z_i$. Given the demand fulfillment vector $\mathbf{d}$, the revenue from regular sales is $r \sum_{i=1}^{n} d_i$, the amount of inventory that will be outdated is equal to $z_1 - d_1$, and the initial state in the next period is $(z_2 - d_2, ..., z_{n-1} - d_{n-1}, q - d_n)$. Let $(\mathbf{z}_t^*, q_t^*)$ denote the optimal solution.

This problem is unique in two ways. First, it has a multidimensional state space. Second, the state and decision variables are economic substitutes. The standard approach used to prove the preservation of structural properties in dynamic programs can not be applied directly to substitutes. Li and Yu (2014) use the concept of multimodularity to establish the structural properties for this problem, as follows.

An n-dimensional set $X \subseteq \mathbb{R}^n$ is called a *multimodular set* if there exist $\mathbf{a}_i \in \mathbb{R}^n$ and $b_i \in \mathbb{R}$ such that $X = \{\mathbf{x} \in \mathbb{R}^n | \mathbf{a}_i \cdot \mathbf{x} \ge b_i, i = 1, 2, .., m\}$ and $\mathbf{a}_i$ has the form $(0, ..., 0, 1, ..., 1, 0, ..., 0)$; that is, the nonzero components of $\mathbf{a}_i$ are either consecutive 1s or consecutive -1s. Let $\mathbf{x} = (x_1, ..., x_n)$. An n-dimensional function $f(\mathbf{x})$ defined on a multimodular set $X \subseteq \mathbb{R}^n$ is *multimodular* (*anti-multimodular*) if $f(x_1 - z, x_2 - x_1, ..., x_n - x_{n-1})$ is submodular (supermodular) in $(\mathbf{x}, z)$.

Anti-multimodularity implies decreasing difference, and it thus can be used to analyze models with substitutable variables. Anti-multimodular functions have some useful properties. A continuous anti-multimodular function is jointly concave. If $g(x)$ is a one-dimensional concave function, then $f(\mathbf{x}) = g(x_1 + x_2 + ... + x_n)$ is anti-multimodular in $\mathbf{x}$. The sum of anti-multimodular functions is still anti-multimodular; that is, if $f(\mathbf{x})$ and $g(\mathbf{x})$ are anti-multimodular, then $f(\mathbf{x}) + g(\mathbf{x})$ is anti-multimodular, and if $f(\mathbf{x}, d)$ is anti-multimodular in $\mathbf{x}$ for any given d and D is a random variable, then $\mathbb{E} f(\mathbf{x}, D)$ is anti-multimodular in $\mathbf{x}$.

Anti-multimodularity is preserved under maximization and the maximizer of an anti-multimodular function has monotonicity properties with bounded sensitivity. This property makes anti-multimodularity a useful tool for identifying structural properties in dynamic programs. To be more specific, if $f(\mathbf{x}, y)$ is an $n+1$ dimensional anti-multimodular function and $\{(\mathbf{x}, y) | \mathbf{x} \in X, y \in Y(\mathbf{x})\}$ is a multimodular set, then

$$g(\mathbf{x}) = \max_{y \in Y(\mathbf{x})} f(\mathbf{x}, y)$$

is anti-multimodular in $\mathbf{x}$. The optimal solution, denoted by $\bar{y}(\mathbf{x})$, satisfies the following inequalities:

$$-1 \leq \Delta_{x_n} \bar{y} \leq \Delta_{x_{n-1}} \bar{y} \leq \ldots \leq \Delta_{x_1} \bar{y} \leq 0.$$

Throughout the remainder of the chapter, the notation $\Delta_{x_i} J(\mathbf{x})$ is used to represent $(J(\mathbf{x} + \delta \mathbf{e}_i) - J(\mathbf{x}))/\delta$, where $\mathbf{e}_i$ is a vector with 1 in its i^{th} component and 0 in all the other components and δ is a small positive number. When $J(\mathbf{x})$ is differentiable, then $\Delta_{x_i} J(\mathbf{x})$ means that $\partial J(\mathbf{x})/\partial x_i$.

Multimodularity is closely related to $L^\natural$-convexity, a stronger notion of complementarity than submodularity (Lu and Song 2005, Zipkin 2008). Multimodular functions and $L^\natural$-convex functions are related through unimodular coordinate transformations. For the models in which the state and decision variables are economic substitutes, to use $L^\natural$-convexity to show structural properties, one must first transform the original variables into complementary variables, then show structural properties with respect to the new variables through showing $L^\natural$-convexity, and finally transform the properties back to those with respect to the original variables. In spite of their mathematical equivalence, they represent two conceptually different paths to the same destination. Whereas one tackles the problems directly, the other takes a detour by transforming them into problems of complementarity.

Using multimodularity, Li and Yu (2014) show that the optimal inventory issuing rule is FIFO, and that both the maximal profit function, π_t, and the objective function, u_t, are anti-multimodular. Note that without the option of clearance sales (e.g., Fries 1975, Nandakumar and Morton 1993), FIFO may not be the optimal issuing rule and multimodularity may not hold. For a given inventory of a certain age, the optimal policy on clearance sales has a clear-down-to structure; that is, there is a clear-down-to level such that a clearance sale will take place if and only if the inventory level is above the clear-down-to level and the clearance sale always reduces the inventory to that level. The details are summarized in Theorem 1 below.

Theorem 1

(i). The functions $u_t(\mathbf{z}, q)$ and $\pi_t(\mathbf{y})$ are anti-multimodular.
(ii). The optimal policy for clearance sales is characterized by $\bar{z}_{t,i}$, where $\bar{z}_{t,i}$ is a decreasing function of $y_{i+1}, y_{i+2}, \ldots, y_{n-1}$ and is independent of $y_1, y_2, \ldots, y_i$. The optimal policy is:

$$z^*_{t,i}(\mathbf{y}) = \begin{cases} y_i & \text{if } y_i \leq \bar{z}_{t,i}; \\ \bar{z}_{t,i} & \text{if } y_i > \bar{z}_{t,i}. \end{cases}$$

In addition, the following inequalities hold:

$$-1 \le \Delta_{y_{i+1}} \bar{z}_{t,i} \le \Delta_{y_{i+2}} \bar{z}_{t,i} \le \ldots \le \Delta_{y_{n-1}} \bar{z}_{t,i} \le 0.$$

(iii). The optimal replenishment quantity $q_t^*(\mathbf{y})$ *is a decreasing function of* $y_1, y_2, \ldots, y_{n-1}$*, and the following inequalities hold:*

$$-1 \le \Delta_{y_{n-1}} q_t^* \le \Delta_{y_{n-2}} q_t^* \le \ldots \le \Delta_{y_1} q_t^* \le 0.$$

The quantities $\bar{z}_{t,i}$ are state-dependent thresholds, and they depend only on inventories that are newer than i. Specifically, the more inventories with a remaining lifetime longer than i, the less inventory with an i-period remaining lifetime should be carried to the next period. In addition, the thresholds $\bar{z}_{t,i}$ are more sensitive to the inventories with a remaining lifetime closer to i. Similarly, the inequalities about the optimal order quantity confirm that the order quantity is more sensitive to newer inventory than to older inventory.

This model is extended by Liu et al. (2019) to allow for two perishable products with a dependent supply. Their research is based on the case of a blood center that periodically collects whole blood and processes it into multiple blood products such as red blood cells and platelets. They similarly characterize the optimal clearance sale and replenishment policies by showing the multimodularity of the objective function.

While both Li and Yu (2014) and Liu et al. (2019) assume that firms control how inventory is depleted and hence that the FIFO issuing rule is optimal, Li et al. (2016) study a problem in which consumers control inventory depletion and hence the LIFO rule is appropriate. In offline retailing, consumers observe expiration dates and decide which items to pick. Therefore, fresher items typically sell first on a LIFO basis. The problem is challenging not only because the state space is large but also because inventory systems under LIFO are known to lack the common technical properties such as concavity needed for analysis, let along stronger properties such as multimodularity.

The sequence of the events is the same as that in the previous models. As the regular demand is fulfilled using the LIFO rule, given an initial state $\mathbf{y} = (y_1, y_2, \ldots, y_{n-1})$, the state transition becomes $\mathbf{Y}(q, \mathbf{z}, D) = (Y_1, Y_2, \ldots, Y_{n-1})$, where for $1 \le i \le n-2$

$$Y_i(q, \mathbf{z}, D) = (z_{i+1} - (D - q - \sum_{k=i+2}^{n-1} z_k)^+)^+$$

and

$$Y_{n-1}(q, \mathbf{z}, D) = (q - D)^+.$$

Here, the notation $x^+ = \max\{x, 0\}$. The quantity of outdated inventory is

$$S(q, \mathbf{z}, D) = (z_1 - (D - q - \sum_{i=2}^{n-1} z_i)^+)^+.$$

The dynamic programming formulation is as follows:

$$\pi_t(\mathbf{y}) = s\sum_{i=1}^{n-1} y_i + \max_{\mathbf{0}\le\mathbf{z}\le\mathbf{y},q\ge 0} u_t(\mathbf{z}, q),$$

where

$$u_t(\mathbf{z}, q) = -(s+h)\sum_{i=1}^{n-1} z_i - \alpha cq + \alpha\mathbb{E}\Big\{ r\min(q+\sum_{i=1}^{n-1} z_i, D) - \theta S(q, \mathbf{z}, D) + \pi_{t+1}(\mathbf{Y}(q, \mathbf{z}, D))\Big\}.$$

The optimal clearance sale policies are shown in Figure 1. Under the LIFO rule, there are two thresholds for each age group of inventory: a lower and an upper threshold. For an age group with a remaining lifetime of two periods or more, if its inventory level is below the lower threshold, then there is no clearance sale; if it is above the upper threshold, then it will be cleared down to the upper threshold.

The optimal policy for the age group with a remaining lifetime of one period is very different, however. Clearance sales may take place if its inventory level is above the upper threshold or below the lower threshold. The lower the initial inventory level, the more new supply is needed to meet demand in the current period. However, the more new supply there is, the less likely it is that the oldest items will be used to meet demand because customers always select the newest items first. The retailer is therefore better off clearing the small number of oldest items to recoup some revenue and avoid outdating. The practice of avoiding having the newest and the oldest items in the system at the same time through clearance sales is unique to the inventory systems under LIFO.

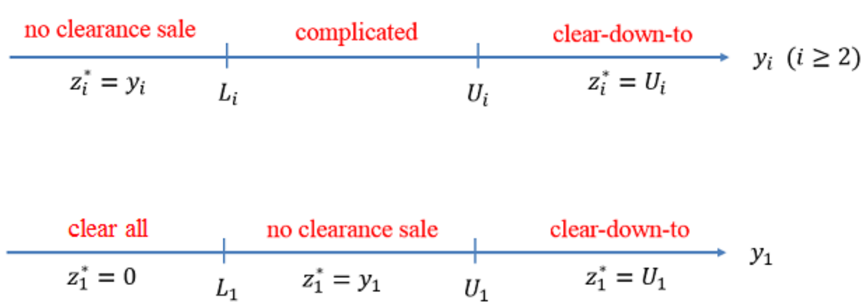

Fig. 1 Optimal clearance sale policies for inventory with different remaining lifetimes under LIFO

To implement the model in practice requires solving a dynamic program with a multi-dimensional state space and a non-concave objective function, which is challenging. Motivated by the structural properties of the optimal policy, Li et al. (2016) consider two myopic heuristics. For both heuristics, the value-to-go function is approximated by $\pi_{t+1}(\mathbf{y}) = \frac{s+\alpha c-h}{2} \sum_{i=1}^{n-1} y_i$ because the marginal values of inventories are bounded between s and $\alpha c - h$. In the first heuristic, in computing the order quantity and clearance sale quantity, all inventories on hand are treated as if they would expire in one period. Let $y = \sum_{i=1}^{n-1} y_i$. The heuristic policies are then derived from the following one-period problem:

$$\max_{\mathbf{0} \le z \le y, q \ge 0} -(s+h)z - \alpha cq + \alpha \mathbb{E}\Big[\ r \min(q+z, D) - \theta(z-(D-q)^+)^+ + \frac{s+\alpha c-h}{2}(q-D)^+\Big].$$

In the second heuristic, in addition to the total inventory level y, information about y_1, which is the inventory level of items with a one-period lifetime remaining, is also needed. The heuristic policies are obtained by solving the following:

$$\max_{z_1, z, q} -(s+h)(z_1+z) - \alpha cq + \alpha \mathbb{E}\Big[\ r \min(q+z_1+z, D) - \theta(z_1-(D-q-z)^+)^+ + \frac{s+\alpha c-h}{2}(q+z-D)^+\Big],$$

subject to the constraints: $0 \le z_1 \le y_1, 0 \le z \le y - y_1, q \ge 0$.

In numerical experiments, the first heuristic may generate as much as 7% less expected profit than the optimal policy. The second heuristic significantly outperforms the first heuristic, and its profit is consistently very close to the optimal profit.

The analysis in Li et al. (2016) demonstrates that inventory with a one-period remaining lifetime (i.e., the oldest inventory) plays a qualitatively different role than inventories of other age groups. First, the optimal order quantity is monotonic in the oldest inventory, but it is not necessarily so in other inventories. Second, the optimal policy on clearance sales with respect to the oldest inventory is to clear all, to not clear, and then to clear down to a certain level when the inventory level increases. The optimal policies with respect to other inventories, however, are different. In particular, clearance sales won't happen when inventories are low enough.

Finally, it is critically important to keep a record of the oldest inventory, and the performance of myopic heuristics that take advantage of that record are consistently close to that of the optimal policy. The value of keeping a record of other inventories, however, is insignificant. There is no age information in the bar codes currently used in retailing. In practice, retailers typically check and remove the expired items manually. Putting items on clearance sales is also done manually. The information about the oldest inventory can be obtained during these manual processes, and the additional effort may not be significant. If this is done, then the second heuristic above can be implemented without the need to include the full age information in bar codes.

3 Multiple Classes of Demand

Firms managing perishable inventory systems usually face multiple classes of customers that differ in their requirements for product freshness. For example, in health care, patients with different diseases may require platelets of different ages. In grocery retailing, some customers may be more sensitive to product freshness than others. In these settings, in addition to the coordination of clearance sales and replenishment decisions discussed earlier, firms also need to determine the optimal allocation of perishable inventory to different classes of customers.

Abouee-Mehrizi et al. (2019) study the problem assuming that unmet demand is back-ordered and that there is a positive lead time for replenishment. Let L denote the lead time. A firm sells a perishable product with a lifetime of $n + L$ periods. Thus, the product has a remaining lifetime of n periods when the firm receives it. Without loss of generality, assume that there are n classes of customers indexed by $k = 1, ..., n$, where class k customers only accept products with a remaining lifetime longer than or equal to k. Let D^k denote the class k demand. Let $\mathbf{D} = (D^1, ..., D^n)$. Let b_k denote the per unit back-order penalty cost for class k demand. Assume that $b_n \geq b_{n-1} \geq ... \geq b_1 > 0$; that is, class i has a higher priority than class j if $i > j$.

The sequence of events is the same as that in the models reviewed in Section 2. Here, the initial state is $\mathbf{y} = (\mathbf{y}_o, \mathbf{y}_p)$, where $\mathbf{y}_o = (y_1, y_2, ..., y_n)$ and $\mathbf{y}_p = (y_{n+1}, y_{n+2}, ..., y_{n+L-1})$. The variable y_i represents the inventory position for on-hand items with a remaining lifetime of i periods for $1 \leq i \leq n$; and y_{n+i} represents the pipeline inventory that will arrive in i periods for $1 \leq i \leq L-1$. The variable q is still used to represent the order quantity and $\mathbf{z} = (z_1, z_2, ..., z_n)$ to denote the amount of on-hand inventory that is carried over to the next period. If $y_i \geq 0$, then $y_i - z_i$ denotes items with a remaining lifetime of i periods that are sold at clearance sales. If $y_i < 0$, that is, if there is unmet back-ordered demand, z_i must be equal to y_i.

Let d_j^k denote the amount of the class k demand that is met by using inventories with a remaining lifetime of j periods. Let $\mathbf{d}^k = (d_1^k, d_2^k, ..., d_n^k)$ and $\mathbf{d} = (\mathbf{d}^1, \mathbf{d}^2, ..., \mathbf{d}^n)$. Denote $O(\mathbf{D}, \mathbf{z}) = \{\mathbf{d} : \sum_{j=k}^n d_j^k \leq D^k - (z_k)^-, \sum_{k=j}^n d_j^k \leq (z_j)^+, d_j^k \geq 0 \text{ for } 1 \leq j, k \leq n\}$. Here, $x^+ = \max\{x, 0\}$ and $x^- = \min\{x, 0\}$. Let r, s, c, h and θ denote the regular price, clearance price, unit purchasing cost, unit holding cost and outdating cost, respectively. The dynamic programming formulation is as follows:

$$\pi_t(\mathbf{y}) = s \sum_{i=1}^{n} y_i + \max_{\substack{(\mathbf{y}_o)^- \leq \mathbf{z} \leq \mathbf{y}_o \\ q \geq 0}} u_t(\mathbf{z}, \mathbf{y}_p, q),$$

where

$$
\begin{aligned}
u_t(\mathbf{z}, \mathbf{y}_p, q) = & -s\sum_{i=1}^{n} z_i - h(\sum_{i=1}^{n} z_i)^+ - \alpha c q \\
& +\alpha\mathbb{E} \max_{\mathbf{d}\in O(\mathbf{D},\mathbf{z})} \Big\{ r(\sum_{k=1}^{n}\sum_{j=1}^{n} d_j^k) - \sum_{k=1}^{n} b_k(D^k - (z_k)^- - \sum_{j=k}^{n} d_j^k)^+ \\
& -\theta(z_1 - d_1^1)^+ + \pi_{t+1}(z_2 - \sum_{k=1}^{2} d_2^k, \dots, z_n - \sum_{k=1}^{n} d_n^k, \mathbf{y}_p, q)\Big\},
\end{aligned}
$$

and the terminal condition is $\pi_{T+L+1}(\mathbf{y}) = 0$.

It can be proved that the value function $\pi_t(\mathbf{y})$ and the objective function $u_t(\mathbf{z}, \mathbf{y}_p, q)$ are both anti-multimodular. The optimal clearance sales strategy follows a clear-down-to structure and the structural properties in Theorem 1 continue to hold in the presence of multiple classes of customers.

The optimal allocation policy is a sequential rationing policy. The demand with the highest priority is first satisfied, then the demand with the second-highest priority, and so on. In fulfilling each demand class, it is optimal to reserve fresher inventories at certain thresholds to meet future demand with a higher priority and use the remainder to fulfill the demand of that class as much as possible following the FIFO rule. The rationing threshold of inventory with a remaining lifetime i when fulfilling class k demand depends on the inventory levels both of products with a remaining lifetime longer than i and products with a remaining lifetime shorter than k. Anti-multimodularity also implies that these thresholds have bounded sensitivities with respect to the state variables.

Finally, through numerical studies, Abouee-Mehrizi et al. (2019) compare three strategies a firm can use to improve the management of perishables: decreasing the lead time, increasing the lifetime of products, and increasing customers' willingness to accept older products. They show that, among these three strategies, decreasing lead time is the most efficient, with a potential cost benefit of 20% in the numerical examples. Increasing the lifetime is more efficient than increasing customers' willingness to accept older products.

Chen et al. (2019a) study a similar problem but under the assumption that the unmet demand is lost and the replenishment lead time is zero. D^k is again used to represent the demand from class k customers who only accept products with a remaining lifetime of at least k periods. Let p_i denote the per unit lost-sales penalty cost for class i demand. Assume that $p_n \geq p_{n-1} \geq \dots \geq p_1 > 0$; that is, class i has a higher priority than class j if $i > j$.

The sequence of events and notation for the state and decision variables are the same as those in Li and Yu (2014). In what follows, q and z_n are used interchangeably. Let d_j^k denote the amount of the class k demand that is met by using inventories with a remaining lifetime of j periods. Let $\mathbf{d}^k = (d_1^k, d_2^k, \dots, d_n^k)$ and $\mathbf{d} = (\mathbf{d}^1, \mathbf{d}^2, \dots, \mathbf{d}^n)$. Denote $O(\mathbf{D}) = \{\mathbf{d} : \sum_{j=k}^{n} d_j^k \leq D^k, \sum_{k=j}^{n} d_j^k \leq z_j, d_j^k \geq 0 \text{ for } 1 \leq j, k \leq n\}$. Let r, s, c, h and θ denote the regular price, clearance price, unit purchasing cost, unit holding cost and outdating cost, respectively. The dynamic programming formulation is as follows:

$$\pi_t(\mathbf{y}) = s\sum_{i=1}^{n-1} y_i + \max_{\substack{\mathbf{0}\le\mathbf{z}\le\mathbf{y}\\ q\ge 0}} u_t(\mathbf{z}, q),$$

where

$$u_t(\mathbf{z}, q) = -(s+h)\sum_{i=1}^{n-1} z_i - \alpha c q + \alpha\mathbb{E}\max_{\mathbf{d}\in O(\mathbf{D})}\Bigg\{r(\sum_{k=1}^{n}\sum_{j=1}^{n} d_j^k) - \theta(z_1 - d_1^1)^+$$
$$-\sum_{k=1}^{n} p_k(D^k - \sum_{j=k}^{n} d_j^k)^+ + \pi_{t+1}(z_2 - \sum_{k=1}^{2} d_2^k, \ldots, z_n - \sum_{k=1}^{n} d_n^k)\Bigg\},$$

and the terminal condition is $\pi_{T+1}(\mathbf{y}) = 0$.

Chen et al. (2019a) prove that Theorem 1 still holds in this model. That is, the value function $\pi_t(\mathbf{y})$ and the objective function $u_t(\mathbf{z}, \mathbf{y}_p, q)$ are both anti-multimodular; the optimal strategy on clearance sales follows a clear-down-to structure and the optimal order quantity is decreasing with inventory level and is more sensitive to changes in newer inventory. The optimal allocation policy is simpler than that in Abouee-Mehrizi et al. (2019) due to the assumption of zero lead time. In particular, demand with a higher priority should be satisfied as much as possible and on a FIFO basis before demand with a lower priority.

Based on the structural properties of the optimal value function, Chen et al. (2019a) then develop an adaptive approximation approach to overcome the curse of dimensionality in solving the dynamic program. The essential idea is to approximate the value function by a linear combination of a one-dimensional function $B_j(x)$, i.e., letting the approximate value function be $\hat{\pi}_{t+1}(\mathbf{y}) = \sum_{j=1}^{n-1}\eta_{t+1,j}B_{t+1,n-1}(\sum_{l=j}^{n-1} y_l)$. Here, for each $j \in \{1, 2, \ldots, n-1\}$, $B_{t,j}(y)$ can be recursively solved by the following one-dimensional dynamic program:

$$B_{t,j}(y) = sy + \max_{\substack{\mathbf{0}\le\mathbf{z}\le y\mathbf{e}_j\\ q\ge 0}} u_t(\mathbf{z}, q),$$

where

$$u_t(\mathbf{z}, q) = -(s+h)\sum_{i=1}^{n-1} z_i - \alpha c q + \alpha\mathbb{E}\max_{\mathbf{d}\in O(\mathbf{D})}\Bigg\{r(\sum_{k=1}^{n}\sum_{j=1}^{n} d_j^k) - \theta(z_1 - d_1^1)^+$$
$$-\sum_{k=1}^{n} p_k(D^k - \sum_{j=k}^{n} d_j^k)^+ + \hat{\pi}_{t+1}(z_2 - \sum_{k=1}^{2} d_2^k, \ldots, z_n - \sum_{k=1}^{n} d_n^k)\Bigg\}.$$

The weights $\eta_{t+1,j}$ for $j \in \{1, 2, \ldots, n-1\}$ are calculated based on the relationship between the marginal values of $B_{t,j}$ and $B_{t,n-1}$. Constructed in this way, $\hat{\pi}_t(\mathbf{y})$ retains the anti-multimodularity of the optimal value function. Fresher inventory has a higher marginal value under this approximate value function.

Under this approximation scheme, the heuristic policy generated by the approximate value function retains the same structural properties as the optimal policy. Numerical studies demonstrate that the proposed approximation approach is nearly optimal, with tan average optimality gap of 0.30%, and significantly outperforms the other heuristics in the literature.

Chen et al. (2019b) apply some of the ideas in Chen et al. (2019a) to a setting in which a blood center faces two classes of age-differentiated demand for platelets from downstream hospitals. Using the tool of multimodularity, they characterize the structure of the optimal policy for whole blood collection, platelet production and inventory issuing, rationing and disposal. Fu et al. (2019) also study perishable inventory systems with multiple classes of demand. However, as they consider product returns and remanufacturing, their model and results are reviewed in Chapter 11.

4 Multiple Locations

Research on non-perishable inventory systems shows that transshipment can balance inventories in different locations and hence simultaneously reduce overages at some locations and shortages at others. This section introduces two papers that provide new insights into the roles and value of transshipment in perishable inventory systems.

Li et al. (2021) explore the idea of transshipment in an offline retailer with a LIFO inventory issuing rule. The retailer owns two outlets, indexed by superscript $i = 1, 2$. The products they sell have an n-period lifetime. The products can be sold at either a regular price p, or a clearance sale price s. Under a regular price, the demand at each outlet is random and is modeled by random variable D^i. The demand under a clearance sale is sufficiently high (or s is sufficiently low) that the products on sale will never go unsold. Assume that D^1 and D^2 are identically distributed but not necessarily independent.

The sequence of events is as follows. 1) At the beginning of a period, the retailer determines how much to order, how much should be sold in a clearance sale and how much and what should be transshipped from one outlet to the other. 2) The random demand for regular sales is realized and satisfied. 3) At the end of the period, the unsold inventory with a remaining lifetime of one period expires. Assume that there is no transshipment cost in the model.

For outlet i, the initial inventory is represented by a vector $\mathbf{x}^i = (x_1^i, x_2^i, \ldots, x_{n-1}^i)$, where x_j^i represents the inventory with a remaining lifetime of j periods at outlet i. Let $x_j = x_j^1 + x_j^2$. The system state is captured by $\mathbf{x} = (x_1, x_2, \ldots, x_{n-1})$. Let q^i be the order quantity of new items at outlet i. Let $\mathbf{z}^i = (z_1^i, .., z_{n-1}^i)$, where z_j^i is the inventory with a remaining lifetime of j periods that retail outlet i has after transshipment and clearance sales. As such, the total amount of inventory with a remaining lifetime of j periods available for regular sale is $z_j^1 + z_j^2$ and the amount sold in clearance sales is $x_j - z_j^1 - z_j^2$. Customers will always choose the freshest products first; that is, inventory leaves the retail shelf on a LIFO basis. Suppose that the system state becomes $\mathbf{Y}^i(q^i, \mathbf{z}^i, D^i) = (Y_1^i, Y_2^i, ..., Y_{n-1}^i)$ in the next period. Then, for $1 \le j \le n-2$

$$Y_j^i(q^i, \mathbf{z}^i, D^i) = (z_{j+1}^i - (D^i - q^i - \sum_{k=j+2}^{n-1} z_k^i)^+)^+$$

and

$$Y_{n-1}^i(q^i, \mathbf{z}^i, D^i) = (q^i - D^i)^+.$$

The amount of outdated inventory is

$$S(q^i, \mathbf{z}^i, D^i) = (z_1^i - (D^i - q^i - \sum_{j=2}^{n-1} z_j^i)^+)^+.$$

Let c, θ, and α be the ordering cost, outdating cost and the discounting factor, respectively. Without loss of generality, assume that there is no holding cost. The dynamic programming formulation is then as follows:

$$J_t(\mathbf{z}^i, q^i) = -s \sum_{j=1}^{n-1} z_j^i - cq^i + p\mathbb{E}\min(q^i + \sum_{j=1}^{n-1} z_j^i, D^i) - \theta\mathbb{E}S(q^i, \mathbf{z}^i, D^i) \qquad (1)$$

and

$$v_t(\mathbf{x}) = s \sum_{j=1}^{n-1} x_j + \max\{J_t(\mathbf{z}^1, q^1) + J_t(\mathbf{z}^2, q^2) + \alpha\mathbb{E}v_{t+1}(\sum_{i=1}^{2} \mathbf{Y}^i(q^i, \mathbf{z}^i, D^i))\}, \quad (2)$$

subject to $z_j^1 + z_j^2 \le x_j$, $z_j^i \ge 0$, $q^i \ge 0$ for all $i = 1, 2$ and $j = 1, 2, ..., n-1$. On the right-hand side of (1), the second term is the purchasing cost, the third term is the revenue from regular sales, and the last term is the outdating cost. The sum of the first terms on the right-hand sides of (1) and (2) represents the revenue from clearance sales. Hence $J_t(\mathbf{z}^i, q^i)$ is the one-period profit generated at outlet i. The planning horizon is T and the terminal condition is $v_{T+1}(\mathbf{x}) = s \sum_{i=1}^{n-1} x_i$. The optimal solution to (2) is denoted by $(\bar{z}_j^i, \bar{q}^i)$, $j = 1, 2, ..., n-1$ and $i = 1, 2$.

The following theorem shows that transshipment plays two roles for perishable inventory systems under the LIFO rule. One is inventory balancing, which is well known in the literature. The other is inventory separation, which is new to the literature. In Theorem 2, we assume that the demand distribution is PF_2, which is a common assumption in the inventory literature (e.g., Porteus 2002, Huggins and Olsen 2010, Li and Yu 2012). The class of PF_2 distributions includes many commonly used distributions such as the exponential, the uniform, the Erlang, the normal and convolutions of such distributions.

Theorem 2 *Suppose that D^i has a PF_2 distribution. If $\bar{z}_i^1 = \bar{z}_i^2$ for $2 \le i \le n-1$, then there is an optimal policy such that at least one of $\bar{z}_1^1$, $\bar{z}_1^2$, $\bar{q}^1$ and $\bar{q}^2$ is zero.*

Theorem 2 includes two special cases. The first case is when $x_i = 0$ for all $i = 2, \ldots, n-1$, and the second is when the lifetime $n = 2$. In both cases, the condition $\bar{z}_i^1 = \bar{z}_i^2$ for $2 \le i \le n-1$ is obviously satisfied. Transshipment allows the retailer to send the oldest inventory to one outlet and the newest inventory to the other (i.e., separation of inventories), and to send inventory from the outlet with excess inventory to the outlet with a shortage (i.e., balance of inventories). Inventory separation occurs when one of $\bar{z}_1^1$, $\bar{z}_1^2$, $\bar{q}^1$ or $\bar{q}^2$ is zero, or when two are zero and one outlet holds only the oldest inventory and the other holds only the newest inventory.

To understand how inventories should be separated and how much benefit transshipment can generate, Li et al. (2021) consider an approximation. Under the approximation, the computation of the optimal policy relies on only two pieces of information, namely, the number of items expiring in one period (old inventory) x_1 and the number of remaining items (new inventory), denoted by $x_{[2]} = \sum_{j=2}^{n-1} x_j$. The profit-to-go is approximated by a linear function. That is, in period t, let $v_{t+1}(\mathbf{x}) = v \sum_{j=1}^{n-1} x_j$, where v is a number bounded by c and s because the marginal value of inventory is bounded by c and s.

Let z_1^i and $z_{[2]}^i$ represent the amount of old inventory and new inventory, respectively, allocated to outlet i for regular sale. Let y^i be the amount of new inventory after ordering at outlet i, that is, the order-up-to level for new inventory at outlet i. Let

$$J(z_1, y) = -sz_1 - cy + p\mathbb{E}\min(D, z_1 + y) - \theta\mathbb{E}(z_1 - (D-y)^+)^+ + \alpha v\mathbb{E}(y - D)^+,$$

and, to find a heuristic policy, solve the following one-period optimization problem:

$$\max\{(c-s)(z_{[2]}^1 + z_{[2]}^2) + J(z_1^1, y^1) + J(z_1^2, y^2)\}$$

subject to $z_1^1 + z_1^2 \le x_1$, $z_{[2]}^1 + z_{[2]}^2 \le x_{[2]}$, $z_1^i \ge 0$, $z_{[2]}^i \ge 0$, $y^i \ge z_{[2]}^i$ for $i = 1, 2$.

Li et al. (2021) then provide a theoretical bound for the gap between the performance of the approximation and the optimal profit. When the demand is compound Poisson, the bound approaches zero as the arrival rate approaches infinity. The optimal policy under the approximation is characterized by two increasing switching curves that divide the entire state space into three regions. In the first region, only one outlet holds old items but both hold new items. In the second, one outlet holds only old items and the other holds only new inventory. In the third, only one outlet holds new items while both hold old items.

Numerical studies show that transshipment and clearance sales are substitutes in terms of both increasing profit and reducing waste. Transshipment can increase profit by as much as several percentage points. It is most valuable in increasing profit when the variable cost of products is high, the outdating cost is high, the clearance sale price is low or the demand variability is high.

Zhang et al. (2018) study the transshipment of perishable inventory under a FIFO rule and exogenous base-stock levels (The model is extended to allow general ordering policies in Zhang et al. 2021 and is reviewed in Chapter 19). Their research is motivated by a platelet (a blood product with a shelf-life of three days) inventory management problem in two hospitals that belong to the same integrated healthcare system.

In the model, there are two locations indexed by superscripts $i = 1, 2$. The product has a lifetime of n periods. The sequence of events is as follows. 1) At the beginning of a generic period, the initial inventory at location i is $\mathbf{x}^i = (x_1^i, ..., x_{n-1}^i)$, where x_j^i is the inventory level of products of age j. The base stock level at location i is S^i, so $x_0^i = (S^i - \sum_{j=1}^{n-1} x_j^i)^+$ items of fresh products are ordered. 2) The random demand D^i at each location is realized and satisfied. 3) Products are transshipped from one location to the other in a FIFO manner; that is, older items are shipped first. Let u denote the total items transshipped from location 1 to location 2. A negative u implies transshipment from location 2 to location 1. Then, $-D^1 \leq u \leq D^2$. 4) After transshipment, the products at each location are issued to satisfy demand in a FIFO manner, and unmet demand is lost. 5) At the end of each period, products reaching age n are disposed of.

Let $\mathbf{X}^i = (X_1^i, ..., X_{n-1}^i)$ denote the initial inventory at location i at the beginning of the next period. Then,

$$X_j^i = (x_{j-1}^i - (D^i + u(-1)^{i+1} - \sum_{k=j}^{n-1} x_k^i)^+)^+.$$

Let p^i, h^i and θ^i denote the unit shortage, holding and outdating cost, respectively, at location i. Denote r^i as the unit transshipment cost from location i to the other location. Without loss of generality, assume that the unit ordering cost at each location is zero. Also assume that the system starts with zero inventory. The one-period cost function $L(\mathbf{x}^1, \mathbf{x}^2, u)$ is then given by:

$$\sum_{i=1}^{2} \Big[p^i(D^i + u(-1)^{i+1} - S^i)^+ + h^i(S^i - D^i + u(-1)^i)^+ \\ +\theta^i(x_{n-1}^i - D^i + u(-1)^i)^+ + r^i(u(-1)^{i+1})^+ \Big].$$

Let $C_t(\mathbf{x}^1, \mathbf{x}^2)$ denote the optimal expected cost-to-go function at period t. The optimality equation is then defined as:

$$C_t(\mathbf{x}^1, \mathbf{x}^2) = \mathbb{E}\left[\min_{-D^1 \leq u \leq D^2} L(\mathbf{x}^1, \mathbf{x}^2, u) + \alpha C_{t+1}(\mathbf{X}^1, \mathbf{X}^2) \right],$$

where α is the discount factor. The terminal condition is $C_{T+1}(\mathbf{x}^1, \mathbf{x}^2) = 0$.

Zhang et al. (2018) first provide a partial characterization of the direction of optimal transshipment. In the case of non-perishable inventory, the direction is determined by whether a location experiences a surplus or a shortage under mild cost conditions. However, for the case of perishable inventory, they show that an important additional factor is the quantity of the oldest inventory x_{n-1}^i, because of the outdating cost.

When the outdating cost is sufficiently small compared with the unit transshipment cost, the optimal transshipment policy for the perishable case is the same as that for the nonperishable case. The details are given in Theorem 3, where $-i$ denotes the location other than i and u^* and u^{N*} denote the optimal transshipment quantity for the perishable and non-perishable cases, respectively.

Theorem 3

(i)

$$u^{N*} = \begin{cases} \min\{(S^1 - D^1)^+, (D^2 - S^2)^+\} & \text{if} \quad S^1 \geq D^1, D^2 \geq S^2 \\ -\min\{(D^1 - S^1)^+, (S^2 - D^2)^+\} & \text{if} \quad D^1 \geq S^1, S^2 \geq D^2 \\ 0 & \text{otherwise.} \end{cases}$$

(ii) $|u^*| \geq |u^{N*}|$.
(iii) If $\theta^i \leq r^i - h^i + h^{-i}$ *for* $i = 1, 2$, *then* $u^* = u^{N*}$.

Theorem 3 shows that, in general, the optimal transshipment quantity for the non-perishable case provides a lower bound on that for the perishable case. Zhang et al. (2018) further present an example to show that this lower bound is non-tight. The implication of these results is that when managing perishable inventory, one should expect transshipments to occur more often or in larger quantities than for nonperishable inventory, because in the perishable case, transshipments are valuable not only for reducing shortages but also for balancing the age of products at different locations, thus reducing outdating.

They then investigate how the optimal transshipment quantity changes with the inventory level at each location. For a special case with a two-period product lifetime, they prove that the optimal cost function is $L^\natural$-convex, which implies that the optimal transshipment quantity is monotonic in the inventory level at each location. They further show via a counterexample that the property of $L^\natural$-convexity, however, does not hold in the general case of longer product lifetimes.

These findings motivate Zhang et al. (2018) to develop a simple transshipment policy that satisfies the monotonicity property. Under this policy, transshipment is triggered when there is either a shortage or immediate outdating. In this case, only the oldest products at each location are transshipped unless there is a shortage at the other location. They then derive approximations of the expected cost functions, which they then use to compute the base-stock levels for both locations. Using real-life data from platelet inventory management in hospitals, they show that the proposed policy performs well and significantly reduces the total cost compared with benchmark policies.

Finally, through numerical studies, Zhang et al. (2018) show that the value of inventory sharing for perishable products is typically higher than for nonperishable products. Interestingly, unlike for non-perishable products, the value of inventory sharing for perishable products can be strictly positive and substantial, even when demand at one location is deterministic, because old perishable products in a location with random demand can be transshipped to a location with deterministic demand to reduce outdates. The implication of this result is that when products are perishable, transshipment should be considered even though the results in the nonperishable inventory literature suggest that it has little or no value.

5 Control of Lifetimes

Firms can control the lifetimes of inventories when they enter the inventory system. For example, retailers of perishable goods are often faced with a choice between more expensive packaging that can extend shelf-life of their products and less expensive packaging that cannot. Other examples include situations in which firms can buy from multiple sources with different product lifetimes and costs.

Li et al. (2017) model a retailer who must determine in each period the optimal order quantities and types of packaging. There are two types of packaging, which they call “regular" and “active". Items in a regular package will perish in one period and the variable cost, which includes the purchasing and packaging costs, is c_1; items in an active package have a two-period lifetime and the variable cost is c_2, which is higher than c_1. Items that perish carry an outdating cost m per unit.

The total demand in each period is independently and identically distributed. Let D be the random demand in a period. It is assumed that customers always prefer an item with a longer remaining lifetime to an item with a shorter remaining lifetime. That is, inventory is depleted on a LIFO basis. When items with a remaining lifetime of two periods are out of stock, a random percentage β of customers are willing to purchase items with a remaining lifetime of one period and the remainder will walk away. Any unmet demand is lost and the penalty cost of not meeting demand is p per unit. The objective is to determine the quantity of items in the two types of packaging in each period that minimizes the total expected cost.

Let q be the quantity of items in active packaging with a two-period remaining lifetime, y be the total number of items with a one-period remaining lifetime, which includes the initial inventory at the beginning of each period and items just ordered but in regular packaging, and $V_t(x)$ be the minimum total expected cost from period t to the end of horizon T, when the initial inventory is x. The costs incurred in future periods are discounted by a discount rate α. Then,

$$V_t(x) = \min_{y \geq x, q \geq 0} J_t(y, q) - c_1 x,$$

where

$$J_t(y, q) = c_1 y + c_2 q + p\mathbb{E}g(y, q, \beta, D) + m\mathbb{E}S(y, q, \beta, D) + \alpha\mathbb{E}V_{t+1}(Y(y, q, \beta, D)).$$

The amount of unmet demand, $g(y, q, \beta, D)$, the initial inventory level in the next period $Y(y, q, \beta, D)$ and the amount of outdating in the current period $S(y, q, \beta, D)$ are given by:

$$
\begin{aligned}
g(y, q, \beta, D) &= (1-\beta)[D-q]^+ + [\beta D - \beta q - y]^+, \\
Y(y, q, \beta, D) &= (q-D)^+, \\
S(y, q, \beta, D) &= [y - \beta(D-q)^+]^+.
\end{aligned}
$$

Here, $[x]^+ = \max\{x, 0\}$. In the above expressions, $[\beta D - \beta q - y]^+$ and $(1-\beta)[D-q]^+$ represent the amount of unmet demand due to the stockout of items with a one-period and two-period remaining lifetime, respectively. The terminal condition is $V_{T+1}(x) = -c_1 x$, which means that any unused inventory at the end of the horizon can be salvaged at a cost of c_1 per unit.

Li et al. (2017) consider two cases that differ depending on the source of uncertainty. In the first case, the total demand D in each period is random but the proportion of customers willing to accept less fresh items β is not. In this case, they show that if the proportion β is high enough, as the initial inventory level increases, the optimal policy changes from using active packaging only to using regular packaging only and finally to ordering nothing.

Note that the retailer here either uses active packaging or regular packaging, but never both at the same time. In deciding on its choice of packaging, the retailer must consider two critical factors. The first is the need to fulfill the demand in the current period, and the second is the likelihood of items with a two-period lifetime being carried over to the next period. Which packaging the retailer should use then depends on the incremental cost. The incremental cost of using active packaging decreases with the quantity of items in active packaging, and is lower than the cost of using regular packaging if and only if the quantity of items in active packaging is sufficiently large that the items in active packaging are highly likely to be carried over to the next period. When the initial inventory x is small, a large amount of extra supply is needed to fulfill the demand in the current period. The retailer should then use active packaging only because then the chance of items with a two-period lifetime being carried over to the next period for such a large order is high. As x increases, the extra supply needed to fulfill the current demand decreases. When x is sufficiently high, the retailer will switch to using regular packaging only. Using active packaging but ordering only a small number of items is suboptimal because the likelihood of a small number of items with a two-period lifetime being carried over to the next period is low.

In the second case, the total demand D in each period is known with certainty but the proportion of customers willing to accept less fresh items β is random. In this case, the optimal policy is to use either active packaging only or regular packaging only, depending on the cost parameters but independent of the initial inventory.

The analysis shows that regardless of the source of demand uncertainty, the optimal policy structure exhibits the same pattern of “separation", that is, never use both types of packaging in the same period. This phenomenon is specific to the LIFO issuing rule and cannot happen under the FIFO issuing rule. The trade-off under FIFO is different. It is worth using active packaging only when there are enough items with a one-period remaining lifetime for there to be a high probability that items with a two-period remaining lifetime will be carried over to the next period. Items with a one-period remaining lifetime can either come from the initial inventory or from a new order in regular packaging in the current period. In other words, the retailer may use both regular and active packaging at the same time under FIFO. The separation phenomenon is reminiscent of the optimal clearance sales policies in Li et al. (2016) and the optimal transshipment policies in Li et al. (2021).

From a practical standpoint, Li et al. (2017) highlight the significance of coordinating inventory decisions and packaging decisions in grocery retailing. In practice, retailers appear to focus on the choice between using only regular packaging and using only active packaging. Some retailers decide to stay with regular packaging because the additional packaging cost does not justify the benefit. Li et al. (2017) argue that retailers should consider the optimal policy, which in general only requires the partial adoption of active packaging and has a lower packaging cost than the policy of active packaging only. Retailers will find it easier to justify the additional cost if they implement the optimal policy. However, from the perspective of waste reduction, Li et al. (2017) show through numerical studies that the optimal policy is almost as good as the policy of using only active packaging. These findings are useful for retail practice.

While Li et al. (2017) focus on grocery retailing, the study by Zhou et al. (2011) is motivated by hospitals' practice of placing expedited orders for platelet inventory in addition to regular replenishments to fulfill demand. They model this problem as a perishable inventory system with dual sourcing. The platelets have a lifetime of three periods. The interval of regular orders is two periods, which is called a cycle. At the beginning of each cycle, the hospital determines the regular order quantity Q and the order-up-to level s for expedited orders in the second period of the cycle. In the analytical model, it is assumed that expedited platelets have a lifetime of two periods. All replenishments have zero leadtimes. Therefore, all platelets in the second period of a cycle are of the same age. These assumptions effectively reduce the dimension of the state space of the dynamic program to one.

Let x denote the inventory level at the beginning of cycle t. Let D^i denote the demand in period i within cycle t, where $i \in \{1, 2\}$. Assume that unmet demand is lost and that the issuing rule is FIFO. The amount of expedited units can then be expressed as $Q_e = (s-(Q-(D^1-x)^+)^+)^+$. The amount of outdated inventory is given by $O = (x - D^1)^+$. The amount of inventory at the beginning of the second period after expediting is $\tilde{X} = \max\{s, (Q - (D^1 - x)^+)^+\}$. Thus, the state at the beginning of the next cycle is $X = (\tilde{X} - D^2)^+$. The total shortage within a cycle is given by $L = (D^1 - x - Q)^+ + (D^2 - \tilde{X})$.

Without loss of generality, the regular ordering cost is assumed to be zero. Let c_e, p and θ denote the expedited unit ordering cost, unit shortage cost and outdating cost, respectively. Let $V_t(x)$ denote the optimal cost from cycle t to the end of planning horizon T. The dynamic program is then given by

$$V_t(x) = \min_{Q,s} c_e\mathbb{E}[Q_e] + \theta\mathbb{E}[O] + p\mathbb{E}[L] + \mathbb{E}[V_{t+1}(X)].$$

The terminal condition is given by $V_{T+1}(x) = \theta\mathbb{E}(x - D^1)^+ + p\mathbb{E}(D^1 - x)^+$.

Zhou et al. (2011) then show that when solving the dynamic program backwards, the optimal solutions Q^* and s^* are uniquely determined by the first-order conditions of the objective function with respect to Q and s.

Using real-life data for platelets, Zhou et al. (2011) then numerically investigate how the optimal cost and optimal decisions vary with the model parameters. In simulation studies, they also incorporate lead times and variable product lifetimes for expedited orders. The numerical results show that the optimal cost is significantly affected by demand uncertainty, lead times, seasonality and the age of expedited orders. The optimal decisions are significantly affected by a change in expected demand but not by a change in demand variance. Furthermore, the expedited order-up-to level is relatively unchanged with respect to demand uncertainty, lead times, seasonality and age. The numerical results also imply that for small hospitals with low average demand but high demand uncertainty, the (Q, s) policy is better than the Q policy where regular orders are placed every period; for large hospitals with low demand uncertainty, the Q policy would be preferred.

In a more recent paper, Chen et al. (2020) extend the model in Zhou et al. (2011) by allowing for both returns and platelet refills during the regular ordering cycle. All of the papers reviewed in this section impose strong assumptions on lifetimes. What is the form of optimal policies when there are two sources of supply with different costs and lifetimes and lifetimes are general finite numbers? This appears to be an open question.

6 Empirical Research

Studies in the literature on perishable inventory control focus on pricing and inventory policies under certain assumptions about consumer behavior and suppliers. Some interesting empirical studies, albeit with different foci, can inform inventory research.

The study by Tsiros and Heilman (2005) examines consumers' behavior with respect to expiration dates for perishable grocery products. In particular, they show that consumers' willingness to pay and their frequency of checking expiration dates depend on their perceived risk associated with expiration, which varies from product to product, their consumption rates, and their ability to take measures to stop or slow the aging process of perishable products. These findings confirm that to bring perishable inventory models closer to current practice, it is necessary to model multiple classes of customers who may have different minimum acceptable remaining product lifetimes and may use different issuing rules.

While Tsiros and Heilman (2005) focus on consumer behavior, Akkas et al. (2019) focus on the supply side. They find that the main sources of product expiration in retail stores are large case sizes relative to daily consumer demand, long lead times, minimum order rules, replenishment workload and manufacturers' incentive programs for the sales force. These findings are useful for managers developing targeted product design and information and incentives design initiatives to reduce waste. For inventory researchers, these findings show that there are opportunities to investigate ideas for managing perishable inventory that involve the whole supply chain, as opposed to only the retailer.

7 Future Research

In the literature, either FIFO or LIFO rule is assumed. The assumption behind the LIFO rule is that consumers are infinitely rational, therefore, in traditional bricks-and-mortar stores where consumers decide which items to pick, the LIFO rule is appropriate. However, the reality is more complex. In bricks-and-mortar stores, older items are usually placed in more convenient reach of customers on the shelves and picking the newest items requires additional effort. While some customers may be willing to make that effort, others may settle for items that are less fresh. One way to capture some of this complexity is to have multiple classes of consumers. For example, one class of consumers chooses items on a LIFO basis, whereas other classes use the FIFO rule but will not select items unless their remaining lifetimes are sufficiently long.

In e-commerce, where retailers control inventory issuance, the FIFO rule is usually assumed because it minimizes outdating. However, it may be suboptimal for retailers if consumer welfare, which is usually important to retailers, is sensitive to the remaining lifetimes of products. In a fuller model that captures this additional issue, the retailer should jointly determine ordering and issuing policies such that the utility, which includes revenue, inventory related costs, and the impact on consumer welfare, is maximized.

One important insight contained in Operations Management textbooks is that consolidating multiple retail outlets into one can reduce the mismatch between supply and demand if demand is not affected by the consolidation, and this is beneficial to retailers. Is this result still true when consumers choose products on a LIFO basis, which is typically the case in physical stores? The foregoing discussion suggests that it may not be true in general. Having all inventories in one location means that consumers will not choose older items unless newer items are sold out. However, if inventories are placed in multiple locations, older items may be sold in some locations if newer items are sold out, even if there might still be newer items in other locations. When consolidation benefits retailers is an interesting and practical question for future research.

References

1. Abouee-Mehrizi H, Baron O, Berman O, Chen D (2019) Managing perishable inventory systems with multiple priority classes. *Production and Operations Management*, **28**(9), 2305-2322.

2. Akkas A, Gaur V, Simchi-Levi D (2019) Drivers of product expiration in consumer packaged goods retailing. *Management Science*, **65**(5), 2179-2195.
3. Chao X, Gong X, Shi C, Yang C, Zhang H, Zhou SX (2018) Approximation algorithms for capacitated perishable inventory systems with positive lead times. *Management Science*, **64**(11), 5038-5061.
4. Chao X, Gong X, Shi C, Zhang H (2015) Approximation algorithms for perishable inventory systems. *Operations Research*, **63**(3), 585-601.
5. Chen K, Song JS, Shang J, Xiao T (2020) Managing hospital platelets inventory with mid-cycle returns and refill. Working paper, available at SSRN 3032926.
6. Chen S, Li Y, Yang Y, Zhou W (2019a) Managing perishable inventory systems with age-differentiated demand. Working paper, City University of Hong Kong.
7. Chen S, Li Y, Zhou W (2019b) Joint decisions for blood collection and platelet inventory control. *Production and Operations Management*, **28**(7), 1674-1691.
8. Chen X, Pang Z, Pan L (2014) Coordinating inventory control and pricing strategies for perishable products. *Operations Research*, **62**(2), 284-300.
9. Fries B (1975) Optimal ordering policy for a perishable commodity with fixed lifetime. *Operations Research*, **23**(1), 46-61.
10. Fu K, Gong X, Liang G (2019) Managing perishable inventory systems with product returns and remanufacturing. *Production and Operations Management*, **28**(6), 1366-1386.
11. Hu P, Shum S, Yu M (2016) Joint inventory and markdown management for perishable goods with strategic consumer behavior. *Operations Research*, **64**(1), 118-134.
12. Huggins EL, Olsen TL (2010) Inventory control with generalized expediting. *Operations Research*, **58**(5), 1414-1426.
13. Karaesmen IZ, Scheller-Wolf A, Deniz B (2011) Managing perishable and aging inventories: review and future research directions. K. Kempf, P. Keskinocak, R. Uzsoy, eds. *Planning Production and Inventories in the Extended Enterprise.* Springer.
14. Li Q, Yu P (2012) On the quasiconcavity of lost-sales inventory models with fixed costs. *Operations Research*, **60**(2), 286-291.
15. Li Q, Yu P (2014) Multimodularity and its applications in three stochastic dynamic inventory problems. *Manufacturing and Service Operations Management*, **16**(3), 455-463.
16. Li Q, Yu P, Du L (2021) Separation of perishable inventories in offline retailing through transshipment. *Operations Research*, forthcoming.
17. Li Q, Yu P, Wu X (2016) Managing perishable inventories in retailing: Replenishment, clearance sales, and segregation. *Operations Research*, **64**(6), 1270-1284.
18. Li Q, Yu P, Wu X (2017) Shelf life extending packaging, inventory control, and grocery retailing. *Production and Operations Management*, **26**(7), 1369-1382.
19. Liu Y, Feng Y, Lai G (2019) Analysis of optimal inventory management for dual blood products with dual sources. Working Paper, available at SSRN 3323514.
20. Lu Y, Song J (2005) Order-based cost optimization in assemble-to-order systems. *Operations Research*, **53**(1), 151-169.
21. Nahmias S (1982) Perishable inventory theory: A review. *Operations Research*, **30**(4), 680-707.
22. Nahmias S (2011) *Perishable Inventory Systems.* Springer.
23. Nandakumar P, Morton TE (1993) Near myopic heuristics for the fixed-life perishability problem. *Management Science*, **39**(12), 1490-1498.
24. Porteus EL (2002) *Foundations of Stochastic Inventory Theory.* Stanford University Press.
25. Prastacos GP (1982) Blood inventory management: An overview of theory and practice. *Management Science*, **30**(7), 777-800.
26. Tsiros M, Heilman CM (2005) The effect of expiration dates and perceived risk on purchasing behavior in grocery store perishable categories. *Journal of Marketing*, **69**(2), 114-129.
27. Zhang C, Ayer T, White CC (2018) Perishable Inventory Sharing in a Two-location System. Working Paper, Duke University.
28. Zhang C, Ayer T, White CC, Bodeker JN, Roback JD (2021) Inventory Sharing for Perishable Products: Application to Platelet Inventory Management in Hospital Blood Banks. Working Paper, Duke University.

29. Zhou D, Leung LC, Pierskalla WP (2011) Inventory management of platelets in hospitals: optimal inventory policy for perishable products with emergency replenishments. *Manufacturing and Service Operations Management*, **13**(4), 420-438.
30. Zipkin P (2008) On the structure of lost-sales inventory models. *Operations Research*, **56**(4), 937-944.