2017

Inventory rotation of medical supplies for emergency response

Quan Spring Zhou
University of Wollongong, szhou@uow.edu.au

Tava Lennon Olsen
University of Auckland

Inventory rotation of medical supplies for emergency response

Abstract

We investigate an inventory control system for a national medical reserve to rotate its long-life perishable product to a hospital. This work is motivated by the serious expiration problem existing in reserves prepared for emergency response. We explicitly consider the perishability of a long-life product, such as latex gloves, and study the joint rotation and ordering decisions. The optimal policy is characterised by two thresholds, and the whole shelf life horizon can be divided into two phases: non-rotation and then rotation after a critical period. We characterise the monotonicity of the order-up-to levels. We find that the optimal policy structure preserves well when extended to scenarios with a capacity constraint and multiple planning horizons. This system possesses an easy-to-implement optimal policy structure, and moreover, implies that we should not always ignore the perishability of long-life products.

Publication Details

Quan Spring Zhou, Tava Lennon Olsen, Inventory rotation of medical supplies for emergency response, European Journal of Operational Research, Volume 257, Issue 3, 2017, Pages 810-821.

Inventory Rotation of Medical Supplies for Emergency Response

Quan Spring Zhou, Tava Lennon Olsen*

The University of Auckland Business School, 12 Grafton Road, Auckland 1142, New Zealand

Abstract

We investigate an inventory control system for a national medical reserve to rotate its long-life perishable product to a hospital. This work is motivated by the serious expiration problem existing in reserves prepared for emergency response. We explicitly consider the perishability of a long-life product, such as latex gloves, and study the joint rotation and ordering decisions. The optimal policy is characterised by two thresholds, and the whole shelf life horizon can be divided into two phases: non-rotation and then rotation after a critical period. We characterise the monotonicity of the order-up-to levels. We find that the optimal policy structure preserves well when extended to scenarios with a capacity constraint and multiple planning horizons. This system possesses an easy-to-implement optimal policy structure, and moreover, implies that we should not always ignore the perishability of long-life products.

Keywords: inventory control, long-life perishable product, rotation, Markov decision process

1. Introduction

To prepare for emergencies, many countries hold back-up medical supplies, which are referred herein as "the reserve." Typical products in the reserve include anti-flu drugs, gloves, gowns, syringes, vaccines, etc. Governments usually require the reserve to maintain a minimum stock level so that it will contain sufficient supplies for most emergencies (e.g., New Zealand Ministry of Health, 2009). This minimum stock level is quite high, because it targets the fulfilment of demand from the whole affected population after an emergency. Although the shelf life of items in the reserve may be as long as several years, the likelihood of a large-scale public health emergency during that period is relatively low. Thus, after several years sitting in the reserve, many medical stocks expire before being used; even non-pharmacy items such as gloves and syringes have expiry dates due to perishability and seals on sterile packaging. This causes substantial waste as the reserve must dispose of expired items and replace them with new items. New Zealand has recently dumped almost 1.5 million doses

*Corresponding author

Email addresses: q.zhou@auckland.ac.nz (Quan Spring Zhou), t.olsen@auckland.ac.nz (Tava Lennon Olsen)

of expired anti-flu drug, valued at $30 million in original cost and $110 million in retail price (Duff, 2014). Australia dumped $200 million of reserve products, which had passed their use-by dates, into landfills (Woodhead, 2014). However, expiration is not limited to these two countries; many countries face serious expiration issues with their reserves (Whybark, 2007).

Observing that hospitals often hold similar supplies and have a regular demand for them, we propose to rotate the reserve items to hospitals before they expire. That is, to transfer old reserve items and use them in hospitals, and at the same time, to replenish the reserve with new items. In such a way, reserve items can be used before their expiry date. This saves the effort and associated costs related to disposing of and replacing expired stocks. However, rotation involves two sets of handling costs (at the reserve and at the hospital) and so the extra cost must be weighed against the benefit of avoiding expiration. We seek to quantify the costs and the savings from rotation, in order to effectively reduce expiration and waste.

Given that rotation incurs extra costs, one may consider combining the two stockpiles and storing them in a hospital. However, this is often not a practical solution for several reasons. First, the reserves are often located outside of main city centres, so that they are protected in case of natural disasters. Further, the volume of reserve stocks is huge, compared to the capacity of most hospitals' on-site warehouses. Since hospitals are usually located in urban areas to be convenient for patients, it is very difficult to justify the cost of expanding the warehouse capacity simply for reserves. Finally, a high safety stock would be likely to degrade operational performance because the hospital will not be able to run a lean inventory system. Therefore, with two separated stockpiles, it is sensible to consider a rotation system as proposed in this paper.

We investigate an inventory rotation system for perishable items with a long shelf life and a minimum volume requirement. We study the stock rotation policy jointly with the hospital's ordering policy, analyse the optimal policy structure, and discuss the implications of the analytical results. Not surprisingly, the optimal policy will spread the rotation over a number of periods. However, it is not a priori obvious whether a hospital's order quantity should increase, decrease, or stay constant as the items in the reserve approach their expiration date. Further, should the hospital order more when ordering from the reserve (versus an outside supplier) or less? We seek to answer these questions.

The rotation system proposed in this paper is different from traditional inventory systems in its assumptions on perishability and shelf life. Typically, inventory management research assumes items with long shelf lives are not perishable; research on perishable inventory tends to study items with very short shelf lives, such as fresh vegetables, blood, etc. (e.g., Nahmias, 2011). However, we cannot always ignore the perishability of reserve items, especially when the reserve needs to constantly maintain a high inventory level and does not have a regular demand. Thus, we consider

the perishability of a long shelf life product that is traditionally assumed to be non-perishable.

Our work contributes to the literature in two key aspects. First, it provides operational insights for reserve stock rotation. To the best of our knowledge, thus far there are no operational guidelines for such a rotation system in the medical reserve; it is not clear how effective rotation could be or how to implement such a system. Second, it derives the optimal policy with a clean structure for rotating long-lifetime perishable inventory. While it is very difficult to implement the optimal policy for fixed life perishable inventory (Nahmias, 2011), the special features of the reserve, with a long shelf life and a minimum stock level, enable us to characterise an easy-to-follow optimal policy structure.

By deriving the optimal policy for the joint ordering and rotation decision, we show that rotation could be effective in reducing expiration. Our model shows that the system has a well-structured optimal policy. A policy with two up-to levels, one for ordering and the other for rotation, is optimal if the rotation cost is linear. The policy structure possesses intuitively appealing results as well as some more surprising properties. We highlight some counter-intuitive monotonicity results and explain the implications that may affect the implementation of the rotation system.

The remainder of this paper is organised as follows. Section 2 reviews the relevant literature. Section 3 describes the model and assumptions. Then, Section 4 characterises the optimal policy and discusses the structural results, and Section 5 extends the model to include more system factors. Section 6 briefly discusses the implications and concludes the paper. All proofs may be found in the Appendices and the Online Companion.

2. Literature Review

We review the literature on perishable inventory systems, inventory rotation, and literature specifically about expiration in the reserve. Perishable inventory has received considerable attention, and a comprehensive review is given by Nahmias (2011). As we consider a fixed life product in this paper, here we briefly review the fixed-life perishable inventory studies.

Perishable inventory models with fixed lifetime are complicated, because the Markov property of the stochastic process describing the number of the items in stock is lost (Schmidt and Nahmias, 1985). Fundamental characterisation of the optimal ordering policy is provided by Nahmias and Pierskalla (1973) in a two-period lifetime setting with zero lead time. Nahmias (1975) and Fries (1975), independently, characterise the structure of the optimal policy for a general fixed lifetime problem with zero lead time and continuous demand assumptions. Though these two papers take alternate approaches, Nahmias (1977) shows that they are essentially identical. Through a lengthy proof, he proves the existence of the optimal order decision $y_n(\mathbf{x})$ when n periods remain in the planning horizon, which is a function of the state vector describing the age distribution for on-hand

inventories, $\mathbf{x}$. Nahmias and Schmidt (1986) then extend the result to the discrete demand case.

Both Nahmias (1975) and Fries (1975) uncovered a fundamental property of perishable inventory systems, that is, the optimal policy is dependent on not only the amount of on-hand inventory, but also its age. However, the dimension of the state vector increases very fast as the shelf life increases, so it is difficult to track the amount of inventory of each age. Since the analysis in these two papers is lengthy and difficult to generalise, it is practically prohibitive to derive and implement the optimal policy, especially when the lifetime is significantly more than two periods (Nahmias, 2011). Due to this complexity, researchers thereafter proposed approximate policies which depend on the number of on-hand inventory only and do not require the age distribution (Schmidt and Nahmias, 1985; Olsson and Tydesjö, 2010). Recently, facilitated by the development of the concept of $L^{\natural}$-concavity, Chen et al. (2014), in their study of the coordinated decision of inventory control and pricing for fixed-life products, significantly shorten their analysis of structural properties, and prove that the optimal order quantity is monotone in the inventory level and is most sensitive to the newly placed order.

By studying the expiration of long-life reserve items, our work differs from these classical perishable inventory studies. The reserve items that we are studying have a very long shelf life that could cover many decision periods, so are not the same as a typical perishable product such as blood or fresh produce. Further, the focus of the reserve rotation system is to control expiration waste. Thus, we tradeoff the waste from expiration with the cost of rotation; holding cost at the reserve does not play any role in our model because the reserve always holds the same quantity of products. Instead, perishable inventory studies generally balance stock-out costs with holding and expiration cost.

The concept of using rotation to reduce expiration is not new, though the meaning of rotation differs with specific scenarios. It is intuitive that rotation is beneficial under some circumstances. The problem is, as we mentioned earlier, that usually there are no systematic guidelines for rotation. Recognising this, Kendall and Lee (1980) formulate blood rotation policies using goal programming to reflect conflicting objectives and priorities. The idea is that regional blood centres collect the remaining blood from each hospital at the end of a day and redistribute on the second day, to rotate the stock from small-volume hospitals to high-volume ones. The concept of rotation is also considered in retailing scenario (Bradford and Sugrue, 1991) and in inventory depletion policies (please refer to Nahmias (2011) for a review of early works). For typical perishable products with a very short life, such as the blood platelets, recent studies take into consideration that products with different ages are valued differently (Chung and Erhun, 2013) or have different demands (Civelek et al., 2015). However, rotation in these studies do not have the same meaning as is in the context of the reserve; the fixed long shelf life and the requirement of a fixed high inventory level make our scenario quite different from previous work on stock rotation.

On the management of the reserve, few papers model fixed-life perishable products. Early work on strategic oil reserves focuses on the decision of the optimal stock size, without considering perishability (e.g., Hanssmann, 1962; Oren and Wan, 1986). Extending Hanssmann (1962)'s model, a recent work by Maddah et al. (2014) considers the reserve deterioration and models the stock level decision assuming an exponential deterioration rate. However, strategic oil reserves are different from the medical reserves in nature: oil products deteriorate gradually due to volatility, but medical reserve products are stable with a fixed guaranteed life so gradual deterioration is not a concern. While the studies on oil reserves focus on the stock size decision, our work is seeking the optimal inventory policy given that the stock size is determined by the population size.

To the best of our knowledge, there is only one paper studying the expiration problem of fixed-life products in the medical reserve with a minimum inventory volume constraint. Shen et al. (2011) investigate the reserve in a Vendor Managed Inventory (VMI) scenario, in which the manufacturer holds the stock and is responsible for meeting the volume requirement. In turn, the reserve allows the manufacturer to sell the items at a predefined date to avoid expiration and spoilage. They model the production and the issuing of the stocked products by building an extended Economic Manufacturing Quantity model, and derive the bounds for production quantity and the length of each cycle.

Our work differs from Shen et al. (2011): we study the reserve expiration from the perspective of the stock holder (the government) and focus on the storage stage, while Shen et al. (2011)'s work is from the manufacturer's perspective and lies in the production process. Further, the VMI scenario considered in Shen et al. (2011)'s work may not be suitable for every country. Countries like New Zealand are far away from suppliers and have relatively small demand, so it is not likely that international suppliers will build a plant there. If the reserve is centrally held by the government, then rotation appears to be a reasonable solution to reduce expiration. Thus, it is important to study a reserve system as proposed in our study, and to derive suitable policies for such systems.

3. Model, Notation, and Assumptions

We consider a periodic review, single fixed-life product inventory system with a reserve stock, a hospital, and a supplier. The reserve needs to constantly keep a very high stock level, but does not have regular demand as it is rarely used. The hospital holding the same product uses its inventory to satisfy patient demand, and is used to ordering only from the supplier. As we model a single hospital, the hospital demand could be the aggregate demand across all hospitals in the region. The product is stable and is taken as equally effective as long as it is within its shelf life; however, it has no value and needs to be replaced if it reaches the expiry date. In order to consume the reserve items before its expiry date, we expect the hospital to order from the reserve – that is what we call "rotation."

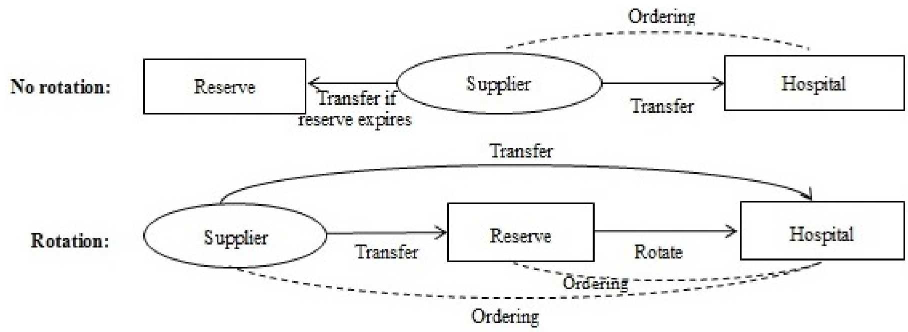

Figure 1: The rotation process flow.

Since the supplier is the only source providing additional stock to the system, all the orders are eventually transferred to the supplier. Every time the hospital rotates from the reserve, the reserve needs to get a replenishment of the same amount from the supplier simultaneously, so as to keep its constant minimum stock level. Both the rotation system and the previous non-rotation operational model are depicted in Figure 2. It is clear that the rotation system includes two decisions for the hospital: how many to order from the supplier and how many to order from the reserve; while in the previous non-rotation system, the hospital only needs to interact with the supplier.

To find the best possible result for this rotation system, we assume that the reserve and the hospital are centralised; that is, the reserve does not charge the hospital for rotation, and they as a whole pay the supplier for any new items. This centralised setting complies with the practice that the central ministry is in charge of both the reserve and the hospital and thus could be the central decision maker, though we note that in practice incentive issues may need to be addressed to achieve the best result. Although the hospital does not need to pay the reserve, rotation incurs costs of transferring inventory and managing the delivery. If items expire in the reserve, then the expired items need to be replaced and incur expiration cost which includes costs of disposal, purchasing, etc. Therefore, the goal is to balance the savings from reduced expiration with the additional costs from rotation. Formally, the objective is to minimise the expected total discounted cost for the centralised system, both the reserve and the hospital, over one shelf life horizon.

We equally divide one shelf life horizon into N periods, which are the decision epochs for ordering and rotation. We index periods within a shelf life as $i \in \{1, \cdots, N\}$, where N is the last ordering period. The hospital demand is an independent and identically distributed (i.i.d) random variable in each period. All unsatisfied demand at the hospital is backlogged. In Subsection 4.1 we discuss the lost sales case, which, though it possesses some structural properties, is less tractable when deriving the optimal policy. However, the backlogging assumption is not too rigid in our setting. The reserve items we consider for rotation here are supplies like gloves, gowns, syringes, and some long-life vaccines, etc. For these items to be used in nonemergency situations, if the hospital does not

have supply on-hand, it is likely to reschedule the surgery or vaccination and wait until the supply is available, especially when it can get the supply reasonably quickly (e.g., from the reserve).

The reserve faces no regular demand, but needs to constantly hold stocks with a very high minimum level, P, which is regulated by the government and so cannot be changed. At the beginning of each period i, the centralised decision maker observes the stock at the hospital, s_i^H, and the number of items replaced in the reserve, s_i^R, and determines the optimal decisions: the order-up-to level in the hospital, y_i, and the total number of new items in the reserve at the end of period i, z_i. Because the reserve needs immediate replenishment of the same amount every time it rotates, the quantity ordered from the supplier is always $y_i - s_i^H$ (of which $y_i - s_i^H - (z_i - s_i^R)$ is for the hospital). The objective is to minimise the expected discounted cost-to-go function.

In each period, purchasing one unit from the supplier costs c_s, rotating one unit from the reserve costs c_r, and the expected inventory holding and stockout cost is $L(y_i)$ if the hospital level is raised to y_i, where $L(\cdot)$ is assumed to be convex. At the end of the horizon, remaining items that have been in the reserve since the start need to be replaced, and this costs c_e per unit; each unit of remaining items in the hospital has a salvage value c_s. The transition of system states is $s_{i+1}^H = y_i - \xi_i$, where ξ_i is the hospital demand, and $s_{i+1}^R = z_i$. A summary of the notation is as follows.

N	the planning horizon,
s_i^H	inventory level in the hospital at the start of period i,
s_i^R	the number of items depleted from the reserve at the start of period i,
y_i	total number of products available in the hospital to satisfy demand in period i,
z_i	total number of items depleted from the reserve at the end of period i,
$L(y)$	inventory holding and stockout cost in one period with y units in the hospital,
c_s	variable unit cost when ordering from supplier,
c_r	variable unit cost of rotation when ordering from the reserve,
c_e	variable unit cost of expiration if an item expires in the reserve,
ξ	period demand for the hospital, with $\Phi(\cdot)$ being the distribution function,
P	the minimum stock level in the reserve,
α	single-period discount factor, $0 \leq \alpha \leq 1$.

Because the shelf life is long and the hospital has regular demand, it is reasonable to assume that the product is not perishable in the hospital but will expire after N periods in the reserve. We assume that, at the beginning of the planning horizon, all the reserve items are fresh with a remaining life of N periods. In Section 5.2 we will relax this assumption and consider that the remaining life may be different and shorter than the initial life. Products delivered from the supplier are fresh, which is a standard assumption and can be easily justified by taking the original life deducted by the delivery

lead time as the shelf life. Then, we define the optimal expected discounted cost function as:

$$C_{N+1}(s^H_{N+1}, s^R_{N+1}) = c_e(P - s^R_{N+1})^+ - c_s s^H_{N+1};\text{ and for } \ i = 1, \cdots, N,$$
$$C_i(s^H_i, s^R_i) = \min_{z_i \geq s^R_i;\ y_i - s^H_i \geq z_i - s^R_i} \left\{ c_s(y_i - s^H_i) + L(y_i) + c_r(z_i - s^R_i) + \alpha \int_0^\infty C_{i+1}(y_i - \xi, z_i)\, d\Phi(\xi) \right\}.$$

Since the terms $-c_s s^H_i$ and $-c_r s^R_i$ in C_i are not affected by the decision of y_i and z_i, we find it convenient to work with a slightly transformed optimal-cost function $\{W_i\}$ defined in Equation (1) (cf., Veinott, 1966). The dynamic programmes $\{W_i\}$ and $\{C_i\}$ are equivalent in terms of the existence and the structure of the optimal policy.

$$W_i(s^H_i, s^R_i) = \min_{z_i \geq s^R_i;\ y_i - s^H_i \geq z_i - s^R_i} \{G_i(y_i, z_i)\}, \qquad \forall i, \quad \text{and} \tag{1}$$
$$G_i(y_i, z_i) = \begin{cases} c_s(1-\alpha)y_i + c_r z_i + L(y_i) + \alpha c_e(P - z_i)^+, & i = N, \\ c_s(1-\alpha)y_i + c_r(1-\alpha)z_i + L(y_i) + \alpha \mathbb{E} W_{i+1}(y_i - \xi_i, z_i), & \text{otherwise.} \end{cases}$$

We now define some threshold values for y. Let $\widetilde{y}$ be the one-period optimal order-up-to level for a system without the option of rotation, and $\widehat{y}$ be the one-period optimal order-up-to level if rotation is the only option for the hospital, as is defined in Equations (2) and (3). Note both inner functions are convex, so it is straightforward that $\widehat{y} \leq \widetilde{y}$ since $c_r(1-\alpha) \geq 0$. This makes sense because the unit cost is higher under rotation, which leads to a lower order-up-to level.

$$\widetilde{y} = \operatorname*{argmin}_y \{c_s(1-\alpha)y + L(y)\}; \tag{2}$$
$$\widehat{y} = \operatorname*{argmin}_y \{c_r(1-\alpha)y + c_s(1-\alpha)y + L(y)\}. \tag{3}$$

Before deriving the optimal ordering and rotation policy, a simple analysis of the rotation decision can reveal some optimal decision properties regarding the reserve state. Let y^*_i and z^*_i denote the optimal policy values for (1). We then have Proposition 1.

Proposition 1. *For each period $i = 1, 2, \cdots, N$,*

1. *If $s^R_i \geq P$, then $z^*_i = s^R_i$, and $y^*_i = \max\left(\widetilde{y}, s^H_i\right)$.*
2. *If $s^R_i \leq P$, then $z^*_i \leq P$.*
3. *If $c_r > \alpha c_e$, then $z^*_i = s^R_i$, and $y^*_i = \max\left(\widetilde{y}, s^H_i\right)$.*

The statements in Proposition 1 are intuitive, and the proof follows from a simple inductive analysis. The first two statements suggest that the system rotates at most P items from the reserve throughout the whole horizon. Once all P products are rotated out, rotation will not be considered anymore in this shelf life horizon (because rotation incurs cost). In that case, it evolves into a standard inventory system with one supplier, and so an order-up-to $\widetilde{y}$ policy is optimal for the hospital. Intuitively, $c_r > \alpha c_e$ indicates that the cost of rotation exceeds the benefit in the last ordering period,

so it is optimal to not consider rotation in that period. Further, it suggests that the unit rotation cost is larger than the marginal benefit in all prior periods. Note that the marginal benefit of rotation is $\alpha^2 c_e$ in period $N-1, \cdots$, and $\alpha^{N-i+1} c_e$ in period i; that is, the marginal benefit decreases as the time goes back towards the beginning of the horizon. Therefore, the condition $c_r > \alpha c_e$ indicates that, for every period, the marginal cost exceeds the marginal benefit of rotation, and so rotation is not a viable option over the whole horizon. It is straightforward that a policy with an order-up-to level $\widetilde{y}$ is optimal in this case. Thus, the scenario of $c_r > \alpha c_e$ is not interesting.

Therefore, in the following we only consider the situation with $c_r \leq \alpha c_e$ and $s_i^R \leq P$ for each period. Further, we confine the feasible interval to be $s_i^R \leq z \leq P$ for the rotation decision for each period starting with $s_i^R \leq P$.

4. The Optimal Policy and Structure

This section derives the optimal policy for the model built in Section 3. We begin by performing a variable substitution and demonstrating some structural properties in Subsection 4.1. Beyond that, this problem possesses features which enable us to actually derive the optimal policy over the horizon: we first obtain the optimal policy for the final period in Subsection 4.2, and then fully characterise the optimal policy in Subsection 4.3. Finally, Subsection 4.4 discusses monotonicity results.

4.1. Variable Substitution and Structural Properties

To simplify the notation, we drop the subscript i representing the period index from variables s_i^R, s_i^H and y_i, z_i, except in cases where we specify i to emphasise that the result is for a specific period i; this should cause no confusion.

Further, we define $w = y - z$ and $s^D = s^H - s^R$, so w is the difference between the hospital's ordering decision and the reserve's rotation decision and s^D is the difference between the hospital state and the reserve state. We use them to replace the old variables; thus, s^D and s^R are the new state variables, and z and w are the new decision variables. Then, the optimisation in (1) is replaced by the one in (4); J_i is equivalent to the problem W_i. Thus, we can solve the dynamic programme J_i and get optimal values of w^* and z^*, and then the optimal value $y^* = w^* + z^*$ is for the original problem C_i. We call J_i the optimal cost function, and g_i the period cost function.

$$J_i(s^D, s^R) = \min_{s^R \leq z \leq P; w \geq s^D} \{g_i(w, z)\}, \quad i = 1, \cdots, N, \text{ where} \tag{4}$$

$$g_N(w, z) = c_r z + \alpha c_e (P - z)^+ + c_s(1-\alpha)(w+z) + L(w+z);$$

$$g_i(w, z) = c_r(1-\alpha)z + c_s(1-\alpha)(w+z) + L(w+z) + f_{i+1}(w, z), \quad i = 1, \cdots, N-1,$$

$$\text{with } f_{i+1}(w, z) = \alpha \mathbb{E} J_{i+1}(w - \xi, z).$$

To solve J_i, we separate it into a bi-level minimisation problem by first minimising over w under a given z. So, we define $\overline{w}_i(z)$ as Equation (5), and use it to facilitate solving the problem.

$$\overline{w}_i(z) = \underset{w}{\operatorname{argmin}}\{c_s(1-\alpha)(w+z) + L(w+z) + f_{i+1}(w,z)\}, \quad i = 1, \cdots, N. \tag{5}$$

Using z and w as decision variables enables us to derive the optimal policy in Subsection 4.3. Before that, we use the concept of $L^\natural$-convexity to prove some structural results, which requires another variable substitution; we include the analysis in Appendix A. In the following, Theorem 1 is implied by the $L^\natural$-convexity property and Theorem A.1, and then in Corollary 1 we obtain more characteristics for $\overline{w}_i(z)$.

Theorem 1. *If $c_r \leq \alpha c_e$, for period $i = 1, \cdots, N$,*

1. *The function $J_i(s^D, s^R)$ is convex and supermodular on $s^D \times s^R$, the function $g_i(w,z)$ is convex and supermodular on $w \times z$.*
2. *The optimal decision $w^*(s^D, s^R)$ is nondecreasing in s^D and nonincreasing in s^R, and $z^*(s^D, s^R)$ is nonincreasing in s^D and nondecreasing in s^R, and for any $\omega \geq 0$,*

$$w_i^*(s^D + \omega, s^R - \omega) \leq w_i^*(s^D, s^R) + \omega, \text{ and } z_i^*(s^D + \omega, s^R - \omega) \geq z_i^*(s^D, s^R) - \omega.$$

3. *Given z, the optimal decision $\overline{w}_i(z)$ is nonincreasing in z and for any $\omega \geq 0$,*

$$\overline{w}_i(z - \omega) \leq \overline{w}_i(z) + \omega.$$

Corollary 1. *For period $i = 1, \cdots, N$,*

1. *$\overline{w}_i(z)$ is strictly decreasing when $z \leq P$, and so is its inverse function $\overline{w}_i^{-1}(w)$.*
2. *Given $z \leq P$, $\overline{w}_i(z) + z$ is increasing in z, and $P + \overline{w}_i(P) = \widetilde{y}$.*
3. *The optimal cost function $J_i(s^D, s^R)$ is increasing in s^D and s^R.*

The monotonicity of $\overline{w}_i(z)$ and $z + \overline{w}_i(z)$ indicates the interaction between the ordering and rotation decisions. As $g_i(w,z)$ is supermodular in $w \times z$, we say $g_i(w,z)$ exhibits cost substitutability (e.g., Topkis, 1998). Cost substitutability means that the rotation decision z and the decision w have opposite directional effects on the cost function. Therefore, in order to minimise the cost, $\overline{w}_i(z)$ moves in an opposite direction with z, that is, $\overline{w}_i(z)$ decreases with z. Following from the definition of $\overline{w}_i(z)$, it is clear that $z + \overline{w}_i(z)$ is the corresponding optimal value of y for a given z. That $z + \overline{w}_i(z)$ is nondecreasing in z indicates that the ordering decision y and the rotation decision z are cost complementary. That is, in order to minimise cost, the ordering decision and the rotation decision need to move in the same direction.

It is intuitive that a higher rotate-up-to level (rotation size) calls for a higher order-up-to level (order size) in the hospital, because rotated items need to go to the hospital. Also, a higher order-

up-to level (order size) has a larger potential to take items from the reserve, and so tends to lead to a higher rotate-up-to level (rotation size). This cost complementarity is bounded by the constraint that the rotation size could not be larger than the order size, i.e., $y - s^H \geq z - s^R$, which links the two decisions together. The interaction between the ordering decision and the rotation decision drives the system and embeds the underlying tradeoffs.

As is demonstrated in Appendix C, the lost sales case also possesses the structural properties in Theorem 1 and Corollary 1. However, due to the different state transition function, the optimal policy for the lost sales case is not tractable. Nevertheless, as previously mentioned, the backlogging assumption is reasonable for reserve items used in nonemergency scenarios and gives helpful insights. Therefore, we focus on the backlogging case and discuss the optimal policy further in the following.

4.2. Optimal Policy for the Final Period

Note that the period cost g_N is different from that of the other periods, so the optimal policy for the last period N is different. Define $\overline{y}$ as Equation (6), and then we have Proposition 2.

$$\overline{y} = \operatorname*{argmin}_{y} \{c_s(1-\alpha)y + L(y) + (c_r - \alpha c_e)y\}. \tag{6}$$

Proposition 2. *When $c_r \leq \alpha c_e$, it is always optimal to rotate in period N. In this case, $z_N^* = \min\left(\overline{y} - s_N^D, P\right)$, $w_N^* = \max\left(\widetilde{y} - P, s_N^D\right)$, and so $y_N^* = z_N^* + w_N^*$ which equals the median value of $\widetilde{y}$, $s_N^D + P$, and $\overline{y}$.*

The proof of Proposition 2 is straightforward by putting $\overline{w}_N(z)$ into the cost function. The condition $\alpha c_e \geq c_r$ guarantees that the marginal benefit is larger than the marginal cost of rotation in period N, so rotation is optimal for that period. Then the concern is the rotation size. Note that $\alpha c_e \geq c_r$ does not say it is beneficial to rotate all the remaining old items in the reserve. Rather, rotating too many and thus leaving the hospital's inventory level too high increases the inventory cost in the hospital, which could outweigh the benefit of rotation.

There are two possible situations in the reserve at the beginning of period N. One is that the reserve has only a few old items left, so it can safely rotate all items to the hospital without increasing the risk of overstock. In this case, the reserve can achieve a rotate-up-to level P, and the hospital may need to order some items from the external supplier as well. The other is that the reserve has too many old items left, so the hospital will get all its orders from the reserve in period N. If that is the case, then $\overline{y}$ as defined in Equation (6) is the optimal order-up-to level, and the corresponding rotate-up-to level is $z_N^* = \min\left(\overline{y} - s_N^D, P\right)$, since it is not optimal to rotate more than P items during the horizon. As a result, the optimal order-up-to level y_N^* is a piece-wise function, and depends on the comparison of $\widetilde{y} - s_N^D$, $\overline{y} - s_N^D$, and P. Thus, the period N optimal policy has an order-up-to level and a rotate-up-to level, and both up-to levels are state dependent.

4.3. Optimal Policy Structure

We now define threshold values $\widehat{w}_i$ and $\widehat{z}_i$ to characterise the optimal policy structure for a generic period i. We will soon show that $f_{i+1}(w,z)$ is a function of only w and so can be reduced to $f_{i+1}(w)$ under certain conditions. For each period $i = 1, \cdots, N-1$, we let $\widehat{w}_i$ be the value of w that satisfies $f'_{i+1}(w) = c_r(1-\alpha)$. Note $\widehat{w}_i$ is only defined when $f_{i+1}(w_i, z_i)$ can be reduced to $f_{i+1}(w_i)$, that is, $z_i \le \overline{w}_{i+1}^{-1}(\widehat{w}_{i+1})$, where $\overline{w}_{i+1}^{-1}(\cdot)$ is the inverse function of $\overline{w}_{i+1}(\cdot)$. When $f_{i+1}(w,z)$ cannot be reduced, or $f_{i+1}(w,z)$ can be reduced to $f_{i+1}(w)$ but there is no such value of w that $f'_{i+1}(w) = c_r(1-\alpha)$, we take $\widehat{w}_i$ as an arbitrarily large value so that $\overline{w}_i^{-1}(\widehat{w}_i)$ is negative. In a slight abuse of language, we will refer to $\widehat{w}_i$ as "not defined" when it is this arbitrary large value and "defined" otherwise. This definition allows us to characterise the optimal policy that is described in Theorem 2.

Let $\widehat{w}_N = \widetilde{y} - P$, and we can combine the period N policy with the other periods. For each $i = 1, \cdots, N$, denote $\widehat{z}_i = \overline{w}_i^{-1}(\widehat{w}_i)$, so that the two thresholds are related by $\overline{w}_i(\widehat{z}_i) = \widehat{w}_i$. The pair $(\widehat{z}_i, \widehat{w}_i)$ makes the global optimal solution for the cost function $g_i(w,z)$, though it is often not the actual optimal solution due to the constraints. As is discussed below, $\widehat{z}_i$ is the threshold value for the rotation decision, and $\widehat{w}_i$ is the threshold to decide whether it is to rotate all or rotate some given the reserve state $s^R \le \widehat{z}_i$.

Theorem 2. *If $c_r \le \alpha c_e$, then for period $i = 1, \cdots, N$, an order-up-to policy is optimal for the ordering decision, and a rotate-up-to policy is optimal for the rotation decision.*

1. *When $\widehat{w}_i \le \min\left(s^D, \overline{w}_i(s^R)\right)$, it is optimal to only use rotation stock, so $w^* = s^D$. Then, $y_i^* = \max(\widehat{y}, s^H)$ for $i = 1, \cdots, N-1$, $y_N^* = \min(\overline{y}, s_N^D + P)$, and $z_i^* = y_i^* - s^D$.*
2. *When $s^D < \widehat{w}_i \le \overline{w}_i(s^R)$, it is optimal to use some rotation stock, such that $w^* = \widehat{w}_i$. Then, $y_i^* = \widehat{y}$ for $i = 1, \cdots, N-1$, $y_N^* = \widetilde{y}$, and $z_i^* = y_i^* - \widehat{w}_i$.*
3. *When $\overline{w}_i(s^R) < \widehat{w}_i$, it is optimal to not rotate, so $z^* = s^R$. Then, $y^* = \max\left(\overline{w}_i(s^R) + s^R, s^H\right)$, and $w^* = y^* - s^R$.*

The proof of Theorem 2 is by induction; we provide the four-step framework of the proof in Appendix B and the detailed proof in the Online Companion. As is shown in Theorem 2, there are three cases: "rotate all" – the hospital orders all from the reserve; "rotate some" – the hospital orders some from the reserve and the rest from the supplier; and "rotate none" – the hospital orders all from the supplier. For period $i = 1, \cdots, N-1$, an order-up-to policy is optimal for the hospital's ordering decision: the order-up-to level is fixed if it rotates from the reserve, and is state dependent if it orders solely from the supplier. There is an optimal rotate-up-to level, $\widehat{z}_i$, which differs in different periods. Note that in the "rotate all" case, it is possible that it ends up both rotating nothing and ordering nothing if $s^H \ge \widehat{y}$; but we still put it under "rotate all" since the resultant w^* level is s^D, which

means every item, if any, ordered by the hospital should come from rotation. We say that "rotation" is optimal when it is either "rotate all" or "rotate some". Figure 2 shows how the state space can be divided under the optimal policy. The two separate areas with "rotate none, order none" are to highlight the different reasons leading to the result: the left area is because the hospital's inventory is high enough to prevent any ordering (and thus prevent any rotation), i.e., $s^H \geq \widehat{y}$, while the right area is because the reserve has already rotated out enough items, i.e., $s^R \geq \hat{z}_i$.

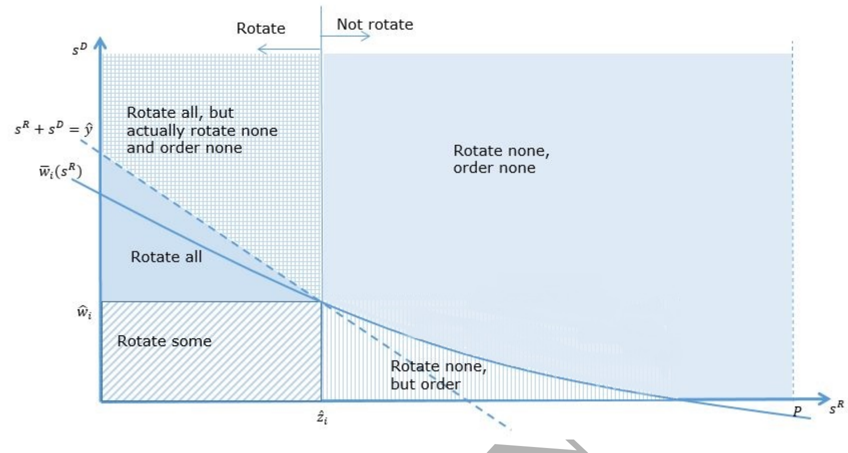

Figure 2: The optimal policy.

The state of the reserve, s^R, determines whether the hospital would order from the reserve, that is, whether rotation is a viable option in that period. There are two possibilities: if $\overline{w}_i(s^R) < \widehat{w}_i$, that is, $s^R > \widehat{z}_i$, then rotation is not viable for period i; if $\overline{w}_i(s^R) \geq \widehat{w}_i$, that is, $s^R \leq \widehat{z}_i$, then it is optimal to consider rotation. This is intuitive. The discounted cost of rotating one unit in period i is $\alpha^{i-1}c_r$ which decreases with i, while for the whole horizon, the savings from reducing expiration are always $\alpha^N c_e$. So, it makes sense to postpone rotation. But expecting rotation to occur only in the last period N is not realistic, because there is a risk that demand in that last period is not big enough to consume all the remaining items. Therefore, the centralised decision maker would like to rotate some in early periods, but does not want to rotate too many too early; so it would be optimal to rotate just enough items to offset the risk of leaving too many in the reserve. Intuitively, there should be some threshold representing the number of old items that should have been rotated out from the reserve for each period. Here, $\widehat{z}_i = \overline{w}_i^{-1}(\widehat{w}_i)$ is the threshold value for period i: once $\widehat{z}_i$ items have been rotated out from the reserve at the beginning of period i, then we can keep the rest of them, because the cost of being left in the reserve is cheaper than that of rotating. Therefore, $s^R > \widehat{z}_i$ indicates it is unnecessary to consider rotation in this period; and on the other hand, $s^R \leq \widehat{z}_i$ means not enough old items have been rotated out and so rotation is necessary.

Order-up-to level when rotating: When rotation is optimal in one period, that is, $s^R \leq \widehat{z}_i$, then the order-up-to level for the hospital is $\widehat{y}$, no matter whether all or some of the items are coming from

the reserve. This is very counterintuitive. One may think that the order-up-to level should be higher when only rotating some, because those obtained from the supplier incur a lower unit cost than those rotated from the reserve, which leads to a lower overall unit cost when rotating some; this argument is based on the difference between the cost of ordering from the supplier and that of rotating from the reserve. However, further analysis can show that the order-up-to level under rotation is not driven by this cost difference, but rather caused by a different tradeoff. From Theorem 2, given the system states, the optimal solution w^* is fixed for both cases of rotation: $w^* = s^D$ when rotating all, and $w^* = \widehat{w}_i$ when rotating part. A fixed w^* level indicates that $y^* - z^*$ is fixed in one period, that is, the difference between the order size and the rotation size should be fixed. If we decide to order one more unit, then we need to get this unit through rotation. Thus, the tradeoff is "ordering and rotating one" versus "not ordering at all", rather than "ordering and rotating one" versus "ordering one but not rotating" as one may initially think. Therefore, $\widehat{y}$ is the optimal value in both cases.

We have seen that the same order-up-to level $\widehat{y}$ is driven by the fixed value of w^*. One may then wonder where the fixed w^* comes from. The answer lies in the strong link between the ordering and rotation decisions. Rotating all means all items on order come from the reserve, so the rotation size equals the ordering size, that is $y^* - s^H = z^* - s^R$ so $w^* = s^D$. When rotating some, the ordering size and the rotation size need not to be equal, but the system would want to raise the reserve state to the rotation threshold by setting $z^* = \widehat{z}_i$, which results in $w^* = \widehat{w}_i$. Therefore, w^* is fixed for both cases, and this leads to the same order-up-to level $\widehat{y}$ no matter whether rotating all or some.

Rotate-up-to level when rotating: While the order-up-to level is fixed under rotation, the rotate-up-to level and thus the actual rotation size are determined by the difference between the hospital state and the reserve state, s^D: if $s^D \leq \widehat{w}_i$, then it is optimal to rotate some; otherwise, it is optimal to rotate all. That is because the ability to rotate is constrained by the system states. Recall that the order size includes the stock destined for rotation, that is, it requires $y - s^H \geq z - s^R$ and thus $w \geq s^D$. As a result, the ideal case, which is to achieve $z^* = \widehat{z}_i$ and so $w^* = \widehat{w}_i$, can only be achieved when $\widehat{w}_i > s^D$. In this case, the hospital orders up to $\widehat{y}$, so the order size is $\widehat{y} - s^H$, of which $\widehat{z}_i - s^R$ (which is also the rotation size) comes from the reserve and the others come from suppliers; thus, it is the case of rotating some. Otherwise, if $\widehat{w}_i \leq s^D$, then the system could not do better than keeping w as s^D. In this case, $w^* = s^D$ indicates $y^* - s^H = z^* - s^R$, that is, the order size $y^* - s^H = \widehat{y} - s^H$ should equal the rotation size $z^* - s^R$: all that the hospital orders come from the reserve; thus, it is the case of rotating all. Therefore, the rotate-up-to levels are different under the two cases: when rotating all, the reserve rotates up to $\widehat{y} - s^D$, which is less than the threshold value $\widehat{z}_i$.

Order-up-to level when not rotating: When rotation is not viable, that is, $s^R > \widehat{z}_i$, then the hospital orders up to $\overline{w}_i(s^R) + s^R$, which is independent of s^D. For a non-rotation period i, the

order-up-to level is increasing in s^R; that is, the hospital tends to order more from the supplier when more reserve items have been rotated. This is intuitive: a high level of rotated items suggests that there is not much pressure to rotate, so the hospital could order more from the supplier. Further, for all $\widehat{z}_i < s^R \leq P$, the order-up-to level, $\overline{w}_i(s^R) + s^R$, is larger than $\widehat{y}$ and smaller than $\widetilde{y}$. When it is optimal to not rotate in one period, the hospital tends to order a little bit more than $\widehat{y}$, because the ordering cost from the supplier is lower than if it is rotating. Note that here the tradeoff is between "ordering and rotating one" and "ordering one but not rotating", as opposed to the different tradeoff when rotating. However, the hospital would not order up to the ordinary level $\widetilde{y}$; rather, it orders up to less than $\widetilde{y}$ from the external supplier so as to have potential capacity for future rotation.

4.4. Monotonicity

In this section, we analyse the policy in different periods as time approaches to the end of the time horizon and provide monotonicity results. Note that we use "decreasing" and "increasing" in a relaxed rather than strict sense.

Proposition 3. *For $i = 1, \cdots, N-1$,*

1. *When $\widehat{w}_i$ is defined, $\widehat{w}_i + \widehat{z}_i = \widehat{y}$.*
2. *For any value of $z \leq P$, $\overline{w}_N(z) \geq \overline{w}_{i-1}(z) \geq \overline{w}_i(z)$.*
3. *When $\widehat{w}_i$ and $\widehat{w}_{i+1}$ are defined, $\widehat{w}_i \geq \widehat{w}_{i+1}$, and $\widehat{z}_i \leq \widehat{z}_{i+1} \leq P$.*

Proposition 3 directly follows from Lemma O.1 in the Online Companion. The result $\widehat{z}_i \leq \widehat{z}_{i+1}$ suggests that the threshold value for old items that should have been rotated out is increasing with time. That is, for a certain s^R level, we may not use rotation in early periods, but will in later periods. This is intuitive because we know that it is rational to postpone rotation. It also implies that for the same s^R level, if it is optimal to use rotation in period i, then it is optimal to use rotation in period $i+1$. Further, the proof of Theorem 2 shows that $z_i^* \leq \widehat{z}_{i+1}$ always stands if rotation is the policy in period i. This suggests that if rotation is the policy for one period, then rotation will be the policy from that period onwards. Put it in another way, it says if the optimal policy in one period is to "rotate none", then it has to be the case that rotation has not been used before. Therefore, the whole horizon can be divided into two phases: rotate none in the first phase, and then start rotation (using rotation for either all or some of the order) in the second phase.

Proposition 4. *Let t_1 be the smallest value of i such that $\overline{w}_i(0) \geq \widehat{w}_i$, then,*

1. *It is optimal to not rotate before period t_1, and consider rotation from period t_1 onwards.*
2. *Let t_2 be the largest value of i such that $c_r > \frac{\alpha^{N-i+1}}{1-\alpha+\alpha^{N-i}} c_e$, then $t_2 \leq t_1$.*
3. *The order-up-to level y_i^* is ordered, with $\widetilde{y} \geq y_1^* \geq \cdots \geq y_{t_1-1}^* \geq \widehat{y}$ and $y_j^* = \widehat{y}$ for $j \geq t_1$.*

Proposition 4 is proved by simple algebra. The critical period t_1 distinguishes the rotation and non-rotation phases. The expression $c_r > \frac{\alpha^{N-i+1}}{1-\alpha+\alpha^{N-i}}c_e$, that is, $c_r(1-\alpha) > \alpha^{N-i}(\alpha c_e - c_r)$, characterises the tradeoff between the increased period cost and the reduced cost-to-go when rotating. Rotating one unit in period i incurs cost $c_r(1-\alpha)$ in the current period, but also reduces the risk of expiration and thus reduces the cost-to-go. The saving in the cost-to-go can be evaluated based on the fixed order-up-to level for rotation periods. Recall that $w + z = \widehat{y}$ always stands when rotating. So, for a rotation period i, rotating one more unit, that is, increasing z by 1, means decreasing w by 1, and so leads to a saving of $\frac{\partial f_{i+1}(w)}{\partial w}$ which has an upper bound since $\frac{\partial f_{i+1}(w)}{\partial w} \leq \alpha^{N-i}(\alpha c_e - c_r)$. If the marginal rotation cost exceeds the upper limit of savings, then it is never optimal to use rotation in that period. Therefore, $c_r > \frac{\alpha^{N-i+1}}{1-\alpha+\alpha^{N-i}}c_e$ implies that rotation is not a viable option. On the other hand, $c_r \leq \frac{\alpha^{N-i+1}}{1-\alpha+\alpha^{N-i}}c_e$ cannot guarantee rotation will be used. Recall that $\widehat{w}_i + \widehat{z}_i = \widehat{y}$. If for a period i with $c_r \leq \frac{\alpha^{N-i+1}}{1-\alpha+\alpha^{N-i}}c_e$, it turns out that $\widehat{w}_i \geq \widehat{y}$, then $\widehat{z}_i$ can only assume the value of 0, which suggests it is optimal to not rotate. Therefore, $c_r \leq \frac{\alpha^{N-i+1}}{1-\alpha+\alpha^{N-i}}c_e$ is a sufficient but not a necessary condition for rotation, and this confirms $t_2 \leq t_1$.

Order-up-to levels as time elapses: Part 3 of Proposition 4 states that the optimal order-up-to level is nonincreasing over the whole time horizon: nonincreasing in the non-rotation phase, and then stays constant in the rotation period, as is shown in Figure 3. One may find this counterintuitive, and expect to rotate more and thus order more in later periods. This unusual trend is actually driven by the different tradeoffs around the decisions.

During the non-rotation periods, y^* is nonincreasing, and is bounded by $\widetilde{y}$ and $\widehat{y}$. It is nonincreasing because the marginal cost of ordering increases as time elapses. We know that $s^R = 0$ and $z = 0$ for all non-rotation periods $i \leq t_1$, so the order-up-to level is $y_i^* = \overline{w}_i(0)$, which decreases with i according to Proposition 3 Part 2. This reflects the perceived cost of ordering from the supplier as opposed to the cost of expiration as time approaches the expiry date. If the rotate-up-to level z, is the same for period i and $i+1$, then the expected expiration cost generated at the end of the time horizon is the same, but the discounted expiration cost is larger in period $i+1$. This means the perceived opportunity cost of ordering from the supplier is larger for period $i+1$ compared to period i. The opportunity cost refers to the cost that could have been saved if the hospital rotated from the reserve rather than ordered from the supplier. As time approaches the expiry date, this opportunity cost gets larger, so the perceived marginal cost of ordering from supplier gets larger. Therefore, when the amount of old items subject to expiration in the reserve is the same for both periods, the optimal value of w is nonincreasing, that is, $\overline{w}_{i+1}(z) \leq \overline{w}_i(z)$. With $y = w + z$, this transforms to the nonincreasing order-up-to levels, that is, $y_{i+1}^* \leq y_i^*$. This decreasing feature holds only for the non-rotation phase, i.e., $i < t_1$. Intuitively, when it gets closer to the rotation phase, the

centralised decision maker would like to leave capacity for rotation, and thus order less from supplier; but the order-up-to level will not be smaller than $\widehat{y}$.

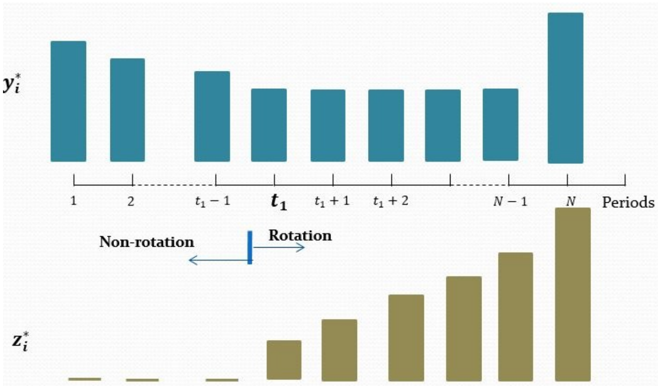

Figure 3: An illustration of the optimal order-up-to and rotate-up-to levels.

When it enters the rotation phase, the order-up-to level stays constant as $\widehat{y}$ no matter whether some or all of the ordered stock is coming from the reserve (except for in the last period N). As we discussed, this results from the fixed value of w^*. In Theorem 2, we give the conditions of when it is optimal to rotate all versus some. However, it is not easy to tell whether it is rotating all or some in the next period, without knowing the realised demand. Note that the comparison of s^D and $\widehat{w}_i$ determines whether it is rotating all or some, and the value of s^D is related to the realised demand in the previous period. If it is optimal to rotate all in period $i \geq t_1$, which suggests $s_i^D \geq \widehat{w}_i$, then period $i+1$ results in $s_{i+1}^D = s_i^D - \xi_i$. Similarly, if it is optimal to rotate some in period $i \geq t_1$, then $s_{i+1}^D = \widehat{w}_i - \xi_i$. In either case, it is not clear whether s_{i+1}^D is bigger or smaller than $\widehat{w}_{i+1}$. However, it suggests that s_{i+1}^D tends to be small and could be smaller than $\widehat{w}_{i+1}$ if demand ξ_i is big. Therefore, in the rotation phase, no matter if it is rotating all or some in a period, if demand is big, then it is likely to rotate some in the next period; otherwise, it is likely to rotate all in the next period. Nevertheless, there is no obvious pattern whether it is to rotate all or some in one arbitrary period.

5. Extensions

This section considers a number of extensions to our model, including the inclusion of a capacity constraint at the hospital and the decreased effective shelf life in a multi-horizon scenario.

5.1. With a Capacity Constraint

Note that $\overline{y}$ defined in Equation (6) could approach infinity, especially when $c_r - \alpha c_e$ is a very small negative number. In this case, the policy will be to rotate all the remaining reserve items in the last period. This may lead to a policy which leaves a whole lot of old reserve items for the last period and results in expiration in the hospital.

To rule out this possibility, we could add a capacity constraint M for the hospital: the hospital can hold at most M units of products. The capacity constraint puts an upper limit of the ordering size which is $M - s^H$, and thus limits the size of rotation. Since $\widetilde{y}$ is the optimal up-to level for the hospital if ordering from the supplier is the only option, it is reasonable to assume that the hospital has a capacity larger than $\widetilde{y}$. Also, it is straightforward that it has no impact on the optimal policy structure if $M \geq \overline{y}$, given that all the optimal order-up-to levels are smaller than $\overline{y}$ even when it is uncapacitated. Therefore, in the following we consider the situation when $\widetilde{y} \leq M \leq \overline{y}$, and show that the optimal policy structure shown in Theorem 2 still holds in this case.

Compared to the base case, we have one more constraint, $w+z \leq M$. We carry the notation from the uncapacitated case, and define the functions and values in a similar manner. To show some values and functions are different from the base case, we use a superscript M to denote the capacitated case when necessary. We use $\widehat{w}_i^M$ and $\widehat{z}_i^M$ to characterise the optimal policy, and present the structural results for the capacitated case in Proposition 5.

$$J_i^M(s^D, s^R) = \min_{s^R \leq z \leq P; w \geq s^D; w+z \leq M} \{g_i(w,z)\}, \quad i = 1, \cdots, N.$$

Proposition 5. *If the hospital has a capacity constraint M with $\widetilde{y} \leq M \leq \overline{y}$, then, with the other assumptions unchanged, by replacing the corresponding values with $\overline{w}_i^M(z)$, $\widehat{w}_i^M$, and $\widehat{z}_i^M$,*

1. *Results in Proposition 1, Proposition 3 and Proposition 4 hold.*
2. *The optimal policy structure is as stated in Theorem 2; a difference lies in period N in which $\overline{y}$ is replaced by M.*
3. *For $i = 1, \cdots, N$, for any value of $z \leq P$, $\overline{w}_i^M(z) \leq \overline{w}_i(z)$; when $\widehat{w}_i^M$ and $\widehat{w}_i$ are defined, $\widehat{w}_i^M \leq \widehat{w}_i$, and $\widehat{z}_i^M \geq \widehat{z}_i$.*
4. *Let t_1^M be the smallest value of i such that $\overline{w}_i^M(0) \geq \widehat{w}_i^M$, then $t_1^M \leq t_1$.*

The proof of the first two statements in Proposition 5 follows a similar approach of those in the base case, and the proof of the last two statements follows from simple algebra based on the result that the derivative of the optimal expected cost function with regard to s^D is larger in the capacitated case, that is, $\frac{\partial f_{i+1}^M(s^D,s^R)}{\partial s^D} \geq \frac{\partial f_{i+1}(s^D,s^R)}{\partial s^D}$.

From Proposition 5, the optimal policy and so the main structural results hold when considering a capacity constraint M as long as $\widetilde{y} \leq M \leq \overline{y}$. The threshold value for the rotation decision is larger in the capacitated case, that is, $\widehat{z}_i^M \geq \widehat{z}_i$. This means, compared to the base case, the capacitated system is more likely to rotate in a period with the same s^R level; to rotate all rather than rotate some; and to have a lower order-up-to level. To sum up, the capacitated system tends to rotate more and start rotation earlier. This is related to the constrained capacity to order and thus to rotate, and could alleviate the pressure and the risk of expiration.

5.2. Multiple Horizon Scenario

If rotation happens earlier than the last period, then, once the first shelf life horizon finishes, not all the reserve items are fresh at the start of the next shelf life horizon. The shortened remaining life poses an extra cost and brings a problem for evaluating the expiration cost: if we could not rotate one item before its actual expiry date, it expires in the reserve.

To keep the tractability of the model, we continue using $c_e(P - z_N)^+$ at the end of a horizon to evaluate the expiration cost, and add a cost factor for each rotation to account for the cost of potential early expiration. For each rotated item, we introduce a period dependent cost factor, $q_i c_e$, with $q_i \geq 0$ convexly decreasing with i. We use $q_i c_e$ to reflect the cost of having expiration happen earlier than the end of next horizon, due to rotating in period i and thus reducing the effective shelf life. So, the total unit rotation cost in period i is now $c_r^i = c_r + q_i c_e$.

Meaning of q_i*:* The extra cost associated with rotating in period i, $q_i c_e$, reflects two effects coming from the replenished items expiring early in the next horizon. First, the probability for one item to reach its actual expiry date before the end of the next horizon decreases with i. Recall that it is optimal to start rotating from period t_1. For an item rotated in period $i \leq t_1$, the replaced one will for sure expire within the next horizon. As i increases, the remaining life at the beginning of the next horizon gets longer, so the probability of expiring decreases. Second, the actual cost incurred by early expiration decreases with i. In practice, when a reserve item expires in period $i \leq N$, there are two options: replace it immediately; or wait until the last period. In either way, the related cost decreases with i. If it is replaced immediately, then this could lead to early expiration again in the following horizons, so the cost of rotating compounds and is higher for early periods. If the reserve waits until the last period to replace all expired items, then the compound effect disappears, but expired items in the reserve incur social cost as the reserve is for emergencies; the longer an expired item stays in the reserve, the higher the cost is. Therefore, with these two effects, the cost of rotating early, $q_i c_e$, convexly decreases with i.

With $c_r^i = c_r + q_i c_e$ as the unit rotation cost, we use superscript V to denote the notation for the multi-horizon scenario when necessary. Now, the optimal expected discounted cost function of period i after the variable substitution evolves to $\{J_i^V\}$ and $\{g_i^V\}$ as Equation (7).

$$J_i^V(s^D, s^R) = \min_{s^R \leq z \leq P; w \geq s^D} \left\{ g_i^V(w, z) \right\}, \quad i = 1, \cdots, N, \text{ with} \tag{7}$$

$$g_N^V(w, z) = (c_r^N - \alpha c_e)z + \alpha c_e P + c_s(1 - \alpha)(w + z) + L(w + z);$$

$$g_i^V(w, z) = (c_r^i - \alpha c_r^{i+1})z + c_s(1 - \alpha)(w + z) + L(w + z) + f_{i+1}^V(w, z), \quad i = 1, \cdots, N - 1.$$

As is shown in Proposition 7, we can extend the results of Proposition 1 to this multi-horizon model, changing the condition $c_r > \alpha c_e$ to be $c_r^N > \alpha c_e$. The intuition is the same: at most P items

will be rotated, and it is always optimal to not rotate if $c_r^N > \alpha c_e$. Note $c_r^N \geq c_r$, so the condition for rotation to be feasible is stricter in the multi-horizon scenario; this is due to the extra rotation cost considered in this scenario. Thus, we will only consider the case $c_r^N \leq \alpha c_e$, that is, $c_r + q_N c_e \leq \alpha c_e$.

Similarly, we define a threshold value $\widehat{w}_i^V$ to characterise the optimal policy structure, and show that $f_{i+1}^V(w, z)$ can be reduced to $f_{i+1}^V(w)$ under certain conditions. For each period $i = 1, \cdots, N-1$, let $\widehat{w}_i^V$ be the value of w such that $f_{i+1}^{V\prime}(\widehat{w}_i^V) = c_r^i - \alpha c_r^{i+1}$. As is in the single horizon model, $\widehat{w}_i^V$ reflects the marginal cost of rotating in period i as opposed to rotating in period $i+1$. Let $\widehat{w}_N^V = \widetilde{y} - P$. Also, define $\widehat{y}_i^V = \operatorname{argmin}_y \left\{(c_r^i - \alpha c_r^{i+1})y + c_s(1-\alpha)y + L(y)\right\}$ and $\overline{y}^V = \operatorname{argmin}_y \left\{(c_r^N - \alpha c_e)y + c_s(1-\alpha)y + L(y)\right\}$. Then we state Proposition 6 and 7.

Proposition 6. *For $i = 1, \cdots, N-1$, $\widehat{y}_i^V \leq \widehat{y}$; $\overline{y}^V \leq \overline{y}$. Further, for $i = 1, \cdots, N-2$, $\widehat{y}_i^V \leq \widehat{y}_{i+1}^V$.*

Proposition 7. *When considering multiple horizons, with the other assumptions unchanged, the results in Proposition 1 hold, upon replacing c_r with c_r^N. Further, if $c_r^N \leq \alpha c_e$, then, by replacing the corresponding values with $\widehat{y}_i^V$, $\overline{y}^V$, $\overline{w}_i^V(z)$, $\widehat{w}_i^V$, and $\widehat{z}_i^V$,*

1. *The optimal policy structure is as stated in Theorem 2.*
2. *Results in Proposition 3 and Proposition 4 Part 1 hold.*
3. *The order-up-to level y_i^{V*} is decreasing in i when $i < t_1^V$, and increasing in i when $i \geq t_1^V$. That is, $\widetilde{y} \geq y_1^{V*} \geq \cdots \geq y_{t_1^V-1}^{V*} \geq \widehat{y}_{t_1^V-1}^V$ and $y_{t_1^V}^{V*} \leq y_{t_1^V+1}^{V*} \leq \cdots \leq y_{N-1}^{V*} \leq \widehat{y}$.*

The proof of Proposition 6 follows from basic algebra, $c_r^N \geq c_r$ and that $c_r^i - \alpha c_r^{i+1}$ decreases with i. The proof of Proposition 7 follows a similar approach of that in the base case. The optimal policy for the multi-horizon model maintains the same structure and some monotonicity properties.

However, there is a difference lying in the trend of order-up-to level, as is shown in Proposition 7 Part 3. The order-up-to level when rotating, $\widehat{y}_i^V$, is increasing with period i, instead of staying constant as in the single horizon model. Analytically, this is a direct result from q_i convexly decreasing and thus unit rotation cost c_r^i decreasing with i. Moreover, this could illustrate how rotation impacts the system. In a single horizon, the timing of rotation does not have much influence, and the total amount of rotation is the only factor that matters – that is the fundamental reason why the order-up-to level is fixed in every period when rotating. However, in the multi-horizon model, rotation in different periods generates different impact on future costs, so, not only the quantity matters, but it also matters when to rotate. As the cost incurred by rotation decreases with i, it tends to rotate more and so have a higher rotate-up-to level in later periods.

Proposition 8. *Comparing the multi-horizon model with the single horizon case, we have:*

1. *For $i = 1, \cdots, N$, for any value of $z \leq P$, $\overline{w}_i^V(z) \geq \overline{w}_i(z)$; when $\widehat{w}_i^V$ and $\widehat{w}_i$ are defined, $\widehat{w}_i^V \geq \widehat{w}_i$, and $\widehat{z}_i^V \leq \widehat{z}_i$.*

2. *For the critical period,* $t_1^V \geq t_1$.
3. *About the order-up-to levels, we have* $y_i^{V*} \geq y_i^*$ *for* $i < t_1$, *and* $y_i^{V*} \leq y_i^*$ *for* $i \geq t_1^V$, *and it is not clear which one is larger when* $t_1 \leq i < t_1^V$.

Proposition 8 shows the comparison of the critical values. The proof follows from simple algebra based on the result that the derivative of the optimal expected cost function with regard to s^D is smaller in the multi-horizon case, that is, $\frac{\partial f_{i+1}^V(s^D,s^R)}{\partial s^D} \leq \frac{\partial f_{i+1}(s^D,s^R)}{\partial s^D}$. The threshold value for the rotation decision is smaller under the multi-horizon scenario, that is, $\widehat{z}_i^V \leq \widehat{z}_i$. This implies, compared to the single horizon case, the multi-horizon model is less likely to rotate in a period with the same s^R level; more likely to rotate some than rotate all; and tends to start rotation later in a horizon. This is due to the cost of future potential expiration brought by rotation in the multi-horizon scenario.

For a rotation period i, the order-up-to level in the multi-horizon scenario is no higher than that in the single horizon scenario, which is confirmed by $\widehat{y}_i^V \leq \widehat{y}$ from Proposition 6. On the contrary, the order-up-to level when not rotating tends to be higher in the multi-horizon case, given $\overline{w}_i^V(s^R) + s^R \geq \overline{w}_i(s^R) + s^R$ from Proposition 8. That is, the multi-horizon scenario tends to order more from the supplier in the non-rotation phase. This is driven by the increased cost of rotation in the multi-horizon model: with potential expiration within a horizon, the benefit of rotation as opposed to ordering from supplier is smaller, which indicates an increased cost of rotation.

To sum up, in the multi-horizon scenario, the optimal policy tends to rotate less often and with smaller rotation size, and start rotation later. This is because the multi-horizon model considers the risk of potential expiration within later horizons and thus the associated cost of rotation is higher.

6. Conclusions

In this paper, we investigate an inventory rotation system for a long-life perishable product, and derive the optimal policy for the rotation decisions and the hospital's ordering decisions. We first consider a base model system with only unit rotation cost in one shelf life horizon, and then extend the model to include a capacity constraint and the decreased effective life in a multi-horizon scenario. We show that the system has an appealing policy structure under these extensions. The well-structured optimal policy shows that the rotation system is applicable in order to control expiration and cost.

In the base model, we show that a policy with a rotate-up-to level and an order-up-to level is optimal and that the whole horizon can be divided into two phases: a non-rotation phase followed by a rotation phase. It is intuitive that the rotate-up-to level $\widehat{z}_i$ is increasing with time. Less intuitively, we find that the order-up-to level at the hospital is first decreasing and then constant as time elapses: it is state dependent and decreasing in the non-rotation phase, and is fixed as $\widehat{y}$ in

the rotation phase. The result is driven by the different tradeoffs around the rotation and ordering decisions and is explained by the underlying cost functions. This policy structure is preserved in the capacitated case and the multi-horizon model, though the threshold values are different.

This paper examines the reserve rotation as a centralised system, assuming a nationalised health-care system. However, in many cases the hospital's interest may not be aligned with the reserve's and the hospital would have no incentive to reduce expiration. It would be of interest to consider a decentralised system, and see how the reserve could motivate the hospital to use the reserve items. A challenge with implementing the optimal centralised policy is that the optimal order level when rotating from the reserve should be lower than the level when not using the reserve. In a decentralised system, one would expect the reserve to lower its price to encourage the hospital to use aging reserve products, which would lead to a higher ordering level and contradicts the optimal policy. A possible mechanism for coordination could be as follows: the reserve asks for a higher unit price to induce the hospital to order the optimal order-up-to amount, but offers a lump sum payment at the end if the hospital has ordered at least a certain amount in total.

While this paper considered a single hospital, there is no reason to believe that the results do not extend to multiple hospitals. The demand in each period could be considered as the aggregate demand across all hospitals and, without the inclusion of fixed costs, the policy will be similar. Multiple hospitals with fixed shipping costs would be an interesting and challenging model, particularly if it included the cost savings from delivering from the reserve to multiple hospitals. As products of different types may influence the rotation policy, it is of interest to investigate the performance of rotating products with different cost and demand parameters, e.g., latex gloves versus anti-flu drugs. Another interesting direction would be to investigate the supplier's reaction to rotation. While the supplier could have regular orders from the hospital without rotation, the reserve is now getting part of the hospital's operational orders. Whether these three parties can coordinate to achieve optimal results for the system is still unknown. Further, as this work studied the rotation policy given a predetermined reserve level P, it would be interesting to investigate how to set an appropriate level P, considering its impact on expiration. These are left as topics for future research.

Acknowledgements

The authors thank the anonymous reviewers for their valuable comments on this paper. Particularly, we thank one of our reviewers for referring us to the concept of $L^\natural$-convexity. This enables us to characterise structural properties in Theorem 1 and Corollary 1 (which also apply to the lost sales case) before deriving the actual optimal policy.

The first author acknowledges the funding received from the University of Auckland Doctoral Scholarship and the Transport Research and Educational Trust Board Grant provided by Chartered Institute of Logistics and Transport (CILT) in New Zealand. The research of the second author was supported in part by Te Pūnaha Matatini Centre of Research Excellence.

Appendix A. Variable Substitution to Enable $L^\natural$-Convexity

$L^\natural$-Convexity

The concept of $L^\natural$-convexity, developed by Murota (2003) in discrete convex analysis, was first introduced to the inventory management literature by Lu and Song (2005). Zipkin (2008) used the concept to establish the structure of lost sales inventory models with positive lead time, providing a new approach to the structural analysis. Later, Huh and Janakiraman (2010) extended the concept to serial inventory systems, and Pang et al. (2012) extended it to inventory-pricing models with positive lead time. Using the development of new preservation properties of $L^\natural$-convexity, Chen et al. (2014) analysed the structure of the joint inventory-pricing decision for perishable products.

The property of $L^\natural$-convexity implies that the function is convex and submodular, plus an additional property related to diagonal dominance (Zipkin, 2008). By definition, a function $f : V \to \mathbb{R}$ is called $L^\natural$-convex if the function $\Psi(v, \zeta) = f(v - \zeta \mathbf{e})$ is submodular on $V \times \mathbb{R}^-$. Zipkin (2008) and Chen et al. (2014) provide the conditions for preservation of $L^\natural$-convexity, the monotonicity structure, and bounded sensitivity of variables. In the following, we employ the concept of $L^\natural$-convexity and its properties to prove some general structural results. The main idea is to transform the state and decision variables.

Transformation

Let $v = -z$ and $s^{NR} = -s^R$, then the programme in (4) is transformed into (A.1).

$$\check{J}_i(s^D, s^{NR}) = \min_{-P \le v \le s^{NR}; w \ge s^D} \{\check{g}_i(w, v)\}, \quad i = 1, \cdots, N, \text{ where} \tag{A.1}$$

$$\check{g}_N(w, v) = -c_r v + \alpha c_e (P + v)^+ + c_s(1 - \alpha)(w - v) + L(w - v);$$

$$\check{g}_i(w, v) = -c_r(1 - \alpha)v + c_s(1 - \alpha)(w - v) + L(w - v) + \check{f}_{i+1}(w, v), \quad i = 1, \cdots, N - 1,$$

$$\text{with } \check{f}_{i+1}(w, v) = \alpha \mathbb{E} \check{J}_{i+1}(w - \xi, v).$$

Similarly, define $\check{w}_i(v)$ as Equation (A.2).

$$\check{w}_i(v) = \operatorname*{argmin}_w \{c_s(1 - \alpha)(w - v) + L(w - v) + \check{f}_{i+1}(w, v)\}, \quad i = 1, \cdots, N. \tag{A.2}$$

For this problem, denote the optimal decisions as $\check{w}^*(s^D, s^{NR})$ and $\check{v}^*(s^D, s^{NR})$. Then, we have Theorem A.1 which gives the properties of the cost functions and optimal decisions based on $L^\natural$-convexity.

Theorem A.1. *If $c_r \le \alpha c_e$, then for period $i = 1, \cdots, N$,*

1. *The function $\check{J}_i(s^D, s^{NR})$ is $L^\natural$-convex in (s^D, s^{NR}), the function $\check{g}_i(w, v)$ is $L^\natural$-convex in (w, v) and thus (w, v, s^D, s^{NR}).*
2. *The optimal decision $\check{w}^*(s^D, s^{NR})$ and $\check{v}^*(s^D, s^{NR})$ are nondecreasing in (s^D, s^{NR}), and for any $\omega \ge 0$,*
$$\check{w}_i^*(s^D + \omega, s^{NR} + \omega) \le \check{w}_i^*(s^D, s^{NR}) + \omega, \text{ and } \check{v}_i^*(s^D + \omega, s^{NR} + \omega) \le \check{v}_i^*(s^D, s^{NR}) + \omega.$$
3. *Given v, the optimal decision $\check{w}_i(v)$ is nondecreasing in v and for any $\omega \ge 0$,*
$$\check{w}_i(v + \omega) \le \check{w}_i(v) + \omega.$$

Proof of Theorem A.1. 1. Note that functions $\check{g}_i(w, v)$ have a common part $c_s(1 - \alpha)(w - v) + L(w - v)$ which is $L^\natural$-convex in (w, v). This is because the function $\Psi((w, v), \zeta) = c_s(1-\alpha)(w-\zeta-v+\zeta)+L(w-\zeta-v+\zeta)$ is independent of ζ and so is submodular in $(w, v) \times \zeta$. Thus, to show the $L^\natural$-convexity of $\check{g}_i(w, v)$, we only need to show the other parts are $L^\natural$-convex, which we can do by induction on i.

The result certainly holds for $\check{g}_N(w, v)$: $\check{g}_N(w, v)$ is $L^\natural$-convex in (w, v) because a single variable function is $L^\natural$-convex (Chen et al., 2014); so, $\check{g}_N(w, v)$ is also $L^\natural$-convex in (w, v, s^D, s^{NR}) because it is separable. By Lemma 2 from Chen et al. (2014), $\check{J}_N(s^D, s^{NR})$ is $L^\natural$-convex in (s^D, s^{NR}), because the constraint set $-P \le v \le s^{NR}; w \ge s^D$ is $L^\natural$-convex.

Assume the $L^\natural$-convexity for $\check{J}_{i+1}(s^D, s^{NR})$, then $\check{f}_{i+1}(w, v)$ is also $L^\natural$-convex because $L^\natural$-convexity is preserved by expectation (Zipkin, 2008). Following that, $\check{g}_i(w, v)$ is also $L^\natural$-convex in (w, v) and thus in (w, v, s^D, s^{NR}), and this in turn implies the $L^\natural$-convexity for $\check{J}_i(s^D, s^{NR})$.

2. By Lemma 3 from Chen et al. (2014), the optimal decision $\check{w}^*$ and $\check{v}^*$ are increasing in (s^D, s^{NR}) and the inequalities hold, because $\check{g}_i(w, v)$ is $L^\natural$-convex in (w, v, s^D, s^{NR}).

3. By Lemma 3 from Chen et al. (2014), $\check{w}_i(v)$ is increasing in v and satisfies the inequality, because $\check{g}_i(w, v)$ is $L^\natural$-convex in (w, v). □

Appendix B. Proofs

Proof of Theorem 1. 1. That the function $\check{J}_i(s^D, s^{NR})$ is $L^\natural$-convex in (s^D, s^{NR}) implies that it is submodular on $s^D \times s^{NR}$, that is, submodular on $s^D \times (-s^R)$. So, the function $J_i(s^D, s^R)$ is supermodular on $s^D \times s^R$. Similarly, $g_i(w, z)$ is supermodular on $w \times z$.

2. The monotonicity follows directly from Theorem A.1, and $w^*(s^D, s^R) = \check{w}^*(s^D, s^{NR})$ and $z^*(s^D, s^R) = -\check{v}^*(s^D, s^{NR})$. Thus, the first inequality in Part 2 of Theorem A.1 transforms to $w_i^*(s^D + \omega, -s^{NR} - \omega) \leq w_i^*(s^D, -s^{NR}) + \omega$ which leads to $w_i^*(s^D + \omega, s^R - \omega) \leq z_i^*(s^D, s^R) + \omega$. Similarly, the second inequality transforms to $-z_i^*(s^D + \omega, -s^{NR} - \omega) \leq -z_i^*(s^D, -s^{NR}) + \omega$ which leads to $z_i^*(s^D + \omega, s^R - \omega) \geq z_i^*(s^D, s^R) - \omega$.

3. Given z, it stands $\overline{w}_i(z) = \check{w}_i(-z)$. So, $\overline{w}_i(z)$ decreases with z following from the monotonicity of $\check{w}_i(v)$. Let $v = -z$ in the inequality in Part 3 of Theorem A.1, then we have $\check{w}_i(-z + \omega) \leq \check{w}_i(-z) + \omega$, which leads to $\overline{w}_i(z - \omega) \leq \overline{w}_i(z) + \omega$. □

Proof of Corollary 1. 1. Theorem 1 shows $\overline{w}_i(z)$ is nonincreasing. When $z \leq P$, g_i is continuously differentiable in w, and $\frac{\partial g_i(w,z)}{\partial w}$ is strictly increasing in z. Therefore, $\overline{w}_i(z)$ is strictly decreasing in z. This shows $\overline{w}_i(z)$ is a bijection and so has an inverse function $\overline{w}_i^{-1}(w)$ which is also decreasing.

2. From Part 2 of Theorem 1, $\overline{w}_i(z - \omega) + z - \omega \leq \overline{w}_i(z) + z$ stands for any $\omega \geq 0$, which implies that $\overline{w}_i(z) + z$ increases with z. From the definition of $\overline{w}_i(z)$, we know

$$c_s(1-\alpha) + L^{'}\left(\overline{w}_i(z) + z\right) + \frac{\partial f_{i+1}\left(\overline{w}_i(z), z\right)}{\partial w} = 0. \tag{B.1}$$

As J_i increases with s^D, $\frac{\partial f_{i+1}(w,z)}{\partial w} \geq 0$, so $c_s(1-\alpha) + L^{'}\left(\overline{w}_i(z) + z\right) \leq 0$. Since $c_s(1-\alpha) + L^{'}(\widetilde{y}) = 0$, we have $\overline{w}_i(z) + z \leq \widetilde{y}$. According to Proposition 1, if $z_i = P$, then an order-up-to $\widetilde{y}$ policy is optimal from period $i+1$ onwards. Therefore, we have $\overline{w}_i(P) + P = \widetilde{y}$ for period i.

3. For any $\sigma > 0$, it stands

$$\begin{aligned} J_i(s^D + \sigma, s^R) &= \min_{s^R \leq z \leq P; w \geq s^D + \sigma} \{g_i(w, z)\} \\ &\geq \min_{s^R \leq z \leq P; w \geq s^D} \{g_i(w, z)\} = J_i(s^D, s^R), \end{aligned}$$

Similarly, $J_i(s^D, s^R + \sigma) \geq J_i(s^D, s^R)$. So, $J_i(s^D, s^R)$ is increasing in s^D and s^R. □

Proof of Proposition 2. Since $\overline{w}_N(z) = \widetilde{y} - z$, $J_N(s^D, s^R) = \min_{s^R \leq z \leq P} \left\{g_N(\max\left(\overline{w}_N(z), s^D\right), z)\right\}$. We need to check $g_N\left(\max\left(\overline{w}_N(z), s^D\right), z\right) = (c_r - \alpha c_e)z + \alpha c_e P + c_s(1-\alpha)\max(\widetilde{y}, s^D + z) + L\left(\max(\widetilde{y}, s^D + z)\right)$. Note $\max(\widetilde{y}, s^D + z)$ is convex in z and $\max(\widetilde{y}, s^D + z) \geq \widetilde{y}$. So, the function g_N is convex in z. Then,

$$\frac{\partial g_N}{\partial z} = \begin{cases} c_r - \alpha c_e, & \text{if } \widetilde{y} \geq s^D + z; \\ c_r - \alpha c_e + c_s(1-\alpha) + L^{'}(s^D + z), & \text{if } \widetilde{y} < s^D + z. \end{cases}$$

Given $c_r \leq \alpha c_e$, g_N is minimised at $z = \overline{y} - s^D$. It is straightforward that $\overline{y} \geq \widetilde{y}$ because $c_r - \alpha c_e \leq 0$. Therefore, $z_N^* = \min(\overline{y} - s_N^D, P)$, and then $w_N^* = \max(s_N^D, \widetilde{y} - z_N^*) = \max(s_N^D, \widetilde{y} - P)$. It follows that $y_N^* = z_N^* + w_N^*$ is the median value of $\widetilde{y}, s_N^D + P$, and $\overline{y}$. □

Proof of Theorem 2. To facilitate the proof, we make use of the results in Corollary 1. Also, we provide Lemma O.1 and a Functional Property List in Online Companion. Based on these, we prove Theorem 2 following four steps. A detailed proof is in the Online Companion.

1. Show the optimal policy and the functional properties for period N.
2. The functional properties of the expected optimal cost function $f_{i+1}(\cdot,\cdot)$ are assumed. We then prove that for period i, the optimal rotate-up-to level is less than the threshold value of the rotation decision for period $i+1$, that is, $z_i^* \leq \widehat{z}_{i+1}$.
3. Solve the period i problem confining the range of the rotation decision z to be $s^R \leq z \leq \widehat{z}_{i+1}$, and establish the period i optimal policy based on Lemma O.1.
4. Show that $f_i(w,z)$ satisfies the properties listed in the Functional Property List. □

Appendix C. Lost Sales Case

In the following, we demonstrate that Theorem A.1, Theorem 1 and Corollary 1 also stand for the lost sales case, because the corresponding optimal cost function $\check{J}_i(s^D, s^{NR})$ is still $L^\natural$-convex. With lost sales, the transition of system state among periods differs from the backlogging case. From Equation (A.1), with the same decision variables w and v, the expected cost function changes to $\check{f}_{i+1}(w,v) = \alpha\mathbb{E}\check{J}_{i+1}\left((w-v-\xi)^{+}+v, v\right)$, with all the other functions having the same formulation. Thus, we have Proposition C.1.

Proposition C.1. *With lost sales case in the hospital, the function $\check{J}_i(s^D, s^{NR})$ is $L^\natural$-convex in (s^D, s^{NR}), the function $\check{g}_i(w,v)$ is $L^\natural$-convex in (w,v) and thus (w,v,s^D,s^{NR}).*

Proof of Proposition C.1. The proof is by induction on i and follows the proof of Theorem 4 in Zipkin (2008). The idea is to reformulate the period i problem with two optimisation steps and show the optimal function given a demand ξ, $\bar{\mathcal{K}}_i(w,v|\xi)$, is $L^\natural$-convex. Thus, at the beginning of period i, before demand, $\check{g}_i(w,v) = \mathbb{E}\left[\bar{\mathcal{K}}_i(w,v|\xi)\right]$ is $L^\natural$-convex in (w,v) because $L^\natural$-convexity is preserved by expectation (Zipkin, 2008; Topkis, 1998), and is also $L^\natural$-convex in (w,v,s^D,s^{NR}) because it is separable, which leads to the $L^\natural$-convexity of $\check{J}_i$.

Interested readers should refer to Zipkin (2008) for a detailed analysis. Note that, though Zipkin (2008) assumes linear holding and penalty cost, the assumption is not essential; the result stands as long as the holding and penalty cost functions do not have items multiplied by both a and r as defined in Zipkin (2008). □

Similar to the backlogging case, Theorem 1 and Corollary 1 imply that the decisions w and z are cost substitutable and the ordering decision y and the rotation decision z are cost complementary.

However, we could not derive the optimal policy for the lost sales case as is for the backlogging case, because the state transition functions are different. The proof for the backlogging case in Theorem 2 largely depends on that $f_i(w,z)$ could be reduced to $f_i(w)$ under certain conditions, but unfortunately this is not true in the lost sales case. As $f_{i+1}(w,z) = \alpha\mathbb{E}J_{i+1}\left((w+z-\xi)^{+}-z, z\right)$ for the lost sales case, it means even though in some cases $J_{i+1}(w,z)$ could be reduced to a single variable function, the expected cost function $f_{i+1}(w,z)$ is still dependent on both variables, and so there is no way to define $\widehat{w}_i$ as we do in the backlogging case. Thus, we could not show the monotonicity of the optimal decisions in different periods as time approaches to the end of the time horizon for the lost sales case.

References

Bradford, J., Sugrue, P., 1991. Inventory rotation policies for slow moving items. Naval Research Logistics 38 (1), 87–105.

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

Chung, Y. T., Erhun, F., 2013. Designing supply contracts for perishable goods with two periods of shelf life. IIE Transactions 45 (1), 53–67.

Civelek, I., Karaesmen, I., Scheller-Wolf, A., 2015. Blood platelet inventory management with protection levels. European Journal of Operational Research 243 (3), 826–838.

Duff, M., 2014. Mass dump of Tamiflu a bitter pill to swallow. http://www.stuff.co.nz/national/ health/9768461/Mass-dump-of-Tamiflu-a-bitter-pill-to-swallow, last accessed on Apr 12, 2014.

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

Hanssmann, F., 1962. Operations research in production and inventory control. John Wiley and Sons, Inc., New York.

Huh, W. T., Janakiraman, G., 2010. On the optimal policy structure in serial inventory systems with lost sales. Operations Research 58 (2), 486–491.

Kendall, K., Lee, S., 1980. Formulating blood rotation policies with multiple objectives. Management Science 26 (11), 1145–1157.

Lu, Y., Song, J.-S., 2005. Order-based cost optimisation in assemble-to-order systems. Operations Research 53 (1), 151–169.

Maddah, B., Yassine, A., Salameh, M., Chatila, L., 2014. Reserve stock models: Deterioration and preventive replenishment. European Journal of Operational Research 232 (1), 64–71.

Murota, K., 2003. Discrete convex analysis. SIAM.

Nahmias, S., 1975. Optimal ordering policies for perishable inventory II. Operations Research 23 (4), 735.

Nahmias, S., 1977. Comparison between two dynamic perishable inventory models. Operations Research 25 (1), 168–172.

Nahmias, S., 2011. Perishable inventory systems. Springer, New York.

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

Nahmias, S., Schmidt, C., 1986. An application of the theory of weak convergence to the dynamic perishable inventory problem with discrete demand. Mathmatics of Operations Research 11 (1), 62–69.

New Zealand Ministry of Health, 2009. National Health Emergency Plan: National Reserve Supplies Management and Usage Policies. New Zealand Ministry of Health, Wellington, 2nd Edition.

Olsson, F., Tydesjö, P., 2010. Inventory problems with perishable items: Fixed lifetimes and backlogging. European Journal of Operational Research 202 (1), 131–137.

Oren, S., Wan, S., 1986. Optimal strategic petroleum reserve policies: A steady state analysis. Management Science 32 (1), 14–29.

Pang, Z., Chen, F. Y., Feng, Y., 2012. Technical note-a note on the structure of joint inventory-pricing control with leadtimes. Operations Research 60 (3), 581–587.

Schmidt, C., Nahmias, S., 1985. $(S-1, S)$ policies for perishable inventory. Management Science 31 (6), 719–728.

Shen, Z., Dessouky, M., Ordonez, F., 2011. Perishable inventory management system with a minimum volume constraint. Journal of Operational Research Society 62 (12), 2063–2082.

Topkis, D., 1998. Supermodularity and Complementarity. Princeton University Press, Princeton.

Veinott, A. J., 1966. On the optimality of (s, S) inventory policies: New conditions and a new proof. SIAM Journal on Applied Mathematics 14 (5), 1067–1083.

Whybark, D. C., 2007. Issues in managing disaster relief inventories. International Journal of Production Economics

108 (12), 228–235.

Woodhead, M., 2014. National medical stockpile dumps $200 million of stock. http://www.australiandoctor.com.au/news/latest-news/national-medical-stockpile-dumps-200m-of-stock, last accessed on Jan 8, 2015.

Zipkin, P., 2008. On the structure of lost-sales inventory models. Operations Research 56 (4), 937–944.