Perishable Inventory Management System With A Minimum Volume Constraint

The federal government maintains large quantities of medical supplies in stock as part of its Strategic National Stockpile (SNS) to protect the American public in case of a public health emergency. Managing these large perishable inventories effectively can help reduce the cost of the SNS and improves national security. In this work, we propose a modified Economic Manufacturing Quantity (EMQ) model for perishable inventory with a minimum volume constraint. We identify when this model is necessary and show that minimizing the cost of maintaining such a system can be formulated as a non-convex non-smooth unconstrained optimization problem. We present an efficient exact algorithm to solve this problem. We illustrate the model and algorithm with numerical results for an example managing stocks of Cipro, an antibiotic to treat people for anthrax infection. We demonstrate the advantage of our proposed model over a standard model and perform sensitivity analysis on the government-controlled system parameters. We show that the government can obtain lower costs or a larger stockpile at the same cost by allowing more freedom in the management of the stocks.

[Keywords: Perishable Inventory Management, EMQ, Emergency Response, non-convex unconstrained optimization problem]

1. Introduction

In most inventory systems, it is assumed that stock items can be stored indefinitely to meet future demands. However, the effects of perishability cannot be ignored for certain types of inventories, which may become partially or entirely unsuitable for consumption as time passes. Typical examples are fresh produce, blood cells, chemicals, photographic films, drugs and other pharmaceuticals. In this work, we investigate a perishable inventory management system with a constant market demand rate in a production environment with a minimum inventory volume ($I_{\min}$) requirement that must be kept at all times.

This work is motivated by the supply chain management for a large-scale emergency response. As part of the national emergency preparedness plan, the federal government maintains a Strategic National Stockpile (SNS). About 20% of the SNS is ready to deploy as push packages and another 80% is in the form of Vendor Managed Inventories (VMIs). For example, in a potential anthrax attack, the stockpile contains enough medicine to treat 10 million people. This stockpile represents enough Cipro, a common antibiotic with a 9 year lifespan that works against anthrax and other infections, to meet regular market demand for several years. The federal government pays pharmaceutical companies to produce and store these large inventories, keeping them ready for use at a moments notice in case of an emergency.

Currently, the SNS policy allows manufacturers to sell the pills at a predefined date prior to expiration rather than let the drugs spoil; however, considering the fact that the size of the stockpile is huge compared with the regular market demand while the drugs are so close to their expiration date, the potential salvage value is low. From the manufacturer's perspective, if it can apply a more sophisticated inventory holding policy which allows the constant usage of the stockpile to meet the regular market demand and refill with new production at the same time to maintain the minimum stockpile requirement, then the firm can save on the total cost in maintaining the stockpile inventory, hence making it possible to further reduce the price charged to the government. From the government's perspective, if it allows firms to sell the pills earlier, there is an opportunity to capture a significant amount of salvage value for the unsold stockpile. The unique challenge of this problem lies in efficiently maintaining a minimum level of perishable inventory.

Since the stockpile contains perishable items which may never be called to use by the federal government (in cases where no terrorist attack occurs), the key challenge from both the government's and manufacturer's perspective is determining (1) how often the stockpile is refreshed and released to the open market, (2) what is a suitable cost effective minimum inventory requirement, and (3) how much should the government pay to the manufacturer for each pill stored in the stockpile. Through sensitivity analysis, we demonstrate how decision makers can use the proposed model to set policy, and illustrate the possiblity of reducing the cost to the government for the same level of VMI by leveraging the regular market demand.

The traditional EMQ (Economic Manufacturing Quantity) model can be readily extended to address the perishability property of a stockpile with no minimum inventory requirement by properly upper-bounding the EMQ cycle to guarantee that the inventory is consumed within its shelf-life. In the case that the minimum inventory required is significantly smaller than the total regular market consumption over the shelf-life of the drug, a tighter limit on the EMQ cycle that also takes into account the freshness of the minimum inventory is sufficient. However, when the minimum inventory is comparable with the total regular market consumption during the shelf-life, a trivial extension to any of the existing perishable inventory policies is no longer adequate. It is therefore imperative to develop a new inventory policy specially geared to the perishable VMI system for the SNS, which not only satisfies the minimum inventory requirement but also minimizes the operational cost of maintaining such a system by incorporating the regular market demand. Hence, in this research, we aim to propose a single inventory system which satisfies the two types of demand: the regular market demand and the minimum inventory requirement for emergency preparedness, which minimizes the operational cost from the manufacturer's perspective.

In this paper, we first review the relevant literature in section 2. In section 3 we discuss the assumptions and policies we adopt and study a straightforward extension of the EMQ model on perishable items with zero or a "small" minimum inventory requirement. Next, we propose a modified EMQ model for perishable items with a minimum inventory constraint. Then we show that to guarantee the freshness of the stockpile a maximum inventory cycle constraint is necessary. In section 4, we present the detailed calculation on the total cost and boundary conditions. For this we decompose the cost of such a system into four components: inventory holding costs, fixed ordering costs, purchasing costs and salvage costs. We can express the four parts of the total cost as a non-convex and non-smooth function of Q. Section 5 covers the exact solution approach and its complexity analysis. We conduct two different numerical experiments on an anthrax attack example in section 6. The first experiment demonstrates the advantage of our proposed model over a standard model, which runs two separate systems to meet the regular market demand and the minimum inventory requirement respectively. In the second experiment, we perform sensitivity analysis on those government controlled system parameters to provide some insights for both parties (the firm and the government) in how to negotiate contract terms. Finally we present some concluding remarks in section 7.

2. Literature Review

In the existing perishable inventory management literature, policies in four different aspects have attracted attention from the research community.

- *Ordering Policy* focuses on *when* and *how much* to order; a well known review is from Nahmias (1982).

- *Issuing Policy* concerns the sequence in which items are removed from a stockpile of finitely many units of varying ages; the most general approach is FIFO and it is proved to be optimal for perishable goods with random supply and demand and fixed life-time under several possible objective functions by Pierskalla and Roach (1972)

- *Disposal Policy* is applied when the strategic disposal of part of the inventory is desirable, such as slow moving stock; the topic on when and how much to dispose under stochastic demand and perishing has been studied by Rosenfield (1989, 1992).

- *Pricing Policy* is closely coupled with the ordering policy in the multi-period newsvendor problem. The price is a decision variable and the forecasted demand is price-sensitive. The pricing and ordering quantity decision can be made either sequentially or simultaneously (Gallego and Van Ryzin, 1994; Abad, 1996; Burnetas and Smith, 2000; Chun, 2003).

For the perishable inventory system, the *ordering policy* is the most researched policy of the four above. Research has been done assuming: fixed or continuously deteriorating lifetime; periodic or continuous review; different distributions of the demand process; lost sale or backorder; etc. Based on different sets of assumptions, various modeling method and solution approaches have been applied.

1. Zero first order derivative point (stationary point) over total cost function: It is usually straightforward to write the governing equation on the inventory level over time and then obtain the inventory carrying cost, along with the fixed ordering cost, the purchasing cost, the shortage cost, and the outdating cost. This can be modeled as an unconstrained nonlinear system. If the total relevant cost function is continuous and second-order differentiable over the decision variable of the current system (in most cases, the ordering quantity or reorder level), then we can obtain the first-order stationary point as the optimum (Ravichandran, 1995; Giri and Chaudhuri, 1998; Liu and Lian, 1999).

2. Heuristics/Approximations: For stochastic demand circumstances, exact optimal policies are not only difficult to compute, but also demanding to implement due to the requirement to keep track of the age distribution of the stock. Nahmias (1982) provided a good summary on the early works of approximated optimal policies and heuristics to obtain the optimal policy parameters. Nandakumar and Morton (1993) developed heuristics from "near myopic" bounds and demonstrated the accuracy with less than 0.1% average error over a wide range of problems. Goh et al. (1993) applied different approximation methods to study a two-stage perishable inventory models.

3. Markovian model: The queuing model with impatient customers has been used to analogize perishable inventory systems. The queue corresponds to the inventory stockpile, service process to the demand, arrival of customers to the replenishment of inventory, and the time a customer will stay in queue before leaving due to impatience corresponds to the shelf-life. Early works (Chazan and Gal, 1977; Graves, 1982) usually wrote the descriptive transition probabilities then obtained the stationary distribution to evaluate the performance measures such as the expected outdating amount. Weiss (1980) and Liu and Lian (1999) also use the performance measures to construct a cost function from which the optimal policy parameter can be computed.

4. Dynamic programming (DP): Since the dynamic economic lot size problem was first proposed by Wagner and Whitin (1958), which reviews the inventory periodically and the demand is deterministic in every review period, dynamic programming techniques have been widely adopted in solving many variations and extensions in the non-perishable inventory context. Hsu (2000) provided a brief review on this topic. There is a well-known zero-inventory property under which no inventory is carried into a production period that is necessary for an optimal solution. Hsu (2000) demonstrated that this zero-inventory property may not hold for any optimal solution with perishable items and proposed a new DP recursion based on an interval division property that solved the problem in polynomial time. Further extension to perishable systems which allow backorder and co-existed stochastic and deterministic demand can be found in Hsu and Lowe (2001) and Sobel and Zhang (2001).

5. Fuzzy theory: Recently, Katagiri and Ishii (2002) introduced fuzzy set theory in a perishable inventory control model with a fuzzy shortage cost and a fuzzy outdating cost; hence the expected profit function is represented with a fuzzy set. The effect of the fuzziness on the obtained ordering quantity is investigated.

There is abundant literature in perishable inventory management to model and solve different types of real-life problems. To the best of our knowledge, there is no prior work with a minimum volume constraint on the inventory size throughout the planning horizon. This extension is trivial when the required minimum volume is not significant compared with the amount consumed by a regular market demand rate within the shelf life as it can be timely and completely refreshed by the regular demand. However, when the minimum volume is huge —comparable with the total regular market demand over the shelf life— then a strategic inventory ordering and disposing policy is needed to guarantee the freshness and readiness of the required minimum inventory as well as the low cost of maintaining such an inventory system. This is the case of the medical stockpiles required for SNS in the large-scale emergency context. In this work, we address this modeling gap by formulating a perishable inventory management model with minimum inventory constraint and providing an exact solution approach to this model.

3. Model

3.1 EMQ Model with Perishability

We assume a single fixed-life perishable item is produced, consumed and stored for an infinite continuous time horizon. We denote T_s as the shelf-life of the inventory (in years). The regular market demand is known with a priori constant rate D per year. The production can start at any time at a constant rate P, which is greater than D; and there is a constant cost A associated with each production setup. The holding cost h, unit purchase cost v and unit salvage cost w are all time invariant. At any point of time, at least a pre-determined non-spoiled minimum inventory volume ($I_{\min}$) must be maintained.

For non-perishable items and no minimum inventory, the classical EMQ model provides an optimal inventory management policy for this problem. This cyclic solution orders a fixed amount, known as the production batch size, Q every $T = Q/D$ units of time. The EMQ cycle T begins with a production phase that lasts $T_1 = \frac{Q}{P} = \frac{D}{P}T$ and is followed by an idle phase lasting $T_2 = \frac{P-D}{P}T$. The optimal batch size of the EMQ model EMQ^* is identified by minimizing a convex inventory cost function.

The proposition below shows that this model can be directly applied to the perishable stockpile with a minimum inventory requirement by adding a constraint to ensure the age of the inventory does not reach expiration. Before we prove the general case, consider a perishable inventory system with $I_{\min} = 0$. The idea is that, if the inventory is consumed following a FIFO policy, the items consumed by time T_1 are produced up to time $T" = \frac{D}{P}T_1$. We obtain an upper bound on the age of the inventory assuming that the next item produced is kept in inventory for the duration of the cycle, which gives an age of $T - T" = \frac{P^2-D^2}{P^2}T$. Ensuring that this upper bound is less than the shelf-life T_s guarantees that the EMQ policy is valid.

Proposition 1. *An EMQ model for a perishable inventory with minimum inventory requirement of $I_{\min} \leq DT_s$ has an optimal production batch size given by $Q^* = \min(EMQ^*, T_s\frac{DP}{P-D} - I_{min}\frac{P+D}{P-D}, (T_s - \frac{I_{\min}}{P})\frac{DP^2}{P^2-D^2})$.*

Proof. When $I_{\min}$ is non-zero and in a small amount, as in Fig. 1, with a FIFO policy, an $I_{\min}$ amount is first consumed to meet the market demand over the period $T_{\min} = \frac{I_{\min}}{D}$.

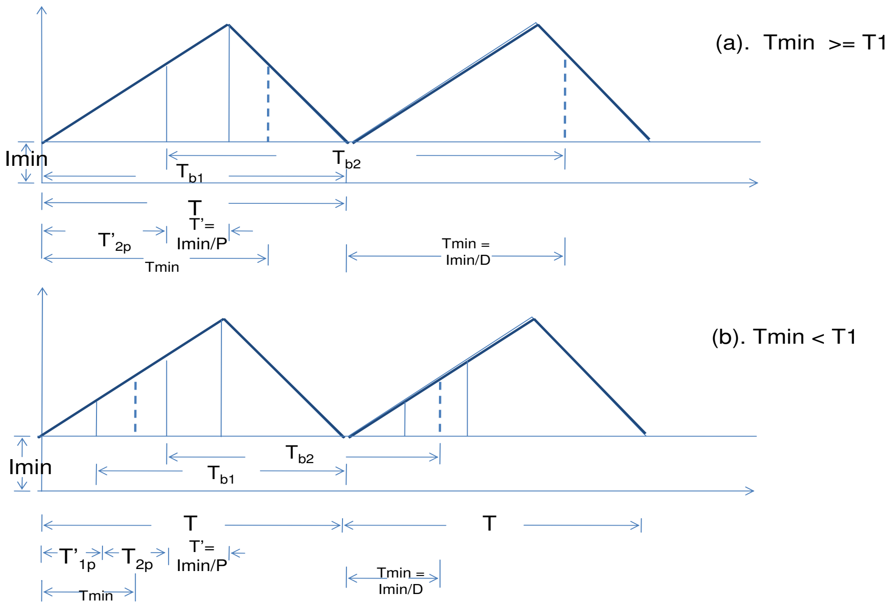

Figure 1: Perishable Inventory System with small $I_{\min}$.

There are two possible upper bounds of the oldest item: one is T_{b1} which is the oldest item consumed at the end of each T; another is T_{b2} which is the oldest item in I_{min} before it completely depletes at the beginning of each T cycle. In Fig. 1 part (a), $T_{\min}$ is longer than T_1. T_1 can be decomposed into two parts: $T' = \frac{I_{\min}}{P}$ which produces the $I_{\min}$ amount consumed at the beginning of the next cycle and $T'_{2p} = \frac{(T-T_{\min})D}{P}$. Therefore, the upper bound of the oldest item should be $T_{b2} = T' + T_2 + T_{\min} = \frac{P-D}{P}T + \frac{P+D}{PD}I_{\min} \leq T_s$. In Fig. 1 part (b), $T_{\min}$ is shorter than T_1 which can be decomposed into three parts: T' with the same definition as in part (a), T_{2p} that produces the items consumed in T_2 period and $T'_{1p} = \frac{(T_1-T_{\min})D}{P}$ that produces the items consumed in T_1 after depleting $I_{\min}$ from the previous cycle. Hence, it shares the same T_{b2} bound as in part (a), we also have another upper bound of the oldest item $T_{b1} = T - T'_{1p} = \frac{P-D}{P}\frac{P+D}{P}T + \frac{I_{\min}}{P} \leq T_s$. Therefore, the *optimal production batch size* for the perishable items with $I_{\min}$ minimum inventory requirement is $Q^* = min(EMQ^*, T_s\frac{DP}{P-D} - I_{min}\frac{P+D}{P-D}, (T_s - \frac{I_{\min}}{P})D\frac{P^2}{P^2-D^2})$. At the same time, $I_{\min}$ is upper bounded by $T_{\min} \leq T_s$, which is $I_{\min} \leq T_s D$. □

This quantifies the situation where a simple extension of the classical EMQ model is applicable to a perishable inventory system. However, when such condition cannot be met, a more sophisticated model is required, which is the focus for the rest of this paper.

3.2 Modified EMQ Model

In this subsection, we introduce the modified EMQ model which incorporates both regular market demand and emergency demand ($I_{\min}$) for a perishable stockpile. We consider an EMQ like model because of the simplicity and wide use of these types of models.

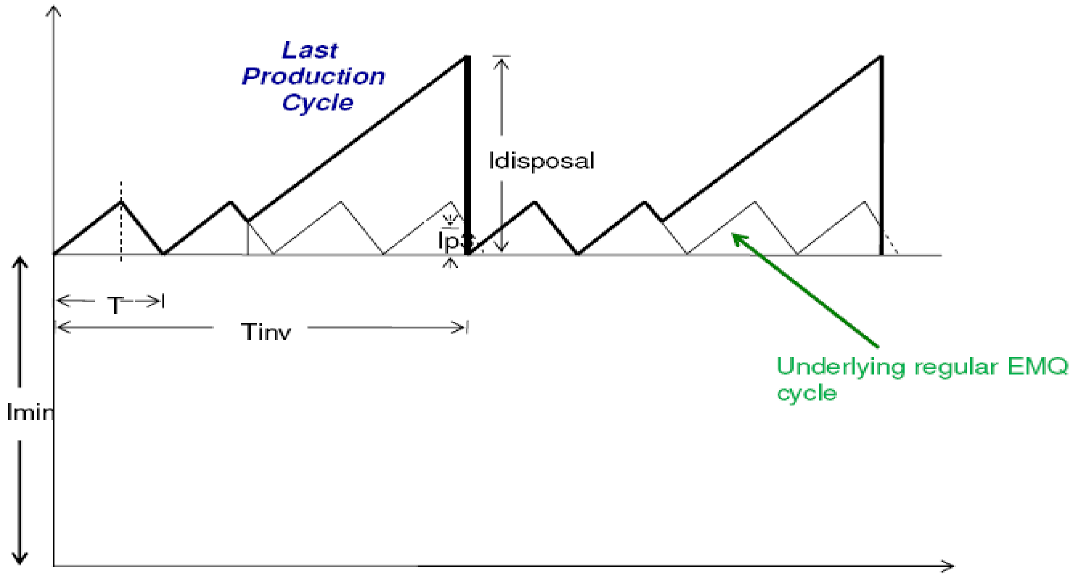

Figure 2: Illustration for the Modified EMQ Model.

Fig. 2 gives an illustration of the inventory plot for the modified EMQ model we propose. We define the *inventory cycle* (T_{inv}) as the minimum length of time that an inventory pattern repeats. We only allow disposing once at the end of each T_{inv}, at which we dispose all the inventory above the $I_{\min}$ level so that the exact $I_{\min}$ is reached at the beginning of the next T_{inv}. We use a FIFO issuing policy such that we always use the oldest items to satisfy the market demand to dispose.

We propose the following ordering policy, illustrated in Fig. 2 to service the regular market demand and maintain a fresh minimum inventory. For any given Q, we initially run a regular EMQ cycle (we call them the "underlying regular EMQ cycles" with cycle length T, where $T = \frac{Q}{D}$) and make some adjustment near the end of the *inventory cycle*. To maintain a fresh inventory in the system, every *inventory cycle*, an $I_{\min}$ amount of inventory must be produced and consumed, either by means of regular market demand or disposal at the end of T_{inv}. These conditions are summarized in the following two equations, which we refer to as the *stability condition* of the system and will be discussed in more detail in the next section.

From a consumption perspective, the $I_{\min}$ level must be greater than the demand rate times T_{inv} to guarantee the necessity of this model.

$$I_{\min} > D \cdot T_{inv}. \tag{1}$$

Otherwise, a traditional EMQ model is sufficient to solve the problem by depleting/refreshing the $I_{\min}$ amount within T_{inv} with the regular market demand rate as presented in section 3.1.

To satisfy the *stability condition* from a production perspective, we require that the $I_{\min}$ level must be less than the production rate times T_{inv} to guarantee production feasibility.

$$I_{\min} \leq P \cdot T_{inv}. \tag{2}$$

In order to produce an $I_{\min}$ amount within T_{inv} while satisfying constraint 1, a *last production cycle* may be required to produce the amount to be disposed at the end of an *inventory cycle*. A similar formula is also stated in the Proposition 1 as the boundary condition between the simple EMQ extension and a more sophisticated system (T_s is used in Proposition 1 in stead of T_{inv} here).

Given the *production batch size* Q, we show in the next section that we can determine all relevant quantities to describe this model. In particular these quantities include, the length of the regular underlying EMQ cycle (given by $T = Q/D$), when to initiate the *last production cycle*, and the length of this *last production cycle*. Therefore, we have a system with a single independent decision variable Q and an unique inventory plot corresponding to the given Q. A total cost relevant to maintaining such a system can be computed with the aid of the inventory plot for a given Q.

Fig. 2 provides an illustrative example on the inventory plot. It demonstrates one case out of the 5 different possible cases. Next, we discuss the 5 cases (see Fig. 3) and the classification criteria.

The classification is based on three criteria:

1. if a *last production cycle* is needed to replace the extra disposing part of the inventory;

2. where the *inventory cycle* ends relative to a regular underlying EMQ cycle (in the production period – the uphill region in the inventory graph, or in the non-production period – the downhill region);

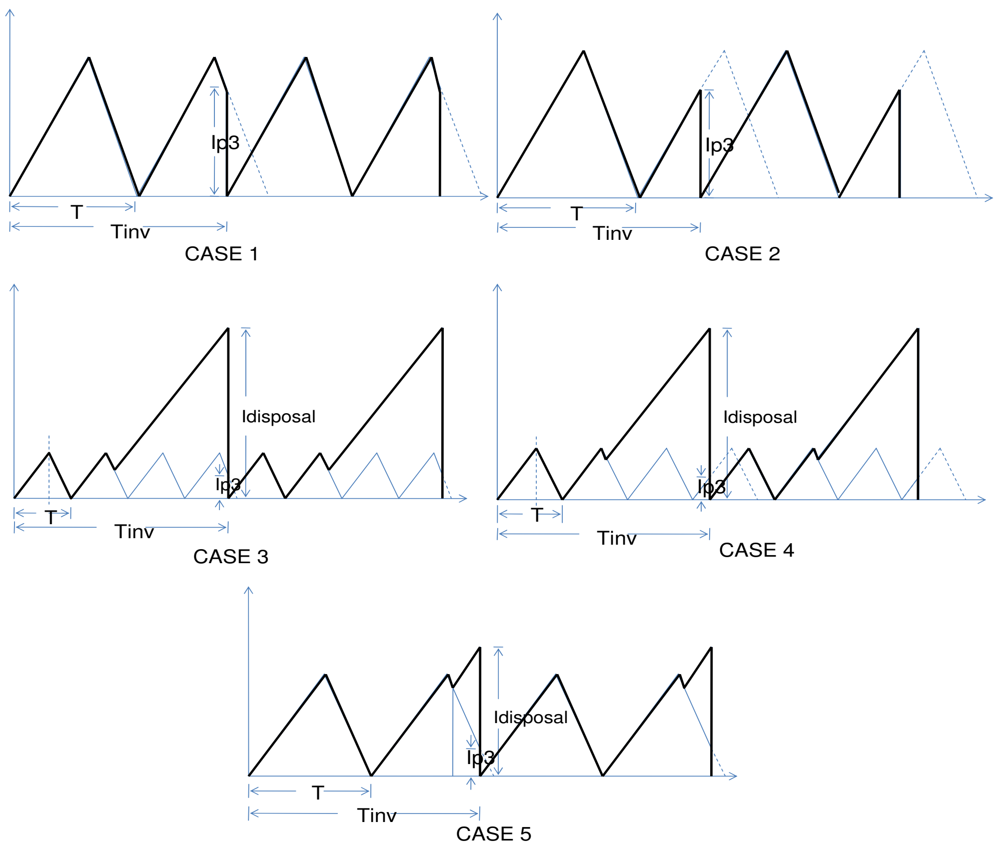

Figure 3: All 5 Cases of Possible Scenarios.

3. where the *last production cycle* starts (if it starts at the same downhill region of a regular underlying EMQ cycle as it ends or not).

The disposing policy of the model requires that we dispose all the inventory above the $I_{\min}$ level at the end of each *inventory cycle* to restart a new cycle with exactly $I_{\min}$. The amount to be disposed is either the part of the minimum inventory ($I_{\min}$) which cannot be consumed by the regular market demand within an inventory cycle (which is $I_{disposal} = I_{\min} - DT_{inv}$) or the part of the inventory above the minimum inventory volume requirement generated by a regular EMQ cycle at the end of an *inventory cycle* (defined as I_{p3}), whichever is bigger. Cases 1 and 2 in Fig. 3 depict the situation that I_{p3} is greater than $I_{disposal}$ and we dispose I_{p3} without a *last production cycle*. Cases 3, 4 and 5 fit the situation where I_{p3} is less than $I_{disposal}$ and we initiate a *last production cycle* to produce the amount $I_{disposal} - I_{p3}$ and dispose exactly $I_{disposal}$ at the end of each T_{inv}. We further classify the cases based on where an *inventory cycle* ends relative to the underlying EMQ cycle. In cases 2 and 4, the *inventory cycle* ends at an uphill region (the production period); otherwise, in cases 1, 3 and 5, the cycle ends in a downhill region (the non-production period). Furthermore, the 3 cases which have a *last production cycle*, are further classified depending on where the last production cycle starts. If it starts at the same non-production period as it ends, it is case 5; otherwise, if it starts at some earlier regular EMQ cycle, we have cases 3 and 4. With the above three criteria, we can distinguish these 5 cases uniquely. These enumerate all possible scenarios.

3.3 Constraint on the Inventory Cycle Length Parameter

In this section, we present a crucial constraint on one important parameter (T_{inv}) in this perishable inventory system, to maintain the whole stockpile within the shelf life throughout the time horizon. Recall that the *stability condition* of the system states that to guarantee the freshness of the stockpile, the system requires refreshing the entire minimum inventory volume within each T_{inv}. That is, for each *inventory cycle*, an $I_{\min}$ amount is consumed and another $I_{\min}$ amount is produced. The amount of the minimum inventory produced in one cycle is completely consumed in the next cycle. At the beginning of each cycle (which is also the end of its previous cycle), the system reaches its minimum inventory volume requirement level by disposing any amount above it. The system runs a regular EMQ model with extra production in the *last production cycle* to produce the remaining amount of stock that cannot be consumed by the market demand and needs to be disposed at the end.

Since we are adopting a FIFO issuing policy, the $I_{\min}$ amount produced in one *inventory cycle* will not be used until the next cycle starts. We assume a continuous and infinite time horizon, and we have a continuous age-distribution for the $I_{\min}$ amount produced in an *inventory cycle* whose age is between 0 and T_{inv} at the beginning of a cycle. The following lemma states a constraint on T_{inv} that is able to guarantee the freshness of the stockpile.

Lemma 1. *Lemma I: (Maximum Inventory Cycle Length) The maximum* inventory cycle *length is at most a half shelf-life $T_{inv} \leq \frac{1}{2}T_s$; this guarantees that*

- *the complete stockpile is always within its shelf-life;*
- *$I_{\min}$ is always younger than half the shelf-life and the salvaged items are aged between the half shelf-life and full shelf-life;*
- *the age distribution of the stockpile repeats itself every T_{inv}.*

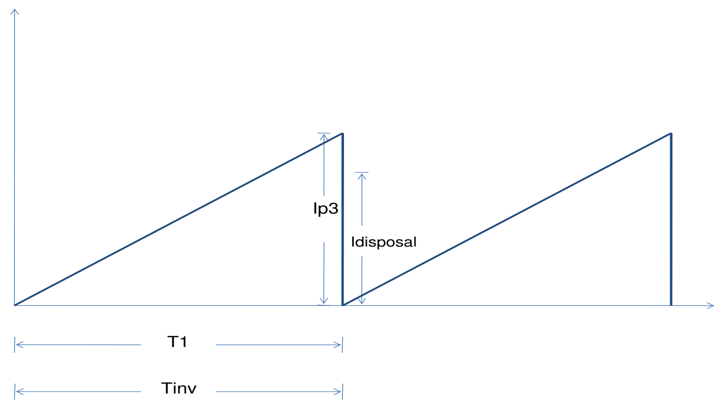

Figure 4: Graph Illustration for an Extreme Boundary Case.

Proof. We prove this by an extreme boundary case. Let us consider the extreme boundary case that the $I_{\min}$ is exactly equal to the production rate times the *inventory cycle*, which means that we have to keep producing all the time to guarantee the amount $I_{\min}$ can be refreshed every T_{inv} (see Fig. 4). That is

$$I_{\min} = P \cdot T_{inv}. \tag{3}$$

Without loss of generality and for the convenience of the proof, we discretize the time horizon. That is, we divide an *inventory cycle* into T_{inv} discrete periods from 0 to $T_{inv} - 1$. Then we can express the age distribution of the $I_{\min}$ at the end of an inventory cycle (we call it cycle 1) as follows, which represents that P units of inventory have been produced at every discrete period in cycle 1:

$$Age = \begin{cases} 0 & \text{P units;} \\ 1 & \text{P units;} \\ ... & ...; \\ T_{inv} - 1 & \text{P units.} \end{cases} \tag{4}$$

We observe the age distribution when it moves to the next inventory cycle (we call it cycle 2). At time 0 in cycle 2, we have P units of age T_{inv} which were produced at time 0 in cycle 1 and we can use D units of them to fulfill the market demand. For this proof, we will deviate from our FIFO service (issuing) policy and leave the $P - D$ unused units in stock. At time 1 in cycle 2, we again get P units of age T_{inv} which were produced at time 1 in cycle 1; same as time 0, we can use D units for the market demand and leave $P - D$ units unused in stock. We can continue this process along the discrete time horizon until $T_{inv} - 1$ in cycle 2. At time $T_{inv} - 1$ in cycle 2, before we dispose the unconsumed inventory produced in cycle 1, we have a total of $T_{inv} * (P - D)$ units of stock to be disposed, with the following age distribution:

$$Age = \begin{cases} T_{inv} & \text{for P-D units;} \\ T_{inv} + 1 & \text{for P-D units;} \\ ... & ...; \\ 2T_{inv} - 1 & \text{for P-D units.} \end{cases} \tag{5}$$

We use the upper bound on cycle length $2T_{inv} \leq T_s$ to ensure that all stocks left after disposing have the age distributed in $[0,\ T_s]$. We claim that this is the upper bound since we break the FIFO issuing policy in the process described above. Using a FIFO policy, we would have a younger inventory to dispose but the first $I_{\min}$ will still have the same age distribution. If we abide with the FIFO issuing policy, at time 1 in cycle 2 we would first use the oldest left-over $P - D$ units of age $T_{inv} + 1$, which is produced at time 0 in cycle 1, to satisfy the market demand before we use the age T_{inv} batch, which is produced at time 1 in cycle 1. Hence, we can claim that the oldest age of the disposal units at the end of cycle 2 will be at most $2T_{inv} - 1$. Thus, a half shelf life is the upper bound for the *inventory cycle* to guarantee the freshness of the stock in the minimum inventory volume ($I_{\min}$ is always younger than half shelf-life) as well as the potential salvage value (the age of salvaged items is always between half shelf-life and full shelf-life). □

4. Total Cost Evaluation and Boundary Conditions

In this section, we first introduce the notation, then discuss the calculation steps required by all 5 cases. In all 5 cases, the total cost depends continuously only with respect to Q.

4.1 Notation

We first introduce the notation used in the calculation. We will continue using the parameters we defined previously and the variables below are notation specially used in our proposed model.

I_{max}:	the maximum inventory level in a regular EMQ cycle
$I_{disposal}$:	the minimum amount to be disposed every T_{inv}
N:	number of complete regular EMQ cycles in a T_{inv}
T_{p3}:	the remainder of T_{inv} divided by T
I_{p3}:	the inventory level of a regular underlying EMQ at the end of a T_{inv}

$T_{disposal}$: production time for the extra inventory to be disposed, which is $max(I_{disposal} - I_{p3}, 0)$

δ: the non-production time in T_{p3}

T_{p1}: from the start of the *last production cycle* to the end of the current underlying EMQ cycle

I_{p1}: the inventory level at the beginning of the *last production cycle*

N_1: number of complete regular EMQ cycles within the *last production cycle*

M: number of regular EMQ orders in a T_{inv}

4.2 Cost Decomposition

Now we are ready to calculate the inventory level for the different cases to prepare the total cost computation. We use the same T, T_1 and T_2 formula as defined in section 3.1. Below are the basic quantities which share the same formula across all 5 cases.

$$I_{max} = Q(1 - \frac{D}{P}) \tag{6}$$

$$I_{disposal} = I_{\min} - T_{inv} \cdot D \tag{7}$$

$$N = \lfloor \frac{T_{inv}}{T} \rfloor \tag{8}$$

$$T_{p3} = T_{inv} \% T = T_{inv} - N \cdot T \tag{9}$$

$$I_{p3} = \begin{cases} D \cdot (T - T_{p3}) & \text{for cases 1, 3, 5; where } T_{inv} \text{ ends at non-production time;} \\ (P - D) \cdot T_{p3} & \text{for cases 2, 4; where } T_{inv} \text{ ends at production time.} \end{cases} \tag{10}$$

For cases 3, 4 or 5, we have the formula for the non-production time in the last incomplete EMQ cycle (T_{p3}):

$$\delta = \max(T_{p3} - T_1, 0) \tag{11}$$

Since we assume that at the end of each *inventory cycle*, the inventory level is set back to $I_{\min}$. At T_{inv}, the inventory level I_{p3} is less than the required disposal amount $I_{disposal}$ for cases 3, 4 and 5. The time that is needed to produce the extra disposal amount ($I_{disposal} - I_{p3}$) is:

$$T_{disposal} = \frac{I_{disposal} - I_{p3}}{P} \tag{12}$$

Note that ($I_{disposal} - I_{p3}$) is just part of the amount produced in the last production cycle which cannot be covered by the production periods in a regular underlying EMQ cycle within an *inventory cycle*. Hence $T_{disposal}$ only occupies part of the last production cycle with the remaining time used to produce items to meet the regular market demand. Another way to express the classification of cases 3, 4 and 5 by δ and $T_{disposal}$ is: if $\delta = 0$, it is case 4; else when $\delta > 0$, if $\delta > T_{disposal}$, it goes to case 5 and otherwise to case 3.

We can decompose the total cost TC of maintaining this perishable inventory system within a single *inventory cycle* into 4 parts: inventory holding cost (TC_{inv}), fixed ordering cost (TC_N), unit purchase cost ($TC_{purchase}$) and salvage cost on the disposal part ($TC_{Salvage}$). That is:

$$TC = TC_{Inv} + TC_N + TC_{Purchase} + TC_{Salvage} \tag{13}$$

We first look at the computation on the purchase cost and salvage cost. The total purchase cost is the amount to produce within an *inventory cycle* times the unit price. The total salvage cost is the amount to dispose at the end of each *inventory cycle* times the unit salvage value. Hence we have the following formulas:

For cases 1 and 2:

$$TC_{Purchase} = (I_{\min} + I_{p3} - I_{disposal}) \cdot v \tag{14}$$

$$TC_{Salvage} = I_{p3} \cdot w \tag{15}$$

For cases 3, 4 and 5:

$$TC_{Purchase} = I_{\min} \cdot v \tag{16}$$

$$TC_{Salvage} = I_{disposal} \cdot w \tag{17}$$

Note that for cases 3, 4 and 5, $TC_{Purchase}$ and $TC_{Salvage}$ are fixed and independent of the *production batch size* (Q); hence they can be removed from the total relevant cost calculation.

Next, we look at the computation on the total ordering cost and inventory holding cost, which are more complicated. We first give the general calculation formula here and then expand them with respect to Q for each case later.

For cases 1 and 2, since there is no *last production cycle*, the number of orders is the number of complete regular EMQ cycles in an *inventory cycle* plus 1. And the total inventory carrying cost can be calculated by the area under the inventory plot. For case 1 where the T_{inv} ends in the downhill region, the area under the inventory plot would be $N + 1$ regular EMQ triangles minus the cut-off small triangle in the shadow, in Fig. 5. For case 2 where the T_{inv} ends in the uphill region, the area would be N regular EMQ triangles plus the small extra triangle in the shadow, in Fig. 12. The formula is as follows:

$$TC_N = (N + 1) \cdot A \tag{18}$$

$$TC_{Inv} = \begin{cases} \frac{1}{2}(N+1)\cdot T \cdot I_{max} - \frac{1}{2}(T - T_{p3}) \cdot I_{p3} & \text{for case 1;} \\ \frac{1}{2}N \cdot T \cdot I_{max} + \frac{1}{2}T_{p3} \cdot I_{p3} & \text{for case 2.} \end{cases} \tag{19}$$

For cases 3 and 4, we can only use the non-production time (T_2) of regular EMQ cycles to produce the $I_{disposal} - I_{p3}$ amount (production time of the EMQ cycle is already in use to satisfy the regular demand D). The number of complete EMQ cycles that would be covered by the last production cycle is:

$$N_1 = \lfloor \frac{T_{disposal} - \delta}{T_2} \rfloor \tag{20}$$

If M is the number of complete regular EMQ cycles before the last production period starts, then there are $M + 1$ orders per inventory cycle.

$$M = N - N_1 \tag{21}$$

$$TC_N = (M+1) \cdot A \tag{22}$$

In cases 3 and 4 (see Fig. 15 and Fig. 13) the last production cycle must start on a non-production period (T_2) and the time of this last production cycle during the current T_2 is:

$$T_{p1} = (T_{disposal} - \delta)\% T_2 = (T_{disposal} - \delta) - N_1 \cdot T_2 \tag{23}$$

Hence we have the height of the short parallel lateral (E_1E_2) of the trapezoid ($E_1E_2E_3E_4$), which is:

$$I_{p1} = T_{p1} \cdot D \tag{24}$$

The total area under the inventory plot for both cases 3 and 4 is M regular EMQ triangles minus the cut-off small shaded triangle plus the area of the trapezoid. The long parallel lateral is exactly the amount to be disposed (which is the length of E_3E_4 which equals to $I_{disposal}$). Hence the total inventory carrying cost can be expressed as:

$$TC_{Inv} = \frac{1}{2}M \cdot T \cdot I_{max} - \frac{1}{2}T_{p1} \cdot I_{p1} + \frac{1}{2}(I_{p1} + I_{disposal}) \cdot (T_{inv} - (M \cdot T - T_{p1})) \tag{25}$$

For case 5 (see Fig. 14), the last production cycle starts at the same downhill slope as the T_{inv} ends and the area under the inventory plot would be similar to the calculation of cases 3 and 4 except for the width of the cut-off small triangle in the shadow and the width of the trapezoid:

$$N_1 = 0 \tag{26}$$

$$I_{p1} = (T - T_{p3} + T_{disposal}) \cdot D \tag{27}$$

$$TC_N = (N+2) \cdot A = (M+1) \cdot A \tag{28}$$

$$TC_{Inv} = \frac{1}{2}(N+1) \cdot T \cdot I_{max} - \frac{1}{2}(T - T_{p3} + T_{disposal}) \cdot I_{p1} + \frac{1}{2}(I_{p1} + I_{disposal}) \cdot T_{disposal} \tag{29}$$

4.3 Total Cost and Boundary Condition of Case 1

Next, based on the results of the previous subsection, we present the detailed calculation on the total cost as a function with respect to Q and the boundary conditions for case 1. This is a relatively simple case. For a more complicated example, please see appendix B which shows the calculation on case 3. The other cases can be found in Shen (2008).

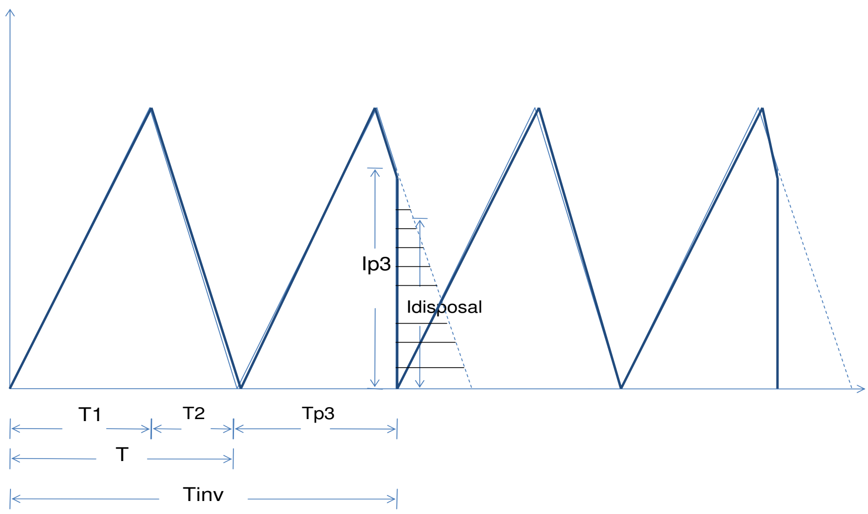

Figure 5: Graph Illustration for Case 1.

$$\begin{aligned} TRC(Q) &= I_{p3} \cdot (v+w) + \frac{1}{2}h[(N+1) \cdot T \cdot I_{max} - (T - T_{p3}) \cdot D(T - T_{p3})] \\ &= (\frac{Q}{D} - T_{inv} + N\frac{Q}{D}) \cdot D \cdot (v+w) + \frac{1}{2}h[(N+1)\frac{Q^2}{D}(1 - \frac{D}{P}) \\ &\quad -D \cdot (\frac{Q}{D} - T_{inv} + N\frac{Q}{D})^2] \end{aligned}$$

We let $TRC(Q) = a \cdot Q^2 + b \cdot Q + c$, then

$$\begin{aligned} a &= \frac{1}{2}h \cdot [(N+1) \cdot (\frac{1}{D} - \frac{1}{P}) - D \cdot \frac{(N+1)^2}{D^2}] \\ &= -\frac{h}{2}\frac{(N+1)}{P} \cdot (1 + N \cdot \frac{P}{D}) \end{aligned}$$

$$\begin{aligned} b &= (N+1)\cdot(v+w) + h\cdot T_{inv}\cdot(N+1) \\ &= (N+1)\cdot[(v+w)+h\cdot T_{inv}] \end{aligned}$$

The $TRC(Q)$ is a quadratic function of Q and the stationary point of the quadratic curve is (where $a < 0$):

$$Q^*_{case1} = -\frac{b}{2a} = \frac{P[(v+w)+h\cdot T_{inv}]}{h(1+N\frac{P}{D})}$$

The boundary points for case 1 on one side is that T_{inv} ends at the downhill region when I_{p3} is exactly the amount of $I_{disposal}$, which is $I_{\min} - T_{inv}D$. That is:

$I_{\min} - T_{inv}D = D[(N+1)\frac{Q}{D} - T_{inv}]$; hence: $Q_{lb} = \frac{I_{\min}}{N+1}$.

The other boundary point lies when T_{inv} ends at the point where $T_{p3} = T_1$. That is:

$T_{inv} = N\frac{Q}{D} + \frac{Q}{P}$; hence: $Q_{ub} = \frac{T_{inv}}{\frac{N}{D}+\frac{1}{P}}$.

5. Solution Approach

5.1 Local and Global Optimality

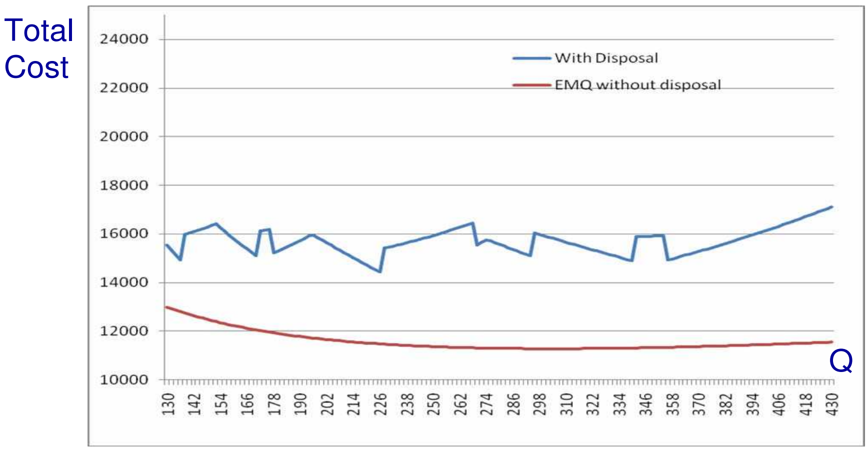

Figure 6: Plot of non-continuous, non-differentiable total cost function with minimum inventory

With the total cost calculation formula in section 4 for the different cases, it is straightforward to plot the total relevant cost with respect to Q (see Fig. 6). The x-axis is the Q value and the y-axis is the total relevant cost. The below smooth plot represents the total relevant cost with respect to Q for a regular EMQ model without the $I_{\min}$ constraint . It is a quadratic function over Q and the optimal Q value can be readily obtained at the stationary point (where the first order derivative is equal to zero). The top irregular plot denotes the total cost with respect to Q for our proposed model. We can see from the plot that the total cost is non-continuous and non-differentiable at those boundary points between two different cases. In this section, we propose an exact solution method to obtain the optimal Q that reaches the minimum total cost with pseudo-polynomial complexity. We first discuss the local optimality property for each segment in the total cost plot. Then we prove the property which guarantees the global optimality. Finally, we present the complete algorithm and demonstrate its complexity.

5.1.1 Local Optimality

Theorem I: (Local Optimality) *Within each case with fixed N and M, the total cost is a quadratic function of Q and the local minimum is either at a boundary point or at the zero first order derivative point.*

Proof. The quadratic property of the total cost with fixed N and M for each case (corresponding to a segment in the plot) is evident from the calculation in section 4. The local minimum lies at either the boundary point with smaller value or the zero first order derivative point (stationary point). □

5.1.2 Global Optimality

After we obtain the local optimality of each segment, we can compare the local minimum values for the different segments and select the lowest one as the global optimum. However, since the number of segments goes to infinity as Q decreases (or N increases), we need to show that only a limited number of segments are required to calculate the local optimal values as potential global optimality candidates.

As in the regular EMQ model, the optimal Q value occurs at a point such that the total ordering cost is at the same level as the total inventory carrying cost. These two components are balanced. In our proposed model, we can use the analogy of the optimality condition for a regular EMQ model to explore the global optimum near the region where the ordering cost and the inventory carrying cost are most balanced. It is straightforward to demonstrate that only when the N value is low, cases 1, 2 and 5 may occur since low N leads to large Q, which also implies large regular underlying EMQ cycles that makes I_{p3} in a scale that is comparable with $I_{disposal}$. When N gets larger, cases 3 and 4 segments alternate between each other. If we prove that the total cost will monotonously increase as N increases after a threshold value $\bar{N}$, then it implies that only a limited number of local optimum need to be computed as global minimum candidates. Next we will prove this property.

Theorem II: (Global Optimality) *When N is greater than a threshold value $\bar{N}$, the total cost will monotonously increase as N increases.*

Proof. We first prove that the total cost of a fixed N value for both cases 3 and 4 is bounded; then we prove that the bound monotonously increases as N increases after a certain threshold value.

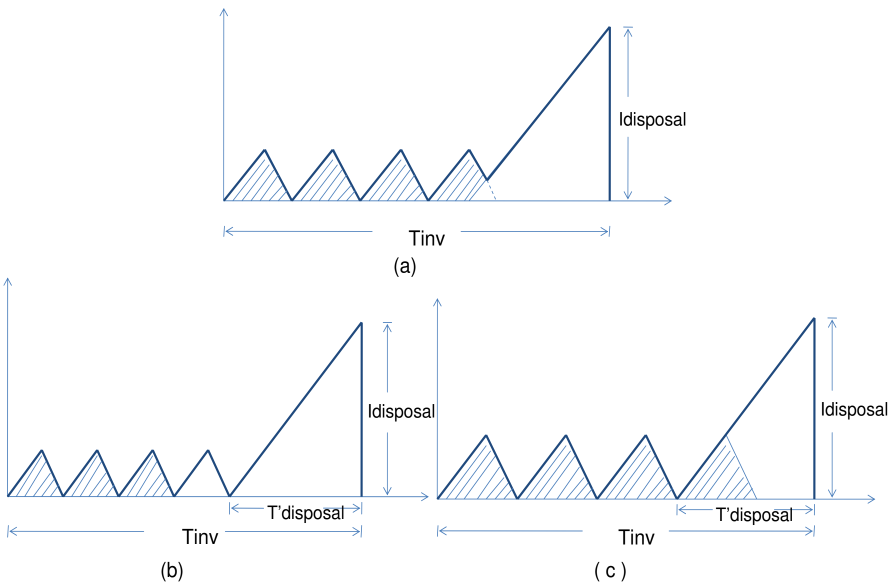

Figure 7: Graph Illustration for the Global Optimum Proof

First, we use the graph to prove the bounds on the total cost with fixed N. Fig. 7 (a) illustrates the general situation of the inventory plot with 4 regular EMQ orderings before the *last production cycle*. If we consider the area under the complete triangle for the *last production cycle* as fixed, then the rest area for the inventory carrying cost would be the area in shadow in Fig. 7 (a). The shadow area is bounded by the shadow area of the three small triangles in Fig. 7 (b) from below; and is bounded by the shadow area of the four large triangles in Fig. 7 (c) from above. This is evident since the *last production cycle* triangle is the same for all three situations (Fig. 7 (a), (b) and (c)) and the time $T_{inv} - T'_{disposal}$ is evenly divided by 3 and 4 respectively in Fig. 7 (c) and Fig. 7 (b). (Please note that $T'_{disposal}$

here represents the length of the *last production cycle*, which is different from the $T_{disposal}$ defined and used in section 3.2 and 4.) Thus the area of one triangle in Fig. 7 (a) is smaller than the one triangle area in Fig. 7 (c) and is larger than the one triangle area in Fig. 7 (b). Since the shadow area in Fig. 7 (a) is between the area for 3 and 4 triangles; the upper and lower bounds on it is proved.

Next, we prove the monotonous increasing property of the bounds after a certain threshold. For both cases 3 and 4, the purchase cost and salvage cost are independent to Q; only ordering cost and inventory carrying cost are sensitive to Q. We let $T_{triangle} = T_{inv} - T'_{disposal}$ denote the time that is occupied by the regular EMQ cycle triangles. The total area for exactly N triangles within $T_{triangle}$ is:
$Area(N) = \frac{1}{2}DT(1-\frac{D}{P})TN$ where $T = \frac{T_{triangle}}{N}$
So: $Area(N) = \frac{D}{2}(1-\frac{D}{P})\frac{T^2_{triangle}}{N}$.
As N increases to $N+1$, the total inventory carrying cost decreases the amount of $Area(N) - Area(N+1)$ and the total ordering cost increases by A (where A is the fixed cost associated with each production setup). Hence it is proved that the total cost at the boundary points (where $T_{triangle}$ can be exactly divided into an integer number of EMQ cycles) is monotonous increasing beyond $\bar{N}$ where $\bar{N}$ is the smallest integer value satisfied by $Area(N) - Area(N+1) \leq A$. □

5.2 Exact Solution Algorithm and Complexity Analysis

Guaranteed by **Theorem I** and **Theorem II**, we have the complete algorithm to reach the global optimum for the modified EMQ model we proposed in section 3.2 with pseudo-polynomial complexity.

Theorem III: (Complexity) *The exact solution algorithm at most needs to explore* $5\bar{N}$ *local minimum points to reach the global minimum, where* $\bar{N} = \lfloor \sqrt{\frac{D}{2}(1-\frac{D}{P})\frac{T^2_{triangle}}{A}} \rfloor$.

Proof. From the proof in **Theorem II**, we have $Area(N) = \frac{D}{2}(1-\frac{D}{P})\frac{T^2_{triangle}}{N}$, and the monotonous increasing threshold is given by $\bar{N}$ where $\bar{N}$ is the smallest integer value satisfied by $Area(N) - Area(N+1) \leq A$.
So we have :
$\frac{D}{2}(1-\frac{D}{P})T^2_{triangle}(\frac{1}{N} - \frac{1}{N+1}) \leq A$
$\frac{D}{2}(1-\frac{D}{P})\frac{T^2_{triangle}}{A} \leq N(N+1) \leq (N+1)^2$
$N+1 \geq \sqrt{\frac{D}{2}(1-\frac{D}{P})\frac{T^2_{triangle}}{A}}$

Hence the smallest integer N which satisfies the above inequality is:
$\lceil\sqrt{\frac{D}{2}(1-\frac{D}{P})\frac{T_{triangle}^2}{A}}\rceil - 1$
which is: $\bar{N} = \lfloor\sqrt{\frac{D}{2}(1-\frac{D}{P})\frac{T_{triangle}^2}{A}}\rfloor$.
For each N, there are at most 5 different cases, thus the maximum number of local minimum points we need to explore before reaching the global optimality is $5\bar{N}$. □

6. Computational Experiments

In this section, we conduct numerical experiments for a potential anthrax attack on our proposed model. It serves two purposes. First, we compare our proposed model which combines two types of demand (an emergency demand – the $I_{\min}$ minimum inventory requirement, and a regular market demand) into a single system, with a standard model which runs the two parts separately; second, we are interested in investigating how the government controlled parameters affect the firm's profit and thus it provides the government a perspective to negotiate parameters with the firms producing the drugs.

This section is organized as follows. We first discuss the parameter estimation in the experiments, then compare two different models (the proposed model and a standard model) and conduct a sensitivity analysis for our proposed model.

6.1 Parameter Estimation

In a potential anthrax attack, the federal government is prepared to treat 10 million exposed persons. This represents a stockpile of 1.2 billion Cipro pills (the treatment regimen is two pills a day for 60 days), as $I_{\min}$ used in our experiments. According to the *Cipro Pharmacy.com* website ($\underline{http://www.ciprofloxacinpharmacy.com/active.html}$): "Cipro has a shelf life of approximately 36 months. However, materiel has been and is currently being tested through the DOD/FDA Shelf Life Extension Program (SLEP) and has received extensions up to 7 1/2 years from original expiration date and some lots have received up to 9 years from original expiration date. Materiel shows no signs of deteriorating based on yearly test." Based on the above statement, we use 9 years (108 months) as the shelf-life of the drug, which indicates that the government would pay for the production and storage of $I_{\min}$ every 9 years. We define a parameter Y to represent the flexibility that the government gives to firms, which is defined as the number of months before the expiration that the government allows $I_{\min}$ to be sold. With the constraint that the maximum cycle length (T_{inv} in our model) must be at most half a shelf-life (we use shelf-life minus the flexibility here to represent the actual allowable on-shelf period). In this section, we use the upper limit that $T_{inv} = \frac{1}{2}(T_s - Y) = \frac{108-Y}{2}$, which dictates the relationship between T_{inv} and Y. We use 36 months for Y in the base case which represents that the government allows firms to resell the required $I_{\min}$ 3 years prior to its expiration. Then $T_{inv} = 108 - 36 \times 2 = 36$ months $=$ 3 years. Please note that the government still pays the firm for producing and keeping the inventory every 9 years (108 months) regardless the flexibility (Y) it allows.

Table 1: Estimated numerical value of the parameters used in the experiment.

Parameter	**Unit**	**Estimation**
$I_{\min}$	million pills	1200 $(= 10 \times 60 \times 2)$
T_{inv}	month	36
Y	month	36 (=108 - 2 $\times$ T_{inv})
p_{gov}	mil \$ /mil pill	0.95
p_{market}	mil \$ /mil pill	4.67
v	mil \$ /mil pill	0.2
w	mil \$ /mil pill	-0.3(3 years)/-0.075(6 years)
D	mil pill / year	300
P	mil pill / year	600
A	mil \$ /time	2 $(= 100 \times h)$
h	mil \$ /mil pill/year	0.02

We also define p_{gov} as the price that the government pays to the firms per pill for production and storage; p_{market} as the price the firms can sell to the regular US domestic market per pill. We estimate p_{gov}, p_{market}, v and w according to Socolar and Sager (2002). Based on the same source, we assume that the government uses a tier pricing strategy to pay firms, 95 cents listed in table 1 represents the price per pill for the first million pills, the second million pills is 10 cent cheaper per pill and from the third million onwards, there is another 10 cent per pill discount. Hence as long as the first million pills' price is given, the total price that the federal government pays to firms is fixed every year. We use 30 cents, the price of generic at Ranbaxy in India (Socolar and Sager, 2002) as the salvage value after 36 months to the secondary overseas market, and 7.5 cents if after 72 months (this will be used in the standard model's calculation only).

The parameters related to the firms' manufacturing and inventory keeping are approximated according to Singh (2001). "In the US alone, Bayer sold \$1.04 billion worth of Cipro in 1999". This is equivalent to about 220 ($\approx \frac{1040}{4.5}$) million pills in year 1999; with a 4% per year growth rate, the demand reaches 300 million in year 2007/2008 and we use this as the demand rate D in our model. "At best, Bayer offered to produce 200 million pills within 60 days". This indicates a maximum production capacity as 1200 million per year, we use half of it – 600 million per year as the production rate allocated for our system. The inventory holding cost (h) mainly comes from the capital cost, with an expected 10% annual investment return rate, $h \approx 0.2 \times 10\% = 0.02$ mil\$/mil pill/year. We assume the ordering cost A is 100 times the h value. In Table 1, we summarize the estimated parameter values as we discussed above and we use them as the base case in our experiments.

As we combine the two parts of demand into our proposed model: the regular market demand with constant rate D and the emergency demand $I_{\min}$ required by the federal government, which are two distinct resources that generate revenues. According to our assumption, both revenues are independent on the *production batch size* (Q); hence the total revenue per year is not a function of Q. As the profit equals to revenue minus cost and the revenue is independent on Q, hence minimizing the cost is the same as maximizing the profit. We emphasize this equivalence relationship because we used minimizing the cost in the analysis in section 4 and we will use the total profit in the following experimental section (6.3) to illustrate how the government controlled parameters ($I_{\min}$, p_{gov} and Y) affect the total profit based on the analysis in 4.

6.2 Model Comparison

We use the base case parameter setting to run our proposed model. We also use the same parameters to run a standard model, which uses two separate systems to meet the two parts of the demand respectively. At the same time, we split the production capacity 600 million pills per year into two parts for these two individual systems. We use 400 million pills per year for the regular market demand with the regular EMQ model and another 200 million pills per year constantly running for the 1.2 billion pills as $I_{\min}$, which is completely refreshed every 6 years and we can obtain 7.5 cents per pill at the secondary market when they are salvaged after 6 years. As the computation is demonstrated below, our model cost 78 million per year to run the combined system which saves 33.45 million a year compared with the standard model. This illustrates the significant benefit (30% cost-saving in a year) by using our proposed model instead of running two separate systems.

For our proposed model, there are four parts in the total cost:

- Inventory carrying cost:
  - Emergency part: $1200 \times 0.02 = \$24$ mil/year;
  - Regular market part: $(\frac{1}{2} \times 150 \times 2 + \frac{1}{2} \times 300) \times 0.02 = 6$ million for 3 years, which is \$2 million per year;
- Ordering cost: 3 times in each T_{inv} (3 years), which is \$2 million per year;
- Salvage cost: $300 \times (-0.3) = -90$ million dollars in 3 years, which is -30 million dollars per year;
- Variable cost: $I_{\min} \times v = 1200 \times 0.2 = 240$ million dollars for 3 years, which is \$80 million per year.

Hence the total annual cost of the proposed system is: $24 + 2 + 2 - 30 + 80 = \$78$ million.

For the standard model, we calculate the cost for the two systems separately:

- Regular market EMQ model:
  - Inventory carrying cost: $\sqrt{2hAD(1 - \frac{D}{P})} = \$2.45$ mil/year;
  - Variable cost: $0.2 \times 300 = \$60$ mil/year;
- Emergency demand:
  - Inventory holding cost: $1200 \times 0.02 = \$24$ mil/year;
  - Variable cost: $\frac{1200}{6} \times 0.2 = \$40$ mil/year;
  - Salvage cost: $-0.075 \times \frac{1200}{6} = -\$15$ mil/year

Hence the total annual cost of the standard model is: $2.45 + 60 + 24 + 40 - 15 = \$111.45$ million.

6.3 Sensitivity Analysis

In the previous subsection, we demonstrated the advantage of our proposed model over a standard model by using the base case parameter settings. Now we are interested in the sensitivity analysis on the government controlled parameters: $I_{\min}$, Y and p_{gov}. We use different salvage strategies to investigate how the profit changes for different $I_{\min}$, Y and p_{gov} values. All figures in this subsection are plotted as the profit from the manufacturer's perspective over the parameters controlled by the government. The government can use these plots to negotiate the parameter settings with the firms.

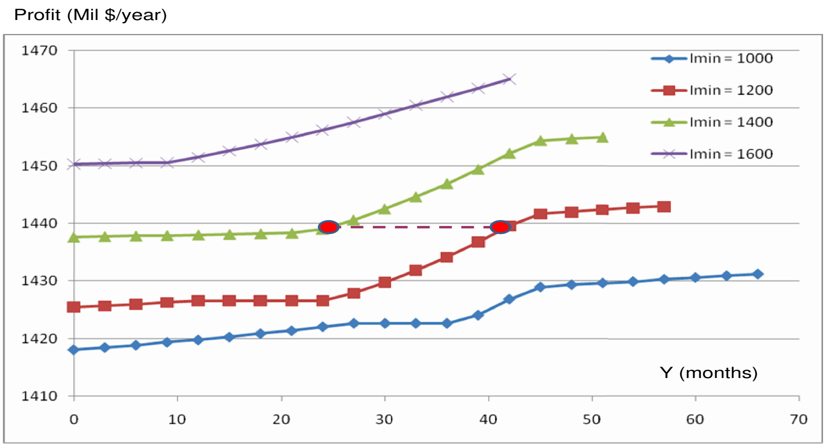

Figure 8: Profit v.s. Y with $Q_{max} = 1600$ at different $I_{\min}$ level

Fig. 8 and Fig. 9 plot the relationship between the total profit and the flexibility (Y) given by the government at different $I_{\min}$ level. The x-axis represents Y which ranges from 0 to 69 months and it corresponds to T_{inv} from 54 to 19.5 months. The y-axis represents the total profit and different plot lines denote different $I_{\min}$ levels.

In Fig. 8, we limit the *maximum production batch size* (Q_{max}) as 1600 million pills. Based on this figure, we have the following observations.

- For any fixed flexibility (Y), the higher $I_{\min}$ level, the more profit the firm can gain; this is due to the extra revenue obtained from the government is higher than the additional cost required for maintaining the extra $I_{\min}$.

- For any fixed $I_{\min}$ level, the more flexibility allowed (larger Y, which means firms can salvage the $I_{\min}$ amount earlier), the higher profit firms can gain; this is due to the extra flexibility given to firms that they can refresh their inventory in a shorter period hence reduce the average running cost of the proposed inventory system. However, the profit increases at different slopes with different Y values. This phenomenon will be explained in a later discussion.

- The two dots, which on two different $I_{\min}$ level plot lines and at the same profit level, demonstrate the trade-off between the high flexibility for the firms and the low minimum inventory requirement for the government. If the government would pay less to the firms by requiring a smaller amount of minimum inventory, it must allow more flexibility (longer time before the expiration to salvage the pills) for the firms to obtain the same level of profitability.

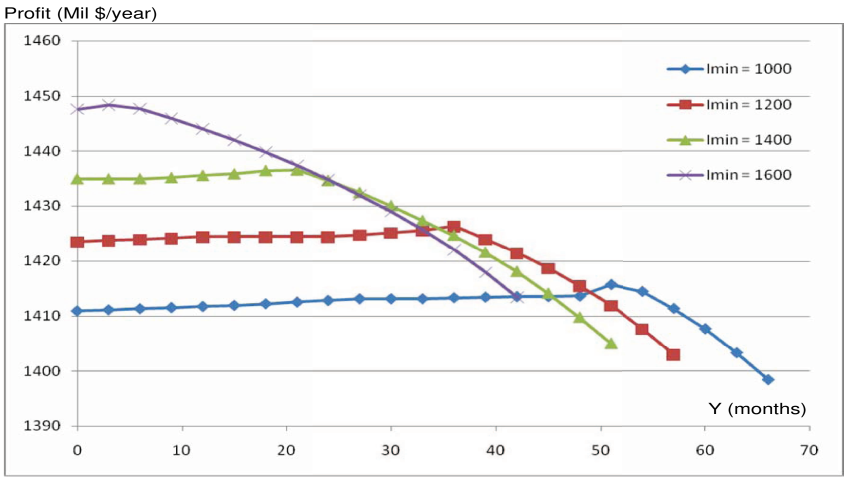

Figure 9: Profit v.s. Y which allows salvage at most 300 million pills

In Fig. 9, we use a different salvage strategy by assuming that we can earn salvage value for the first 300 million disposed pills only. Additional pills will be disposed at a zero salvage value. Compared with Fig. 8, this figure shares the same plot shape when Y is low. But the plots start to have a negative slope at the points when the salvage amount reaches 300 million pills. The turning points are at different Y values for different $I_{\min}$ plots; the lower the $I_{\min}$, the larger the turning point's Y value. This is explained by the fact that a less minimum inventory requirement will reach a given salvage amount with higher flexibility (or, a larger Y); or from the mathematical formula: the salvage amount $=$ $I_{\min} - DT_{inv} = I_{\min} - D\frac{9-Y}{2} = I_{\min} - 4.5D + \frac{DY}{2}$. For a fixed salvage amount, a smaller $I_{\min}$ comes with a larger Y. When $I_{\min}$ is fixed, as the flexibility Y increases, the inventory cycle (T_{inv}) is reduced; this results in more salvage amount at the end of each T_{inv}. If the salvage amount exceeds the 300 million limit, only the first 300 million is valuable; however, we still need to pay the unit cost for those salvaged stocks above the 300 million. Thus, after the turning point, the more we are forced to salvage, the less profit we gain.

Fig. 10 illustrates the relationship between the profit and the flexibility (Y) at different prices the government would pay firms. Again we limit the *maximum production batch size*

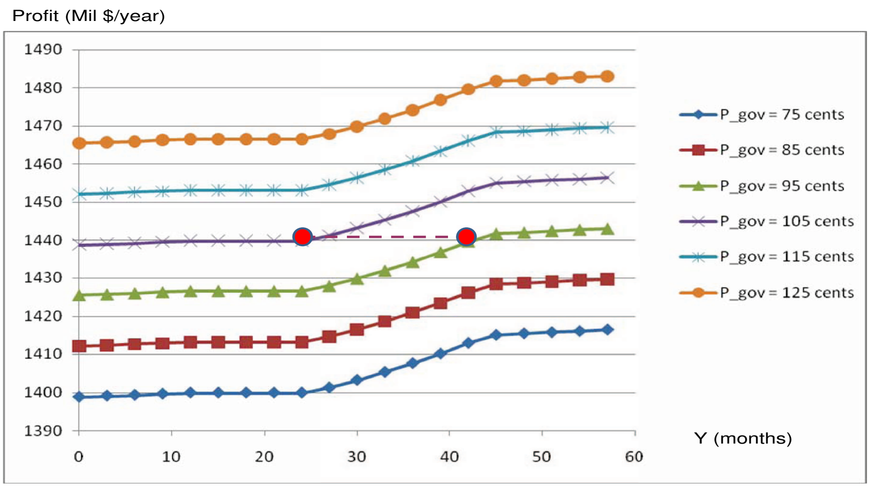

Figure 10: Profit v.s. Y with $Q_{max} = 1600$ at different unit government-paid price

(Q_{max}) as 1600 million pills and use the base case $I_{\min}$ as 1.2 billion pills. Similar to Fig. 8, we have the following observations.

- The higher unit price the government pays, the higher profit the firm gains.
- The two dots, which on two different p_{gov} level plot lines and at the same profit level, demonstrate the trade-off between the high flexibility for the firms and the low price for the government. If the government would pay less to the firms by reducing the unit price of a pill, it must allow more flexibility (longer time before the expiration to salvage the pills) for firms to obtain the same level of profitability.

Fig. 11 provides a graph explanation on the different slopes at different segments on the profit v.s. flexibility (Y) plot. It uses the base case parameters and gives the inventory plots at the *optimal production batch size* Q^* for different Y values. Through the sensitivity analysis, the marginal profit gained over a fixed small $\triangle Y$ in (a), (b) and (d) segments all come from the cost saving on the regular EMQ part, which is trivial, compared with the marginal profit gain over a fixed small $\triangle Y$ at the (c) segment, which is from the additional gain due to the extra salvage amount.

In summary, in this experimental section, we estimate the model parameters for a potential anthrax attack scenario and use them to quantitatively demonstrate the advantage of our proposed model over a standard model. We also perform a sensitivity analysis on the

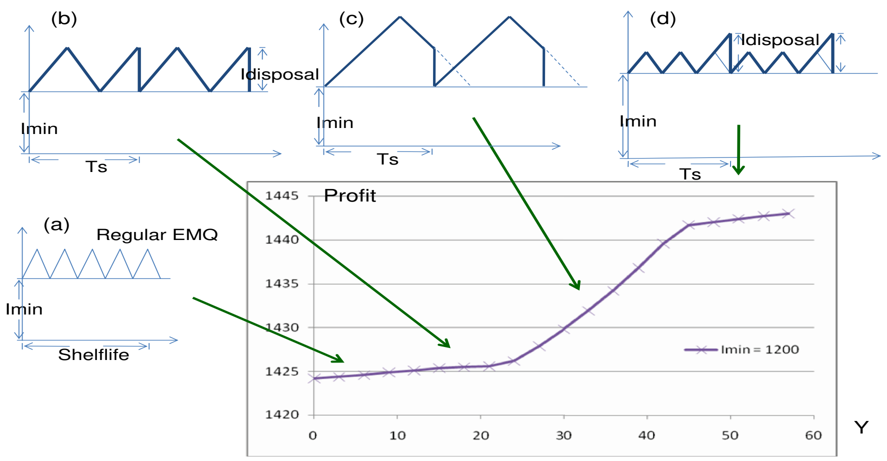

Figure 11: Illustration of the different slopes in the profit v.s. Y plot

government controlled parameters to provide the aid for the negotiation on the parameter settings between the federal government and the firms.

7. Conclusion and Future Research

For the perishable inventory management system, a trivial extension of a regular EMQ model is adequate when the required minimum inventory is not significant compared with the amount consumed by the regular market demand rate within the shelf-life. Since it can be timely and completely depleted and refreshed by the regular market demand. However, when we consider the VMIs in the SNS for the large-scale emergency setting, the minimum inventory requirement is in a scale which is comparable with the total market consumption within the shelf-life. A more sophisticated inventory management strategy is required to provide the fresh and massive stockpile throughout the time horizon in the system. Hence in this work, we modeled the perishable inventory management problem with a minimum inventory volume constraint as a modified economic manufacturing quantity (EMQ) model. We discussed the policies and assumptions adopted in this model from both the regular perishable inventory management context and the special constraints on the minimum stock size and maximum inventory cycle length enforced by the large-scale emergency response context. Different possible scenarios are discussed and the calculation on the total cost and boundary conditions for each scenario is presented. The total cost is decomposed into four components: inventory holding costs, fixed ordering costs, purchasing costs and salvage costs.

With the aid of inventory plots, we formulated the problem to minimize the total relevant cost w.r.t. the production batch size as an unconstrained non-continuous non-differentiable optimization problem. We proved the existence of the local as well as global minimum of the total cost with respect to the order quantity. Hence, an exact solution procedure is proposed and its complexity is proved to be pseudo-polynomial. We estimated the parameters in the modified EMQ model for a potential anthrax attack scenario from various sources and used them to compare the proposed model with a standard model to show a significant cost saving of running our system as around 33 million US dollars per year, which saved about 30% of the cost to a standard model. We performed sensitivity analysis on some government controlled parameters in the system and observed that at a given profitability level of the firm, there are trade-offs between the less amount paid by the government to firms (either by reducing the $I_{\min}$ requirement or by reducing the unit price p_{gov}) with the higher flexibility the government allows to firms (longer time before the expiration to salvage the pills).

For the inventory management problem, our current model assumed uniform unit price on the items sold to the regular market and uniform unit salvage value on the items disposed at the end of each inventory cycle regardless of their age. An interest in the future research is to combine the inventory model with a revenue model to address the potential economic impact with more sophisticated issuing, pricing and disposing strategies which incorporate the age distribution of the stockpile. Another interesting direction in the future work is to extend the deterministic demand rate to a stochastic one. Since it is more realistic to assume that the demand is random and it follows a certain probabilistic distribution and address it with some stochastic analysis techniques.

Appendix A: Detailed Graph Illustration of cases 2, 4 and 5

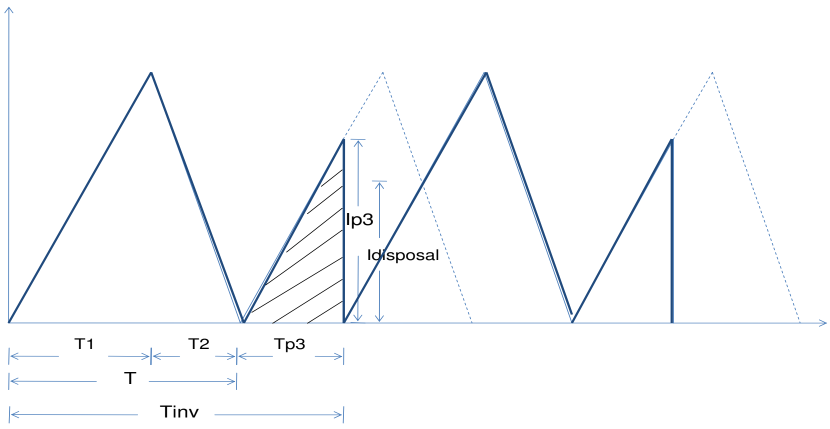

Figure 12: Graph Illustration for Case 2.

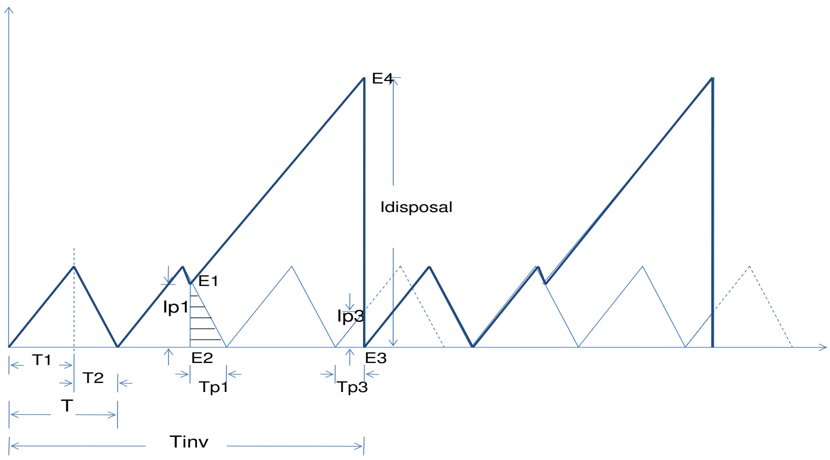

Figure 13: Graph Illustration for Case 4.

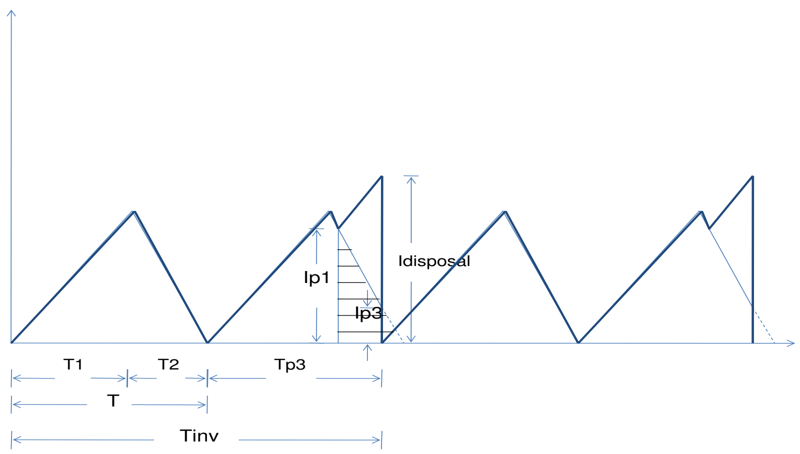

Figure 14: Graph Illustration for Case 5.

Appendix B: Total Cost Calculation and Boundary Conditions for case 3

$$TRC(Q) = \frac{1}{2}h[M \cdot T \cdot I_{max} - T_{p1} \cdot DT_{p1}) + (I_{disposal} + T_{p1} \cdot D) \cdot (T_{inv} - M \cdot T + T_{p1})]$$

We first compute T_{p1}, which can be expressed as a linear function of Q:

$$\begin{aligned} T_{p1} &= T_{disposal} - (T_{p3} - T_1) - N_1 \cdot T_2 \\ &= \frac{I_{disposal} + T_{inv} \cdot (D - P)}{P} + (N + 1 - N_1) \cdot (\frac{1}{D} - \frac{1}{P}) \cdot Q \\ &= \beta_1 + \alpha_1 \cdot Q \end{aligned}$$

Where we define:

$$\alpha_1 = (M + 1) \cdot (\frac{1}{D} - \frac{1}{P}) \qquad \beta_1 = \tfrac{I_{\min}}{P} - T_{inv}$$

Hence we have:

$$\begin{aligned} TRC(Q) &= \frac{1}{2}h[M \cdot \frac{Q^2}{D} \cdot (1 - \frac{D}{P}) - (\alpha \cdot Q + \beta)^2 \cdot D + \\ &\quad (I_{disposal} + (\alpha \cdot Q + \beta) \cdot D) \cdot (T_{inv} - M \cdot \frac{Q}{D} + \alpha \cdot Q + \beta)] \end{aligned}$$

We let $TRC(Q) = a \cdot Q^2 + b \cdot Q + c$, then

$$\begin{aligned} a &= \frac{1}{2}h \cdot [M \cdot (\frac{1}{D} - \frac{1}{P}) - D \cdot \alpha^2 + \alpha \cdot D \cdot (\alpha - \frac{M}{D})] \\ &= -\frac{h}{2} \cdot M^2 \cdot (\frac{1}{D} - \frac{1}{P}) < 0 \end{aligned}$$

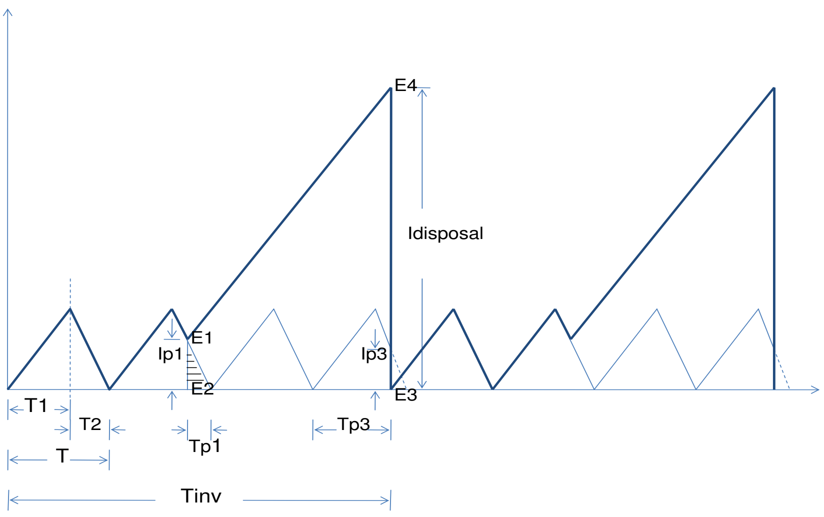

Figure 15: Graph Illustration for Case 3.

$$b \quad = \quad \frac{h}{2} \cdot [(2M+1) \cdot (T_{inv} - \frac{I_{disposal}}{P} - \frac{D}{P} \cdot T_{inv}) + \frac{I_{disposal}}{D}]$$

The $TRC(Q)$ is a quadratic function of Q and the stationary point of the quadratic curve is (where $a < 0$):

$$Q^*_{case3} = -\frac{b}{2a} = \frac{(2M+1) \cdot (T_{inv} - \frac{I_{disposal}}{P} - \frac{D}{P} \cdot T_{inv}) + \frac{I_{disposal}}{D}}{2M^2 \cdot (\frac{1}{D} - \frac{1}{P})}$$

Next, we show the calculation of the boundary point values for case 3. Without loss of generality, we assume that the N and M values are both fixed and as shown in Fig. 16, the T_{inv} period ends at somewhere in the downhill slope between C and D; the *last production cycle* starts somewhere in the downhill region between A and B. From Fig. 16, we can infer that only if $I_{disposal}$ falls in range 1 or range 2 or range 3, it is feasible for case 3 with the parameter M and N valid (that is, T_{inv} ends between C and D and the *last production cycle* starts between A and B). Otherwise, if $I_{disposal}$ is too short (below range 3) or too high (above range 2) with fixed N, the value of parameter M would increase or decrease accordingly.

When $I_{disposal}$ falls into range 1, the two boundary points of case 3 with M and N are:
(1) T_{inv} ends at D, and the *last production cycle* starts at $SD1$;
(2) T_{inv} ends at C, and the *last production cycle* starts at $SC1$.

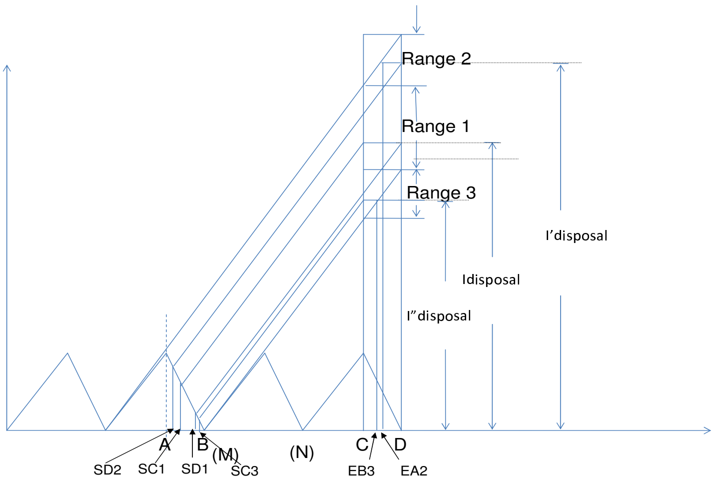

Figure 16: Graph Illustration for the Boundary Value of Case 3.

Or when $I_{disposal}$ falls into range 2, the two boundary points are:

(1) T_{inv} ends at D, and the *last production cycle* starts at $SD2$;

(2) T_{inv} ends at $EA2$, and the *last production cycle* starts at A.

Or when $I_{disposal}$ falls into range 3, the two boundary points are:

(1) T_{inv} ends at $EB3$, and the *last production cycle* starts at B;

(2) T_{inv} ends at C, and the *last production cycle* starts at $SC3$.

Next, we will calculate the Q values corresponding to these boundary points.

For fixed N and M values, when T_{inv} ends at point D, we have $(N+1)\cdot T = T_{inv}$. Since $T = \frac{Q_D}{D}$, we have $(N+1)\cdot\frac{Q_D}{D} = T_{inv}$. Hence $Q_D = \frac{T_{inv}\cdot D}{N+1}$, where Q_D is the *production batch size* when T_{inv} ends at point D.

When T_{inv} ends at point C, we have $N\cdot T + T_1 = T_{inv}$, so $N\cdot\frac{Q_C}{D} + \frac{Q}{P} = T_{inv}$. Hence, $Q_C = \frac{T_{inv}}{\frac{N}{D}+\frac{1}{P}}$, where Q_C is the *production batch size* when T_{inv} ends at point C.

If T_{inv} ends somewhere between C and D and the *last production cycle* starts at A, we have

$$\frac{I_{disposal} - I_{p3}}{P} = (N-M+1)\cdot T_2 + [T_{inv} - (N\cdot T + T_1)],$$

If we use Q_A to represent the *production batch size* when the *last production cycle* starts from A, then:

$$\frac{I_{disposal} - [(N+1) \cdot \frac{Q_A}{D} - T_{inv}] \cdot D}{P} = (N-M+1) \cdot Q_A \cdot (\frac{1}{D} - \frac{1}{P}) + [T_{inv} - (N \cdot \frac{Q_A}{D} + \frac{Q_A}{P})]$$

Hence we have:

$$Q_A = \frac{(P-D) \cdot T_{inv} - I_{disposal}}{P \cdot (M-1) \cdot (\frac{1}{D} - \frac{1}{P})}$$

If T_{inv} ends somewhere between C and D and the *last production cycle* starts at B, we have

$$\frac{I_{disposal} - I_{p3}}{P} = (N-M) \cdot T_2 + [T_{inv} - (N \cdot T + T_1)],$$

If we use Q_B to represent the *production batch size* when the *last production cycle* starts from B, with the similar calculation for Q_A, we have:

$$Q_B = \frac{(P-D) \cdot T_{inv} - I_{disposal}}{P \cdot M \cdot (\frac{1}{D} - \frac{1}{P})}$$

References

Abad, P. (1996). Optimal pricing and lot-sizing under conditions of perishability and partial backordering. *Management Science 42*(8), 1093–1104.

Burnetas, A. and C. Smith (2004). Adaptive ordering and pricing for perishable products. *Operations Research 48*(3), 436–443.

Chazan, D. and S. Gal (1977). A markovian model for a perishable product inventory. *Management Science 23*(5), 512–521.

Chun, Y. H. (2003). Optimal pricing and ordering policies for perishable commodities. *European Journal of Operational Research 144*(1), 62–82.

Gallego, G. and G. van Ryzin (1994). Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science 40*(8), 999–1020.

Giri, B. and K. Chaudhuri (1998). Determinisitic models of perishable inventory with stock-dependent demand rate and nonlinear holding cost. *European Journal of Operational Research 105*(3), 467–474.

Goh, C.-H., B. Greenberg, and H. Matsuo (1993). Two-stage perishable inventory models. *Management Science 39*(5), 633–649.

Graves, S. (1982). The application of queuing theory to continuous perishable inventory systems. *Management Science 28*(4), 400–406.

Hsu, V. N. (2000). Dynamic economic lot size model with perishable inventory. *Management Science 46*(8), 1159–1169.

Hsu, V. N. and T. Lowe (2001). Dynamic economic lot size models with period-pair-dependent backorder and inventory costs. *Operations Research 49*(2), 316–321.

Katagiri, H. and H. Ishii (2002). Fuzzy inventory problems for perishable commodities. *European Journal of Operational Research 138*(3), 545–553.

Liu, L. and Z. Lian (1999). (s, s) continuous review models for products with fixed lifetimes. *Operations Research 47*(1), 150–158.

Nahmias, S. (1982). Perishable inventory theory: A review. *Operations Research 30*(4), 680–708.

Nandakumar, P. and T. Morton (1993). Near myopic heuristics for the fixed-life perishability problem. *Management Science 39*(12), 1490–1498.

Pierskalla, W. and C. Roach (1972). Optimal issuing policies for perishable inventory. *Management Science 18*(11), 603–614.

Ravichandran, N. (1995). Stochastic analysis of a continuous review perishable inventory system with positive lead time and poisson demand. *European Journal of Operational Research 84*(2), 444–457.

Rosenfield, D. (1989). Disposal of excess inventory. *Operations Research 37*(3), 404–409.

Rosenfield, D. (1992). Optimality of myopic policies in disposing excess inventory. *Operations Research 40*(4), 800–803.

Shen, Z. (2008). *Routing and Inventory Models for Emergency Response to Minimize Unmet Demand*. Ph. D. thesis, University of Southern California.

Singh, K. (2001, Nov). War profiteering: Bayer, anthrax and international trade. http://www.corpwatch.org/article.php?id=723.

Sobel, M. and R. Zhang (2001). Inventory policies for systems with stochastic and deterministic demand. *Operations Research 49*(1), 157–162.

Socolar, D. and A. Sager (2002, June). Cipro prices expose windfall profits. Health Services Department, Boston University School of Public Health.

Wagner, H. M. and T. M. Whitin (1958). Dynamic version of the economic lot size model. *Management Science 5*(1), 89–96.

Weiss, H. (1980). Optimal ordering policies for continuous review perishable inventory models. *Operations Research 28*(2), 365–374.