Inventory control with multiple setup costs

Citation for published version (APA):

Alp, O., Huh, W. T., & Tan, T. (2014). Inventory control with multiple setup costs. *Manufacturing & Service Operations Management*, 16(1), 89-103. <https://doi.org/10.1287/msom.2013.0450>

DOI:

[10.1287/msom.2013.0450](https://doi.org/10.1287/msom.2013.0450)

Document status and date:

Published: 01/01/2014

Document Version:

Accepted manuscript including changes made at the peer-review stage

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Inventory Control with Multiple Set-Up Costs

Osman Alp*

Woonghee Tim Huh[†]

Tarkan Tan[‡]

July 4, 2009; Revised: July 5, 2010, June 7, 2012, January 2, 2013, April 3, 2013

Abstract

We consider an infinite-horizon, periodic-review, single-item production/inventory system with random demand and back-ordering, where multiple set-ups are allowed in any period, and a separate fixed cost is associated for each set-up. Contrary to majority of the literature on this topic, we do not restrict the order quantities to be integer multiples of the exogenously-given batch size and instead allow the possibility of partial batches, in which case the fixed cost for ordering the batch is still fully charged. We build a model that particularly takes the batch ordering cost structure into account. We introduce an alternative cost accounting scheme to analyze the problem, which we use for developing a computationally efficient optimal solution method and several properties of the optimal solution. In addition, we propose two heuristic policies, both of which perform extremely well computationally.

1 Introduction and Related Literature

In manufacturing environments, batch production occurs frequently. For example, when industrial ovens are used for production, a number of items are processed simultaneously upon a single setup. Similarly, a fixed cost component is incurred for the initiation of the production of each batch, for reasons such as the usage of materials (e.g. a wafer that has to be consumed once the production of a batch of electronic chips is initiated), dies, and forms. In these cases, the batch setting becomes an important consideration for highly customized products with relatively low demands. Another example is for procurement environments, where a powerful supplier manifests a price scheme that reflects his cost structure which involves fixed costs for reasons such as the ones mentioned above. In retail industry, on the other hand, suppliers generally ship items with trucks or in containers.

*Industrial Engineering Department, Bilkent University, 06800 Ankara, Turkey. osmanalp@bilkent.edu.tr.

[†]Sauder School of Business, University of British Columbia, Vancouver BC Canada, V6T 1Z2. Research partially supported by the National Science Foundation of the United States (DMS-0732169), and the Natural Sciences and Engineering Research Council of Canada. tim.huh@sauder.ubc.ca.

[‡]School of Industrial Engineering, Eindhoven University of Technology, P.O. Box 513, 5600MB Eindhoven, The Netherlands. T.Tan@tue.nl.

If the cost of such shipments is born by the retailer, batch ordering is practiced as the usage of each additional truck is associated with a fixed cost which is independent of the quantity loaded in the truck. For managing transportation within a distribution system, decisions are made daily whether to send a truck that is partially full or whether it should be sent another day to fill it up – as illustrated by a Harvard Business School case on Merloni (Hammond and Kelly, 1990).

In line with these examples, Lippman (1969a, 1971) also demonstrates the applicability of such a cost structure in the transportation and production context and characterizes the optimal solutions with deterministic demand. When demand is stochastic, however, the treatment of batch ordering is more analytically challenging (see Lippman 1969b). Typical approaches for modeling the procurement cost in the literature are either (i) to ignore the set-up costs altogether by considering only the linear cost, (ii) to include a single set-up cost for ordering, i.e., to assume an infinite batch size, (iii) to cap the total order quantity in each period, or (iv) to impose additional requirement such as full batch sizes (also referred to as “full truck loads” or “full containers”). We note each of these models is appropriate in many application settings, and has been analyzed well. The results for the first three variants are widely known, characterized by the optimality of the base-stock policy (Scarf, 1960; Veinott, 1966), the (s, S) policy (Scarf, 1960; Veinott, 1966), and the modified base-stock policy (Federgruen and Zipkin, 1986a and 1986b), respectively.

In the literature, the phrase “batch ordering” typically refers to the problem with the full batch size ordering restriction, i.e., the ordering quantity in each period must be an integer multiple of the full batch size (or capacity). This problem is introduced by Veinott (1965) who shows the optimality of a threshold-type policy, where the order quantity in each period is the smallest multiple of the full batch size that will bring the inventory level above a certain level. Iwaniec (1979) identifies a set of conditions under which full batch size ordering policy is optimal. Later, Zheng and Chen (1992) develop an algorithm to compute the optimal parameters within the class of the full batch ordering policy. With the full-batch restriction, the optimality of the threshold-type policy extends to multi-echelon systems as shown by Chen (2000) and Huh and Janakiraman (2010). Related to this, in a single-stage problem with at most one full batch in each period, Gallego and Toktay (2004) show that a straightforward modification of the above threshold-type policy remains optimal. None of the aforementioned papers on batch ordering allows the possibility of a partial batch.

Iwaniec (1979) shows that full batch size ordering policy is optimal when the ratio of the fixed ordering cost to the batch size is high, which is consistent with our numerical result. Hence, the restriction of ordering full batches may be reasonable when the fixed cost of ordering is dominant. However, with moderate fixed costs, rounding up or down the order quantity to a multiple of the full batch size may not be optimal. The batch ordering problem with the possibility of partial batches is a difficult problem to analyze since the ordering cost as a function of ordering quantity has multiple points of discontinuity. The manager needs to decide how many batches to order as well as whether all the batches should be full or not. It involves not only comparing the risks of having too much inventory and too little inventory, but also striking a balance between maximizing output for each set-up and exercising the flexibility to calibrate the amount of inventory. When the problem parameters are such that this balance is less intricate, traditional approaches will suffice. For example, when the batch capacity is abundant, in comparison to demand, the problem can be approximated by an infinite capacity problem. (If the batch size is infinite, then an optimal policy is an (s, S) policy – see, for example, Scarf (1960).) Similarly, when the fixed cost component is dominant or the fixed batch size is small, searching the solution within the full batch ordering policies will be sufficient, as discussed above. (Note that if the batch size is 1 unit, then every non-fractional order is “full batch” ordering.) Finally, when the fixed cost is small, the optimal myopic policy will also perform well for the infinite-horizon problem – since myopic policy is optimal when the fixed cost is zero. From our numerical computation (Section 4), even the best of these traditional policies can result in substantial errors (up to 91% in our test bed). Furthermore, it is difficult to know *a priori* which policy should be applied in which case and finding the optimal parameters of the policies such as (s, S) is time-consuming. Therefore, there is a need of efficient and better solutions in some contexts.

The partial batch ordering flexibility, a key feature of this paper, has received relatively little attention in the literature. The classical reference for partial batch ordering under stochastic demand in a single-stage and single-item setting is found in Lippman (1969b), who considers subadditive ordering costs for which the partial batch ordering problem is an example. Using the myopic-optimality framework of Veinott (1965), Lippman derives some partial properties of the optimal policy. The class of ordering cost structure that Lippman (1969b) considers is a subadditive func-

tion, which is more general than our cost structure; however, we provide a sharper characterization of the optimal policy in the sense that we provide bounds that are tighter and easily computable from input data (see Section 3.3 for details). In addition, we note that, like the single-period problem of Lippman (1969b), the single-period problem that we identify also has a two-threshold structure in the optimal policy, but we also show that the myopic policy based on our single-period cost performs well computationally as a heuristic policy for the infinite-horizon problem.

While a single-stage problem with partial batch ordering is already a difficult problem, we are aware of only two papers, Cachon (2001) and Tanrikulu et al. (2010), that contain this feature as a part of a more complex multi-item joint-replenishment context where heuristic methods are proposed and evaluated. Similar to Lippman (1969b), we focus on a simpler yet still challenging single-stage and single-item problem. In addition, we mention some other papers that are related to this paper. Chao and Zhou (2009) consider the multi-echelon full-batch ordering problem with *minimum setup time* (time between consecutive orders). Alp et al. (2003) consider a general version of our problem under *deterministic* demand settings and show some optimality properties.

We mention two recent papers that generalize the cost structure of this paper by allowing the possibility that multiple fixed costs in a period may not be identical. These papers include analytical results that are more broadly applicable than our setting, but we restrict our attention to the case of identical fixed costs. Caliskan-Demirag et al. (2010) use the notion of the (C, K) -convexity (Shaoxiang, 2004) to obtain a result on after-ordering inventory levels that is related to Theorem 7 of our paper. Their result can be thought of as an extension of Shaoxiang and Lambrecht (1996), where one can order at most one batch per period. Zhang and Çetinkaya (2011) further refine the partial characterization of the optimal policy by recognizing a monotonicity property of after-order inventory levels. (The structural property that Zhang and Çetinkaya identify has an arbitrary number of thresholds, making the resulting dynamic program computationally challenging to solve.) Their approach is based on a novel notion called “non- (Δ, C) -decreasing”, which is related to the (C, K) -convexity. We note that while the primary focus of these papers as well as Lippman (1969b) is identifying the *structure* of the optimal policy, we seek to identify how to solve the problem easily (both optimally and heuristically). These three papers consider a finite-horizon problem while we consider the long-run average cost criterion taking advantage of the alternate cost accounting

scheme, which is one of our contributions.

Our main contributions can be summarized as follows: (1) To facilitate our analysis of the partial batch ordering problem, we introduce an alternative cost accounting scheme equivalent to the original cost structure. Here, instead of charging the per-batch cost to every batch ordered, we recognize that an unnecessary additional ordering cost is incurred when a batch is *less than* full and impose an appropriate penalty cost for such batches. (2) We provide a complete characterization of the optimal solution for the single-period problem with the alternative cost accounting scheme. (3) We also study a relaxed version of the multi-period problem and describe how to find an optimal solution for the relaxed problem, which provides a lower bound on the optimal cost of the original problem. (4) Making use of the alternative cost accounting scheme, we develop an efficient optimal solution method and some properties of the optimal solution. (5) We examine two policies that can be used to solve the problem in a heuristic manner. These policies are based on our analysis of the single-period problem and the lower bound mentioned above, and they perform very well in a wide range of problem parameters.

2 Model

2.1 Description

Demand is stochastic and unmet demand is assumed to be fully backlogged. The relevant costs in our environment are inventory holding costs, backorder costs and fixed costs associated with set-ups. We ignore unit ordering costs without loss of generality because of our interest in a long-run average cost. We assume full availability of the ordered quantities, and that the lead times can be neglected. The capacity of each batch is fixed, and we refer to it as a *base* quantity; we do not restrict the order quantities to be integer multiples of the base quantity, but the set-up cost is a function of the number of set-ups.

Let $t \in \{1, 2, \dots\}$ index the time periods in a forward manner. The following sequence of events takes place in each period t . (1) At the beginning of the period, the manager observes the current inventory level denoted by x_t . Positive x_t corresponds to excess inventory, and negative x_t corresponds to outstanding backlog. (2) The manager then orders $q_t \geq 0$ units, which incurs a

set-up cost of $\tilde{c}(q_t)$ given by

$$\tilde{c}(q) = K \cdot \lceil q/Q \rceil$$

where $K \geq 0$ represents the fixed cost per set-up, $Q > 0$ is an integer denoting the fixed base quantity, and $\lceil \cdot \rceil$ is the smallest integer greater than or equal to the argument inside (thus, $\lceil q/Q \rceil$ denotes the number of batches (the number of set-ups) required to order q units). Note that $\tilde{c}(q)$ is a right-continuous step function where the non-continuous increments are identical and equally spaced. Since we assume that order replenishment is instantaneous, these q_t units arrive immediately. It is convenient to denote the after-ordering inventory level by y_t , i.e., $y_t = x_t + q_t$. Clearly, $y_t \geq x_t$. (3) Then, demand D_t is realized. We assume that the sequence of demands $(D_1, D_2, \dots)$ are independent and identically distributed, and we denote the common distribution by D . We assume that demands are discrete with integer supports. (We note that the setting in Lippman (1969b) allows for continuous demand.) Furthermore, we assume $E[D] > 0$ (otherwise, the problem becomes trivial). An appropriate linear overage or underage cost of h or b per unit is charged, respectively. (4) The excess demand, if any, is backlogged, and therefore, $x_{t+1} = y_t - D_t$. The expected overage and underage cost in each period t depends only on y_t , given by

$$\mathcal{L}(y) = E_D [h \cdot [y - D]^+ + b \cdot [D - y]^+] , \quad (1)$$

which is convex. The expected cost in period t can be written as $\tilde{C}(x_t, y_t) = \tilde{c}(y_t - x_t) + \mathcal{L}(y_t)$.

In this paper, we consider the infinite-horizon problem with the objective of minimizing the long-run average cost, i.e., $E \left[\limsup_{T \rightarrow \infty} \sum_{t=1}^T \tilde{C}(x_t, y_t) / T \right]$. (While there is some debate regarding whether the objective function should employ the supremum or infimum when it is not known *a priori* whether the limit exists or not (Puterman, 1994; Section 8.1), we adopt the former.) Let $\tilde{C}^*$ denote the optimal long-run average cost. In addition, it is useful to consider a finite planning horizon of T periods, and to define the T -period average cost by

$$\frac{1}{T} \sum_{t=1}^T E \left[\tilde{C}(x_t, y_t) \right] - \frac{1}{T} \cdot v \cdot x_{T+1} , \quad (2)$$

where v can be interpreted as the unit salvage value if the terminal inventory x_{T+1} is nonnegative, or the unit penalty cost if x_{T+1} is negative. Note that the salvage value becomes insignificant for the long-run average cost as the horizon becomes arbitrarily long. This observation allows us formulate a single-period cost that is different from the one used by Lippman (1969b).

2.2 Dynamic Program and Equivalent Alternate Cost Formulation

The finite-horizon T -period problem can be formulated by using dynamic programming:

$$f_t(x_t) = \min_{y_t \geq x_t} \tilde{c}(y_t - x_t) + \mathcal{L}(y_t) + E_{D_t} [f_{t+1}(y_t - D_t)] \text{ for } t = 1, \dots, T, \text{ and } f_{T+1}(x_{T+1}) = -vx_{T+1}. \quad (3)$$

This formulation is essentially the same as the dynamic program given in Section 3 of Lippman (1969b) when the batch size is constant across all periods. Note that per-period average cost of the T -period model converges to the long-run average cost as $T \rightarrow \infty$. Here, since the cost function $\tilde{c}$ is a step function (i.e. neither convex nor concave in the number of units ordered), it poses difficulties in identifying structures that are preserved through dynamic programming, and the optimal solution for the dynamic program is likely to be quite complex in general. Therefore, we introduce an alternative cost accounting scheme that would be easier to work with in our analysis.

Define

$$c(q) = K \cdot (\lceil q/Q \rceil - q/Q), \quad (4)$$

which is the portion of the set-up cost that can be accounted for under-utilization. That is, if q is not an integer multiple of the base quantity Q , then an amount of $Q \cdot (\lceil q/Q \rceil - q/Q)$ could have been produced without an additional set-up. If q is indeed an integer multiple of Q , then there is no under-utilization of the set-up, and hence the above expression is zero. Also let

$$C(x_t, y_t) = c(y_t - x_t) + \mathcal{L}(y_t). \quad (5)$$

The following theorem states that minimizing the long-run average cost in terms of C is equivalent to minimizing the long-run average of the original cost $\tilde{C}$ up to an additive constant. Let C^* denote the optimal long-run average cost in terms of the C function. All proofs are in the appendix. The main idea behind the proof of Theorem 1 is that the average ordering cost is at least K/Q per unit (thus, at least $(K/Q) \cdot E[D]$ per period) regardless of the ordering policy.

Theorem 1. *There exists an average-cost optimal policy satisfying $\lim_{T \rightarrow \infty} \sum_{t=1}^T (y_t - x_t)/T = E[D]$. For any policy satisfying $\lim_{T \rightarrow \infty} \sum_{t=1}^T (y_t - x_t)/T = E[D]$, we have*

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T E[\tilde{C}(x_t, y_t)] = \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T E[C(x_t, y_t)] + \frac{K}{Q} \cdot E[D].$$

Moreover, $\tilde{C}^* = C^* + E[D] \cdot K/Q$.

	Original Cost Formulation	Alternate Cost Formulation	Difference
Underage and Overage Cost	$\mathcal{L}(y)$	$\mathcal{L}(y)$	0
Ordering Cost	$\tilde{c}(q) = K \cdot \lceil q/Q \rceil$	$c(q) = K \cdot (\lceil q/Q \rceil - q/Q)$	$(K/Q) \cdot q$
Single-Period Cost	$\tilde{C}(x, y) = \mathcal{L}(y) + \tilde{c}(y - x)$	$C(x, y) = \mathcal{L}(y) + c(y - x)$	$(K/Q) \cdot (y - x)$
Average Cost	$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T E[\tilde{C}(x_t, y_t)]$	$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T E[C(x_t, y_t)]$	$(K/Q) \cdot E[D]$
Optimal Average Cost	$\tilde{C}^*$	C^*	$(K/Q) \cdot E[D]$

Table 1: Comparison of the Original and Alternate Cost Formulations

We can interpret this alternative cost accounting scheme as follows. If each order was an integer multiple of the base quantity Q , then the long-run average set-up cost would have been $E[D] \cdot K/Q$, as every unit procured is eventually sold and the long-run average number of ordered units is the average demand. This is captured by the term $(K/Q) \cdot E[D]$. However, in general, the average set-up cost would be higher, as the orders might include partial batches, and such an additional set-up cost is accounted for in both $\tilde{C}$ and C above. From Theorem 1, we adopt the objective of minimizing the long run average of the alternate cost function C instead of $\tilde{C}$. A comparison between the original and alternate cost formulation is summarized in Table 1.

3 Analysis

We first focus on the single-period problem, for which we show the optimality of a threshold-type policy (Section 3.1). Then, in an infinite-horizon setting, we consider a variant problem, and show how we compute the optimal solution for this problem, which provides a lower bound on the optimal cost (Section 3.2). Then, for the original infinite-horizon formulation, we partially characterize the optimal solution of the original problem and we present an efficient method to solve the problem to optimality (Section 3.3). The analysis of this section is the basis for heuristic policies for the original problem to be described in Section 4.

3.1 The Single-Period Problem and the Myopic Policy

We now consider the *single-period* problem, which is given by

$$\min_{y \geq x} C(x, y) = c(y - x) + \mathcal{L}(y) . \quad (6)$$

We let $y(x)$ denote any inventory policy for the above problem, $y^*(x)$ being the optimal one; this notation makes it explicit that the action y depends on the current state x . Since our ordering cost penalizes for unused space in the batch instead of charging a flat-fee for each batch, we note that

this problem is *not* the same as minimizing the original single-period cost $\tilde{C}(x, y)$, but it can be shown that it is equivalent to minimizing the original last-period problem given in (3) with $t = T$ and $v = K/Q$. Furthermore, the single-period problem considered in Lippman (1969b) differs from our single-period problem since his problem adopts $\tilde{c}(\cdot)$ instead of $c(\cdot)$ in (6).

The reason why we study the single-period problem with respect to the alternate cost scheme $C(x, y)$ as opposed to the original cost scheme $\tilde{C}(x, y)$ is three-fold: (i) The alternate scheme is a better representation of the infinite-horizon long-run average cost, as the trade-offs that the original cost structure brings in for the single-period problem do not necessarily reflect the long-run impact of under-utilizing a batch, namely not having sufficient units for future use. Our numerical results (in Section 4.1) support this point. (ii) The single-period problem based on the alternate cost scheme exhibits desirable structural results based on thresholds (Theorem 2). We note that Lippman (1969b)'s single-period analysis establishes similar structure properties but for the original cost function. (iii) The notation introduced here for the single-period problem will become useful in defining policies and bounds for the multi-period problem, providing a direct link between the single-period and multi-period problems. Such a connection has not been established in the earlier literature.

Since the base quantity is Q , the modular arithmetic based on Q plays an important role in our analysis. For any integer z , we define $[z] = z \bmod Q$ such that $[z] \in \{0, 1, \dots, Q-1\}$, i.e., $z = [z] + kQ$ for some integer k . Mathematically, $[\cdot]$ defines a set of equivalence classes.

Recall that the expected holding and backorder cost $\mathcal{L}$ defined in (1) is a convex function, satisfying $\mathcal{L}(y) \rightarrow \infty$ as $|y| \rightarrow \infty$. Let y° be the minimizer of $\mathcal{L}$. (If $\mathcal{L}$ has multiple minimizers, then fix y° at the largest minimizer, to avoid ambiguity in the definition.) We note that y° is an integer due to our integer demand assumption. Also, for each equivalent class $[z]$, let $\theta^{[z]}$ be the largest member of this class not exceeding y° , i.e.,

$$\theta^{[z]} = \max\{w \mid w \leq y^\circ, [w] = [z]\} .$$

Clearly, $y^\circ - Q < \theta^{[z]} \leq y^\circ$ for any z . It is useful to define the following sets of size Q :

$$\mathcal{S}^L = \{y^\circ - Q + 1, y^\circ - Q + 2, \dots, y^\circ\} \quad \text{and} \quad \mathcal{S}^R = \{y^\circ, y^\circ + 1, \dots, y^\circ + Q - 1\} .$$

Notation	Description
$\mathcal{L}(\cdot)$	Expected holding and backorder cost
y°	Minimizer of $\mathcal{L}$
$\mathcal{S}^L$	$\{y^\circ - Q + 1, y^\circ - Q + 2, \dots, y^\circ\}$
$\mathcal{S}^R$	$\{y^\circ, y^\circ + 1, \dots, y^\circ + Q - 1\}$
$[z]$	$z \bmod Q$
$\theta^{[z]}$	$\max\{w \mid w \leq y^\circ, [w] = [z]\}$
$\mathcal{Y}$	Consecutive integers corresponding to the Q smallest values of $\mathcal{L}$
$y^{[z]}$	Unique member y of $\mathcal{Y}$ such that $[y] = [z]$

Table 2: Notation Introduced in Section 3.1

Note that the set $\{\theta^{[z]} \mid 0 \leq z < Q\}$ is the same as $\mathcal{S}^L$. Also, define $\mathcal{Y}$ to be a set of Q consecutive integers that correspond to the Q smallest values of the convex function $\mathcal{L}$, i.e., $|\mathcal{Y}| = Q$, and $\mathcal{L}(y') \leq \mathcal{L}(y'')$ for any $y' \in \mathcal{Y}$ and $y'' \notin \mathcal{Y}$. Note that $\mathcal{Y} \subset \mathcal{S}^L \cup \mathcal{S}^R$. For any integer z , let $y^{[z]}$ be the unique member y of $\mathcal{Y}$ such that $[y] = [z]$. (See Table 2 and Figure 1.) For example, if demand is discrete uniform with support $\{3, 4, 5, 6\}$ and $h = 1$ and $b = 2$, then $\mathcal{L}(y)$ given in (1) assumes the values of 7, 5, 3, 1.75, 1.25, 1.5, 2.5 and 3.5 for $y = 1, \dots, 8$, respectively. Thus, if $z = 1$, then $y^{[z]} = \theta^{[z]} = 5$. If $z = 2$, then $y^{[z]} = 6$ and $\theta^{[z]} = 2$.

If there is no batch ordering (i.e., $Q = 1$), then the optimal policy would have been simply ordering up to y° policy for each period. However, with the batch setting, there would not be a single order-up-to point; instead, the myopic policy would be ordering up to the interval that contains Q lowest values of $\mathcal{L}$, i.e., the set $\mathcal{Y} = \{y^{[z]} \mid 0 \leq z < Q\}$, and $y^{[z]}$ represents the member of this set that can be reached from z using full batches. In the infinite-horizon problem where the ending inventory level in a period becomes a lower bound for the after-ordering inventory level in the next period, it may be prudent to order less than the myopic solution, and thus we also consider the set $\mathcal{S}^L = \{\theta^{[z]} \mid 0 \leq z < Q\}$.

The main result of this section is the characterization of the optimal policy for the single-period problem (Theorem 2 below). In proving this result, we establish a partial characterization of the optimal policy (see Proposition 9 in the appendix) – that it is optimal not to order if the beginning inventory level exceeds y° , and otherwise the after-ordering inventory level must belong to a specific subset of $\mathcal{S}^L \cup \mathcal{S}^R = \{y^\circ - Q + 1, y^\circ - Q + 2, \dots, y^\circ + Q - 1\}$. It is shown that it suffices to specify the optimal policy for a subset of the state space in (6), namely the set $\mathcal{S}^L = \{y^\circ - Q + 1, \dots, y^\circ\}$. (If $x \geq y^\circ$, then it is optimal not to order; if $x \leq y^\circ - Q$, then the order-up-to level is the same as

Figure 1: Illustration for the definitions of sets $\mathcal{S}^L$, $\mathcal{S}^R$ and $\mathcal{L}$ when $Q = 4$.

Figure 2: Illustration for the definitions of $\tilde{\theta}$ and $\underline{\theta}$ when $Q = 4$.

Figure 3: Illustration of the optimal policy for the single-period problem given in Theorem 2.

that of x' where $x' \in \mathcal{S}^L$ and $[x] = [x']$.) Furthermore, for any $x \in \mathcal{S}^L$, it is adequate to consider a restricted action space given by $y(x) \leq Q + x$, i.e., ordering at most one batch.

Based on the above discussion, we focus on the case of $x \in \mathcal{S}^L$, and it suffices to search the optimal after-ordering inventory level y within the range of $\{x, x+1, \dots, x+Q\}$. Then, $c(y-x)$ defined in (4) can be written as

$$c(y-x) = \begin{cases} (K/Q) \cdot (x+Q-y) & \text{if } y > x \\ 0 & \text{if } y = x \end{cases} = K \mathbf{1}_{\{y > x\}} - \frac{K}{Q}(y-x). \quad (7)$$

Thus, the single-period problem of (6) can be written as follows: For any $x \in \mathcal{S}^L$,

$$\min_{y \in \{x, \dots, x+Q\}} \mathcal{L}(y) + c(y-x). \quad (8)$$

While $\mathcal{L}(y)$ is convex, $c(y-x)$ is concave in y . Thus, the objective function in (8) is neither convex nor concave. However, this function is K -convex, and the optimal solution to this problem can be characterized with two thresholds, as we elaborate below in Theorem 2.

We note that (8) is similar to a capacitated inventory model with a fixed ordering cost studied in the literature, for example, Shaoxiang and Lambrecht (1996) and Shaoxiang (2004), but an important difference arises from the cost structure since (7) is not a typical “fixed ordering cost”. However, the structural properties shown in Theorem 2 are somewhat reminiscent of these papers.

The minimization problem given in (8) has the feasible region of $\{x, \dots, x+Q\}$. The optimal solution to this problem will be either (i) one of the boundary solutions (x or $x+Q$, depending on the smaller of $\mathcal{L}(x)$ and $\mathcal{L}(x+Q)$) or (ii) an interior solution, and we will discuss how to determine the optimal solution between them. Consider first the *boundary* solutions. Since $x \in \mathcal{S}^L$ and $\mathcal{Y}$ is the set of consecutive integers corresponding to the Q smallest values of $\mathcal{L}$, it follows from the definitions of $\mathcal{S}^L$ and $\mathcal{Y}$ that only one of x and $x+Q$ belongs to $\mathcal{Y}$, i.e., $x = y^{[x]}$ or $x+Q = y^{[x]}$ (see Figure 1). Thus, in either case, the cost associated with the boundary solution is $\mathcal{L}(y^{[x]}) + c(y^{[x]} - x) = \mathcal{L}(y^{[x]})$. Meanwhile, if the optimal solution is an interior solution, then it also solves a less-constrained version of this problem given in (9) below. Since that expression (8) depends on y only through $\mathcal{L}(y)$ and $(K/Q) \cdot y$, it is useful to define

$$\tilde{\theta} = \arg \min_{y \in \mathcal{Z}^+} \mathcal{L}(y) - (K/Q) \cdot y, \quad (9)$$

where $\mathcal{Z}^+$ is the set of nonnegative integers. This minimization problem given in (9) is a convex

function minimization problem. We note that since y° is the minimizer of a convex function $\mathcal{L}(\cdot)$, $\tilde{\theta}$ is bounded below by y° (i.e., $y^\circ \leq \tilde{\theta}$). However, $\tilde{\theta}$ may or may not belong to the set $\mathcal{Y}$ or $\mathcal{S}^R$. (See Figure 2(a).)

Thus, the optimal solution to (8) is the boundary solution $y^{[x]}$ if $\tilde{\theta}$ lies outside the feasible region of (8); otherwise, it is obtained by comparing the cost of ordering-up-to $y^{[x]}$ to the cost associated with $\tilde{\theta}$. The optimal solution to (8) depends on, as we shall see, whether the value of $y^{[x]} \in \mathcal{Y}$ exceeds certain thresholds or not. Before we provide a formal statement, we define $\underline{\theta}$ as follows:

$$\underline{\theta} = \begin{cases} \max \left\{ \min \mathcal{Y}, \min \{ \theta \in \mathcal{Z}^+ : \mathcal{L}(\theta) \leq \mathcal{L}(\tilde{\theta}) + \frac{K}{Q} \cdot (\theta + Q - \tilde{\theta}) \} \right\} & \text{if } \tilde{\theta} \leq \max \mathcal{Y}, \\ -\infty & \text{if } \tilde{\theta} > \max \mathcal{Y}. \end{cases} \quad (10)$$

(See Figure 2(b) in for illustration.) It can be verified easily that the quantity $\underline{\theta}$ is well-defined (the inequality in the above minimum operator is satisfied, for example, if $\underline{\theta} = \tilde{\theta}$), and that $\underline{\theta}$ satisfies $\underline{\theta} \leq y^\circ \leq \tilde{\theta}$.

We establish the structural properties of the optimal policy for the single-period problem, which we refer to as the *myopic* policy. The optimal policy given in the following theorem thresholds: an order-up-to level $\tilde{\theta}$ (as defined by (9)) that is valid *if* a less-than-full batch is used, and another parameter $\underline{\theta}$ (as defined by (10)) that dictates *when* a less-than-full batch should be used, ensuring a large enough minimum batch size.

Theorem 2. *For the single-period problem based on (3), the following policy is optimal:*

$$y(x) = \begin{cases} x & \text{if } x > y^\circ \\ y^{[x]} & \text{if } x \leq y^\circ \text{ and } y^{[x]} \in \{\underline{\theta}, \dots, \tilde{\theta}\} \\ \tilde{\theta} & \text{if } x \leq y^\circ \text{ and } y^{[x]} \notin \{\underline{\theta}, \dots, \tilde{\theta}\}. \end{cases}$$

We discuss the properties of the solution given in Theorem 2. The first case ($x > y^\circ$) corresponds to the case where there is too much inventory initially, and thus no additional units are ordered. In the remaining two cases, the optimal order quantity corresponds to either a full-batch solution, $y^{[x]}$, or a partial-batch solution, $\tilde{\theta}$. (Recall the definition of $\tilde{\theta}$ in (9) that accounts for the partial-batch penalty. It may be the case that $y^{[x]} = \tilde{\theta}$, in which case we have a full-batch solution.) These two cases are illustrated in Figure 3. We first identify $y^{[x]}$ by adding enough multiples of the base quantity Q to x such that $y^{[x]}$ belongs to the set $\mathcal{Y}$. If $y^{[x]}$ is between $\underline{\theta}$ and $\tilde{\theta}$, then order $y^{[x]} - x$ (indicated by Arrow 2 in the figure). Otherwise, the ordering quantity is adjusted such that the after-ordering inventory level becomes $\tilde{\theta}$ (Arrow 1 or 3).

We note that there is an asymmetry in the structure of the optimal policy around $\tilde{\theta}$. If $y^{[x]} = \tilde{\theta} + 1$, decreasing the after-ordering inventory level from $y^{[x]}$ to $\tilde{\theta}$ costs only K/Q (see Equation (4)). However, if $y^{[x]} = \tilde{\theta} - 1$, then increasing the after-ordering inventory level from $y^{[x]}$ to $\tilde{\theta}$ costs $(K/Q) \cdot (Q - 1)$ which can be much larger than K/Q , and thus it is not as attractive to change the inventory level from $y^{[x]}$. Hence, it is optimal to round *down* the order-up-to level to $\tilde{\theta}$ if the full-batch solution $y^{[x]}$ is slightly above the target level $\tilde{\theta}$; however, if $y^{[x]}$ is slightly below the target level $\tilde{\theta}$, then it is optimal to keep $y^{[x]}$ and *not* round up the order-up-to level to $\tilde{\theta}$.

The following result is a corollary of Theorem 2, and shows structural relations between the parameters of the optimal policy, $\underline{\theta}$ and $\tilde{\theta}$, and the set-up cost K . In particular, when K is sufficiently small, then it is optimal to follow the policy of ordering up to $\tilde{\theta}$, which itself converges to y° . This is the basic base-stock policy, which is optimal when there is no batch constraint. However, as K becomes arbitrarily large, it is not optimal to order a partial batch, and the optimal policy is order-up-to $y^{[x]}$ for any x . This is exactly the (R, nQ) policy for the batch-ordering problem with full-batches only (see, for example, Veinott, 1965).

Corollary 3. *The value of $\tilde{\theta}$ is non-decreasing in K , and the value of $\underline{\theta}$ is non-increasing in K . Furthermore, there exist nonnegative numbers $\underline{K}$ and $\overline{K}$ such that, for the single-period problem with starting inventory level x , the order-up-to $\tilde{\theta}$ policy is optimal if $K \leq \underline{K}$ and the order-up-to $y^{[x]}$ policy is optimal if $K \geq \overline{K}$.*

The myopic policy of this section will be used later in Section 4.1 to motivate a heuristic policy for the infinite-horizon long-run average cost problem.

3.2 Relaxed Problem and the Reduced MDP Approach

The infinite-horizon problem is difficult to analyze because of the lack of convexity-like structures in the ordering cost. In this section, we introduce a relaxation of the original problem that is relatively straightforward to analyze and solve. This relaxation provides a lower bound on the cost for the original problem. Furthermore, it motivates the development of a heuristic policy, which will be further discussed in Section 4.1 (this policy, as we shall see, performs very well computationally).

For this relaxation, we no longer impose the constraint $y_t \geq x_t$ in each period, and allow the possibility that inventory can be scrapped. The cost of scrapping inventory is also given by the

same c function defined in (4), even for the negative q values; as a result, scrapping 1 unit has the same cost as ordering $Q - 1$ units in our relaxed model. This modification is motivated by (i) how modular arithmetic plays an important role in analyzing inventory problems with batches, and (ii) an observation that ordering $Q - 1$ units has the same cost as order $2Q - 1$ units (corresponding to the amortized fixed cost associated with a single unit). We refer to this problem as the *infinite-horizon relaxed problem*. The following lemma partially characterizes the optimal policy. (The proof of this lemma is obtained from the fact that $\mathcal{Y}$ contains the Q lowest values of the $\mathcal{L}$ function.)

Lemma 4. *For the infinite-horizon relaxed problem, there exists an optimal policy such that $y_t(x_t) \in \mathcal{Y}$ for any starting inventory level x_t and period t .*

Lemma 4 shows that, for the reduced problem, the optimal action $y_t(x_t)$ in any period t can be restricted to the set $\mathcal{Y}$ for any starting inventory level x_t . Furthermore, since increasing or decreasing inventory by multiples of Q is “costless”, the optimal actions associated with two starting inventory levels are the same, provided that these starting inventory levels differ by a multiple of Q , i.e., $y_t(x_t) = y_t(\hat{x}_t)$ if $[x_t] = [\hat{x}_t]$. Thus, the optimal action depends on x_t only through the modular arithmetic class to which it belongs, i.e., $[x_t]$, and the state space can be thought of any set of Q consecutive integers, for example, the set $\mathcal{Y}$.

Therefore, we can define a *reduced* Markov Decision Process (MDP) with the state space corresponding to $\mathcal{Y}$ (for the inventory level before ordering) in the modular arithmetic and the action space $\mathcal{Y}$ (for the inventory level after ordering). (Recall that $\mathcal{Y}$ is a set of consecutive integers attaining the Q smallest values of $\mathcal{L}$.) This MDP has Q states, and there are Q possible actions available at each state. We now specify the components of this MDP in detail. The cost function associated with ordering from x_t to $y_t \in \mathcal{Y}$ is given by

$$C([x_t], y_t) = c([y_t - x_t]) + \mathcal{L}(y_t) \quad \text{where} \quad c([y_t - x_t]) = (K/Q) \cdot ([x_t - y_t]) .$$

For the transition probability, we define $p^{[i]} = P\{(D \bmod Q) = i\}$, which is the probability that the demand belongs to the set $\{i + kQ \mid k = 0, 1, 2, \dots\}$. Clearly, $p^{[0]} + p^{[1]} + \dots + p^{[Q-1]} = 1$. Then, the probability that the next state is $[x_{t+1}]$ given the current state-action pair of $([x_t], y_t)$ is $p^{[y_t - x_{t+1}]}$. We refer to this reduced MDP as $\mathcal{M}$. Let $\hat{C}$ be the steady-state cost of $\mathcal{M}$, and let $\hat{y}([x])$ denote the optimal action for the state $[x]$ in $\mathcal{M}$. (The optimal action is independent of t since

a stationary policy is optimal under long-run average-cost criterion when the set of state-action pairs is finite by Theorem 4-14 of Heyman and Sobel (1984).) The next theorem states that $\hat{C}$ is a lower bound on the optimal “alternate cost” C^* , and this result follows since the reduced MDP disregards one of the original constraints $y_t \geq x_t$. (Recall that $\tilde{C}^*$ is the optimal original cost.)

Theorem 5. *For the infinite-horizon relaxed problem, the optimal policy is given by $y_t(x_t) = \hat{y}([x_t])$, and the long-run average cost is $\hat{C}$. Furthermore, $\hat{C} \leq C^* = \tilde{C}^* - E[D] \cdot K/Q$.*

We conclude this section with the following observation that the myopic policy of Section 3.1 and the reduced MDP policy of this section coincide under a certain technical condition. This condition, that the demand belongs to equivalence classes with equal probabilities, is a reasonable approximation, especially for platykurtic distributions (negative excess kurtosis), if the support of demand is large compared to the base quantity Q . For example, with $Q = 3$ and the Poisson distribution with mean 5, the maximum deviation of $p^{[i]}$ from $1/Q$ is less than 1%, and it is less than 0.01% with mean 10.

Proposition 6. *Suppose that $p^{[i]} = 1/Q$ holds for each $i \in \{0, 1, \dots, Q - 1\}$. Then, the myopic policy is optimal for $\mathcal{M}$, i.e., $\hat{y}([x]) = y_T(\tilde{x})$ where $y_T(\cdot)$ is the optimal myopic solution given in Theorem 2 and $\tilde{x}$ is the largest number such that $\tilde{x} \leq y^\circ$ and $[x] = [\tilde{x}]$.*

The proof of the above proposition follows since the probability distribution of the ending inventory in terms of the equivalent classes is independent of the starting inventory level. (Note that the transition probability from the state-action pair $([x_t], y_t)$ to $[x_{t+1}]$ is given by $p^{[(y_t - x_{t+1}) \bmod Q]}$, which is $1/Q$ under the condition of Proposition 6; since this probability is independent of the action y_t , the optimal choice of y_t is to minimize the single-period cost in the current period.)

3.3 The Original Problem and Optimal Solution

We now analyze the original multi-period problem. In particular, we first state a proposition that partially characterizes the optimal solution, and present an efficient method to solve the problem to optimality, which makes use of this partial characterization. Finally, we elaborate on the optimal policy structure, discussing why it does not resemble any of the traditional inventory control policies.

The following theorem establishes the upper and lower bounds on the optimal action $y^*(x)$ and states some optimality properties. The proof of this theorem is based on a sample path argument

that if an optimal policy violates one of these properties, then we can construct another policy that either orders zero or more full batches or mimics the original policy, such that the cumulative cost of the new policy is at least as low as the cumulative cost of the original policy.

Theorem 7. *There exists an optimal stationary policy $y^*(x)$ for the original infinite-horizon problem satisfying the following properties:*

- (a) *For any $x \geq y^\circ$, $y^*(x) = x$.*
- (b) *For any $x < y^\circ$, $y^*(x) \in \{y_l(x), y_l(x)+1, \dots, y_u(x)\}$ where $y_l(x) = \theta^{[x]}$ and $y_u(x) = \min\{\max \mathcal{Y}, \theta^{[x]} + Q\}$.*
- (c) *For any pair of $x', x'' < y^\circ$ satisfying $[x'] = [x'']$, $y^*(x') = y^*(x'')$.*

This theorem sharpens an earlier partial characterization of Lippman (1969b, Theorem 8) which is based on finite horizon dynamic programming: (i) All the quantities used in the statement of Theorem 7 are given in terms of the elementary data of the problem specification whereas the statement of Lippman’s result includes the solution to the dynamic programming formulation and thus it is not readily accessible. (ii) Theorem 7 implies that it is optimal to order if the starting inventory is below a threshold $y^\circ - Q$ and not to order if the starting inventory is above another threshold y° . In comparison, while Theorem 8 of Lippman establishes an upper threshold (above which is optimal not to order) under a general condition, a lower threshold (Lemma 9 of Lippman) is provided only when a certain relationship between $\mathcal{L}$ and Q is satisfied. (This relationship, in our setting, corresponds to $b > K/Q$, which is reasonable but not superfluous.) Furthermore, we show that the gap between the two thresholds is at most Q , a result not established by Lippman’s paper. (iii) Our result shows that the order-up-to level belongs to the set of consecutive integers $\{\min_x y_l(x), \dots, \max_x y_u(x)\}$, which in turn is a subset of $\{\theta^{[x]}, \dots, \theta^{[x]} + Q\}$ containing $Q+1$ elements, while Lippman’s result cannot easily provide the set of potential order-up-to levels.

We note that a result similar to Theorem 7 is shown by Caliskan-Demirag et al. (2010) for a finite-horizon case. Their lower and upper bounds, corresponding to $\theta^{[x]}$ and $\theta^{[x]} + Q$ in part (b), are not as tight as ours. Zhang and Çetinkaya (2011) have the same bounds, but their characterization of the optimal policy may have an arbitrary number of reorder and order-up-to points.

In what follows, we propose an MDP approach to solve the original problem to optimality, making use of Theorem 7 to define the state space and action space, each of which is restricted to the $O(Q)$ in size. Note that parts (a) and (b) of Theorem 7 imply that the action space at each state has at most $Q+1$ elements. Furthermore, it suffices to consider an MDP in which the state space is bounded below by $y^\circ - Q$ (by part (c) of Theorem 7) and above by $\max \mathcal{Y}$ (by parts (a) and (b)). It follows from the definitions of y° and $\mathcal{Y}$ that the number of integers between these two bounds are at most $2Q$. Consequently, the state-action pair space of this MDP has at most $(2Q) \cdot (Q+1)$ elements, and we can solve it by using value iteration, for example via a linear program with these many variables (Chapter 4 of Heyman and Sobel, 1984). In this MDP, the cost function associated with ordering from x_t to y_t is given by $c(y_t - x_t)$ and $C(x_t, y_t)$ as given by Equations (4) and (5). We prescribe the transition probabilities referring to $p^{[i]}$ as defined in Section 3.2. The transition probability to x_{t+1} from y_t is $p^{[y_t - x_{t+1}]}$ if $x_{t+1} \geq \min \mathcal{Y}$. For $x_{t+1} < \min \mathcal{Y}$, the transition probability is $p^{[y_t - x_{t+1}]}$ if $y_t - x_t \leq Q$, and $(p^{[y_t - x_t]} - P\{D \equiv [y_t - x_t]\})$ otherwise.

While our method for the original problem can be carried out easily, the optimal policy for this problem does not exhibit a simple structure; see Figure 4(a), showing that the optimal policy is a mixture of re-order level, (multiple) order-up-to level, and full-batch ordering policies. Thus, such a policy cannot be mimicked by simple policies such as the base-stock policy, (s, S) policy, (R, nQ) policy (full batch ordering policy, FBO), or a newsboy-type policy, which we refer to as the traditional policies (Figure 4(a)). In fact, these traditional policies do not perform well numerically for the whole range of parameters considered in our experiment described later in Section 4. The presence of multiple order-up-to levels appears to be positively related to the degree of demand certainty based on our numerical investigation (described later in Section 4) – the multiple order-up-to levels in the optimal order policy have been found in 24, 33, 35 and 54 out of 144 problem instances for each demand coefficient of variation of 0.15, 0.10, 0.05 and 0.01, respectively.

4 Heuristics and Traditional Approaches

Even though the solution method in Section 3.3 is efficient, one may want to improve this computational efficiency further to obtain close-to-optimal solutions almost instantaneously. Such an efficiency may be essential to solve the problem on a frequent basis (e.g., the Merloni case) or to

Figure 4: The Optimal and Other Policies for a Multiple-Period Problem: Negative binomial demand distribution with mean 25 and coefficient of variation 0.21; $h = 1$, $b = 50$, $K = 146$, and $Q = 91$. Order-up-to points of the policies are depicted only when different from the optimal.

conduct sensitivity analysis for a sizable set of problem parameters. To examine this, we propose in Section 4.1 heuristics based on the alternate cost accounting scheme (which we refer to as the “heuristic policies”), and test their performances. In Section 4.2, we investigate some traditional ordering policies that are frequently adapted in practice and studied in the inventory literature, and identify when it is safe to operate with these traditional policies.

4.1 Heuristic Policies

The heuristics that we propose in this section are based on our discussions in Sections 3.1 and 3.2. They provide insights and perform extremely well, coinciding with the optimal solution under some conditions. For each of these heuristic policies, the inventory level after ordering belongs to $\mathcal{Y}$, an interval of length Q , but they differ regarding which element of $\mathcal{Y}$ is chosen depending on the starting inventory level in each period.

(i) INTERVAL-BASED POLICY. This policy is motivated by the analysis of the single-period problem in Section 3.1. It is characterized by two parameters θ_L and θ_U , both of which belong to $\mathcal{Y}$. We denote this policy by $IB(\theta_L, \theta_U)$, and specify as follows:

$$y_t(x_t) = \begin{cases} x_t & \text{if } x_t > y^\circ, \\ y^{[x_t]} & \text{if } x_t \leq y^\circ \text{ and } y^{[x_t]} \in [\theta_L, \theta_U], \\ \theta_U & \text{if } x_t \leq y^\circ \text{ and } y^{[x_t]} \notin [\theta_L, \theta_U]. \end{cases}$$

(For a pictorial representation, see Figure 3, where $\underline{\theta}$ is replaced with θ_L and $\tilde{\theta}$ with θ_U .) We interpret this policy as a generalization of the base-stock policy. If the initial inventory exceeds y° ,

then we do not order. Otherwise, we consider θ_U as a target order-up-to level. If the batch size to reach θ_U is at least of size $\theta_U - \theta_L$, then we order-up-to θ_U ; otherwise, we do not want to incur the set-up cost K for a small number of units, and we order up to $y^{[x_t]}$ instead. We refer to this policy as the IB policy. Note that the IB policy becomes a base-stock policy if $\theta_L = \theta_U$. Another example of the IB policy is the myopic policy, which is optimal for the single period problem in Section 3.1; in the myopic policy, we set $\theta_L = \underline{\theta}$ and $\theta_U = \tilde{\theta}$. The optimal choice of parameters (θ_L, θ_U) for the IB policy can be attained by solving a two-dimensional optimization problem; in this paper, we use complete enumeration.

We consider the IB policy since threshold-type policies have been widely used in inventory management for their implementability. Furthermore, this policy has been mentioned in the literature for a single-period problem (e.g., Lippman (1969b)). It is equivalent to the single-product version of the (Q, S) policy in Cachon (2001), where a minimum load is imposed. In our paper, the motivation of this policy comes from its optimality for the single period problem; note that we have reached the same policy structure using a different approach, strengthening the attractiveness of this class of policies.

(ii) REDUCED MDP-BASED POLICY. This policy is motivated by the relaxed MDP problem of Section 3.2, which provides a lower bound $\hat{C}$. It is possible that the order-up-to level specified by the *relaxed* problem may not be attainable because scraping inventory is not allowed in the original problem. In this case, we take the optimal policy of the relaxed problem as a target base-stock level for the original problem. We refer to this policy as the RMB policy, which is specified as follows:

$$y_t(x_t) = \max\{\hat{y}([x_t]), x_t\} .$$

To implement this policy, we first need to solve the relaxed MDP problem $\mathcal{M}$, which has Q states, each with Q possible actions. Unlike the IB policy, this policy does not have any parameter and thus there is no need to search through the parameter space. While the inventory decision in each period makes an impact on future decisions by constraining the after-ordering inventory of the next period as well as changing the *modularity* of the starting inventory in the next period, the reduced MDP-based policy incorporates the latter impact, ignoring the former. As we will see later, this policy performs quite well, suggesting that modularity issues play an important role in our problem.

We have conducted a numerical experiment evaluating the performance of these two policies, (i) and (ii) above, as well as the myopic policy introduced in Section 3.1. We set the mean demand to 25 units per period, and vary the coefficient of variation (CV) among $\{0.05, 0.25, 0.5, 1.0, 1.5\}$, and let $h = 1$, $b \in \{2, 5, 10, 50, 100\}$, $K \in \{2, 5, 10, 50, 100, 200\}$, and $Q \in \{5, 10, 25, 50, 100, 200\}$. We model demand distribution using Negative Binomial whenever possible, and we make the only exception when $CV=0.05$ by using discretized Gamma, since this CV is not attainable with Negative Binomial when the mean demand is 25. We have tested all possible combinations of these parameters resulting in $5 \times 5 \times 6 \times 6 = 900$ problem instances. For each problem, we compare the average cost C of each heuristic to the average cost of the optimal policy using the alternate cost function, and report the *relative cost error* of each heuristic (which is the difference in costs divided by the optimal cost), denoted by Δ , where a subscript indicates the type of a heuristic policy. (Using the original cost $\tilde{C}$ would have included the “minimum fixed cost”, and would have made the relative error smaller.)

The IB policy, whose parameter values (θ_L, θ_U) have been optimized through search, provides near-optimal results. The relative cost error of the IB policy, Δ_{IBP} , is 0.04% on average, and 2.92% in the worst instance (Table 3). In fact, the IB policy coincides with the optimal policy in 853 of the total 900 problem instances. Our experiments show that the RMB policy also performs exceptionally well, even better than the IB policy, resulting in the relative cost error, Δ_{RMBP} , of 0.003% on average and 0.353% at maximum (Table 4). Furthermore, the RMB policy has the same cost as the optimal policy in all but 29 cases among 900 problem instances. Thus, we conclude that the IB policy and the RMB policy have both performed extremely well in our experiments.

We now state the following theorem that relates heuristic policies with the optimal policy. While the technical condition of the following theorem does not hold in our experiments, it can be argued that these conditions “approximately” hold in many settings, thus providing a plausible explanation of why our heuristic policies perform very well.

Theorem 8. *Suppose that $D \geq Q$ with probability 1. Then, the following statements are true for the original problem: (a) The RMB policy is optimal; (b) If $p^{[i]} = 1/Q$ holds for each $i \in \{0, 1, \dots, Q-1\}$, then the myopic policy is optimal.*

Part (a) of the above theorem is based on an observation that the condition $D \geq Q$ ensures that

Table 3: Aggregate Summary of Δ_{IBP} Values (in %) Over All b and CV Values

Average [min, max]	$Q = 5$	$Q = 10$	$Q = 25$	$Q = 50$	$Q = 100$	$Q = 200$	Overall
$K = 2$	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]
$K = 5$	0.00 [0.00, 0.00]	0.05 [0.00, 1.05]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.01 [0.00, 1.05]
$K = 10$	0.00 [0.00, 0.00]	0.08 [0.00, 0.92]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.01 [0.00, 0.92]
$K = 50$	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.54 [0.00, 2.92]	0.00 [0.00, 0.00]	0.00 [0.00, 0.02]	0.09 [0.00, 2.92]
$K = 100$	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.22 [0.00, 1.38]	0.00 [0.00, 0.00]	0.01 [0.00, 0.16]	0.04 [0.00, 1.38]
$K = 200$	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.06 [0.00, 1.54]	0.08 [0.00, 0.51]	0.26 [0.00, 1.81]	0.18 [0.00, 2.15]	0.10 [0.00, 2.15]
Overall	0.00 [0.00, 0.00]	0.02 [0.00, 1.05]	0.01 [0.00, 1.54]	0.14 [0.00, 2.92]	0.04 [0.00, 1.81]	0.03 [0.00, 2.15]	0.04 [0.00, 2.92]

Table 4: Aggregate Summary of Δ_{RMBP} Values (in %) Over All b and CV Values

Average [min, max]	$Q = 5$	$Q = 10$	$Q = 25$	$Q = 50$	$Q = 100$	$Q = 200$	Overall
$K = 2$	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]
$K = 5$	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]
$K = 10$	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]
$K = 50$	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.01 [0.00, 0.21]	0.01 [0.00, 0.16]	0.01 [0.00, 0.14]	0.00 [0.00, 0.21]
$K = 100$	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.02 [0.00, 0.31]	0.02 [0.00, 0.23]	0.01 [0.00, 0.31]
$K = 200$	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.01 [0.00, 0.15]	0.02 [0.00, 0.35]	0.01 [0.00, 0.33]	0.01 [0.00, 0.35]
Overall	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.00 [0.00, 0.21]	0.01 [0.00, 0.35]	0.01 [0.00, 0.33]	0.00 [0.00, 0.35]

Table 5: Aggregate Summary of Δ_{MP} Values (in %) Over All b and CV Values

Average [min, max]	$Q = 5$	$Q = 10$	$Q = 25$	$Q = 50$	$Q = 100$	$Q = 200$	Overall
$K = 2$	0.30 [0.00, 1.73]	0.00 [0.00, 0.01]	0.16 [0.00, 2.34]	0.02 [0.00, 0.23]	0.03 [0.00, 0.22]	0.02 [0.00, 0.08]	0.09 [0.00, 2.34]
$K = 5$	1.91 [0.00, 7.70]	0.84 [0.00, 6.68]	1.26 [0.00, 19.81]	0.04 [0.00, 0.19]	0.06 [0.00, 0.26]	0.09 [0.00, 0.50]	0.70 [0.00, 19.81]
$K = 10$	2.22 [0.00, 16.11]	2.23 [0.00, 15.95]	2.48 [0.00, 23.37]	0.18 [0.00, 1.53]	0.14 [0.00, 0.73]	0.28 [0.00, 1.13]	1.25 [0.00, 23.37]
$K = 50$	11.23 [0.00, 83.43]	8.45 [0.00, 73.36]	30.62 [0.00, 168.57]	8.41 [0.01, 43.05]	16.16 [0.29, 49.88]	15.27 [0.79, 37.79]	15.02 [0.00, 168.57]
$K = 100$	22.49 [0.00, 167.57]	16.72 [0.00, 145.12]	22.68 [0.00, 122.82]	7.18 [0.00, 34.67]	6.97 [0.30, 31.24]	31.57 [1.02, 88.49]	17.93 [0.00, 167.57]
$K = 200$	45.01 [0.00, 335.85]	33.27 [0.00, 288.65]	16.68 [0.00, 84.22]	6.77 [0.00, 26.05]	5.10 [0.00, 65.53]	20.28 [3.37, 40.51]	21.18 [0.00, 335.85]
Overall	13.86 [0.00, 335.85]	10.25 [0.00, 288.65]	12.31 [0.00, 168.57]	3.77 [0.00, 43.05]	4.74 [0.00, 65.53]	11.25 [0.00, 88.49]	9.36 [0.00, 335.85]

the constraint that $y_t \geq x_t$ is always satisfied in the reduced MDP of Section 3.2. Furthermore, part (b) follows from the logic used in Proposition 6 that the ordering decision in a period does not constrain the optimal decision in the next period, and thus it makes sense to minimize the single-period cost in the current period. A similar argument is used by Zhang and Çetinkaya (2011) to show the structure of the optimal policy when demand is uniformly distributed between 0 and an integer multiple of the batch size.

Note that when the condition of (b) in Theorem 8 is applicable, the IB policy is also optimal since the myopic policy is a special case of the IB policy. While both IB policy and RMB policy give close-to-optimal solutions, we have tested the myopic policy (MP) which might be a preferred alternative in practice as it requires resolution of simpler mathematical expressions. Our experiments show that the myopic policy results in the average relative cost error Δ_{MP} of 9.36% over all 900 problem instances (see Table 5 for details). We note that if K is large, myopic policy does not perform well, and therefore it is not recommendable; meanwhile, myopic policy performs well for lower values of K , which might be expected since the myopic policy is optimal for the special case with $K = 0$.

We have compared the three heuristic policies to the optimal policy, as in Figure 4(b). We have observed that IBP and MP tend to track the optimal policy better than MP or traditional policies.

4.2 Traditional Policies

In this section, we test the performance of the full batch ordering (FBO) policy, (s, S) policy, and “original myopic” (OM) policy, i.e. myopic policy with the original cost primitives, and we assess their capability in providing a satisfactory solution. We study FBO policy since many of the papers addressing batch ordering in the literature optimizes within this class of policy. The best possible FBO policy is considered here by a rather straightforward adaption of the dynamic programming model that we presented in Section 2.2. Note that the Δ_{FBO} values reported here are the best achievable bounds within FBO policies, and the full batching ordering policy implemented in practice may have worse performance. We study (s, S) policy since it is commonly implemented in problem environments with fixed a cost component. We optimize both parameters of this policy. Finally, we study myopic policy based on the original cost function, $\tilde{C}(x_t, y_t)$ in Section 2.1, with the modification that the order-up-to point is constrained to be non-negative. (The reason for this

modification is to overrule the possibility of infinite backordering when the backorder cost is not high enough to recoup the set-up (and the holding) cost, as discussed in Section 3.1.)

(i) FBO POLICY. Table 6(a) shows that FBO coincides with the optimal policy when Q is low and K is high, as expected (due to increased importance of high utilization), but it is not sensible for the other extreme of high Q and low K . In the table, we have restrained ourselves from reporting the relative cost errors larger than 50%, which we denote by N/A. We note that this insight has been confirmed by the experiments conducted by Caliskan-Demirag et al. (2010), and it is also consistent with Corollary 3 which shows that as K becomes sufficiently large, the full batch ordering becomes asymptotically optimal. (By the same logic, the full batch ordering is optimal as Q becomes very small.)

We observe from Table 7(a) that, as CV increases, FBO performs much better. This is because the optimal policy tends to order in full batches for high CV values since the single period cost function tends to be flatter. In comparison, the optimal policy consists of multiple order-up-to points for low CV values which is not possible to mimic using the FBO policy. To verify this, we have measured the *weighted batch utilization* values (average base quantity occupied by an order relative to the full base quantity, weighted by the base quantities) – see Table 7(b).

In practice, full batch ordering policy is favored when the backorder costs are relatively high (Tanrikulu et al., 2010). The performance of FBO in Table 7(a) displays a trend of improvement with respect to b . This may be explained by the fact that the inventory-related cost (holding and backorder cost) increases in b . However, the interaction among the problem parameters are intricate, and several problem instances do not demonstrate this monotonicity with respect to b , and the above intuition is not necessarily correct. Finally, we comment that, our observations with respect to CV and b are not unique to the FBO case, and hold for all the other policies.

(ii) (s,S) POLICY. Table 6(b) shows that (s, S) policy performs very well when Q is high and K is relatively low. Note that (s, S) policy is the optimal policy when Q is sufficiently large – when Q is high enough, the optimal order quantity is likely to fit within one base quantity. However, as K increases, high batch utilization becomes more and more important, even for large base quantities, which cannot be assured by (s, S) . Therefore, when the problem parameters are such that high

Table 6: Relative Cost Error (in %) Averaged over all b and CV Values

(a) Δ_{FBO} Values (in %)							(b) $\Delta_{(s,S)}$ Values (in %)						
K	Q						K	Q					
	5	10	25	50	100	200		5	10	25	50	100	200
2	3.70	11.59	32.94	N/A	N/A	N/A	2	3.59	1.42	2.34	0.06	0.01	0.00
5	0.81	2.03	13.88	N/A	N/A	N/A	5	11.12	5.32	7.93	0.24	0.03	0.00
10	0.00	0.41	6.52	45.87	N/A	N/A	10	26.59	20.09	17.71	0.67	0.08	0.01
50	0.00	0.00	0.00	22.12	5.23	7.60	50	N/A	N/A	N/A	24.56	0.92	0.05
100	0.00	0.00	0.00	14.70	0.00	5.84	100	N/A	N/A	N/A	N/A	3.32	0.16
200	0.00	0.00	0.00	8.39	0.00	4.36	200	N/A	N/A	N/A	N/A	14.25	0.52

(c) Δ_{OM} Values (in %)						
K	Q					
	5	10	25	50	100	200
2	2.43	0.64	0.29	0.04	0.03	0.03
5	14.78	3.50	2.46	0.19	0.14	0.13
10	N/A	19.29	6.86	0.63	0.45	0.42
50	N/A	N/A	N/A	32.38	21.51	18.05
100	N/A	N/A	N/A	N/A	N/A	40.91
200	N/A	N/A	N/A	N/A	N/A	N/A

Table 7: Statistics Useful for Understanding FBO (Averaged Over All K and Q Values)

(a) Δ_{FBO} Values (in %)						(b) Batch Utilization (in %): Optimal Policy					
b	CV					b	CV				
	0.05	0.25	0.5	1	1.5		0.05	0.25	0.5	1	1.5
2	N/A	N/A	29.27	14.22	10.07	2	58.33	58.84	65.00	76.79	82.83
5	N/A	N/A	26.42	10.24	5.47	5	56.84	57.40	62.44	76.55	84.11
10	N/A	N/A	23.43	7.95	3.59	10	55.98	56.85	61.66	76.78	84.76
50	N/A	47.34	18.02	5.08	1.87	50	56.58	56.43	60.52	76.53	85.18
100	N/A	44.33	16.39	4.38	1.53	100	56.53	56.48	60.33	76.56	85.38

Table 8: Best Performing Policy and the Corresponding Relative Cost Error (in %)

(a) CV = 1.5 and $b = 100$											
(Best Policy, Δ)	$Q = 5$		$Q = 10$		$Q = 25$		$Q = 50$		$Q = 100$		$Q = 200$
$K = 2$	FBO	0.00	FBO	0.00	FBO	0.08	(s, S)	0.04	(s, S)	0.01	(s, S) 0.00
$K = 5$	FBO	0.00	FBO	0.00	FBO	0.01	(s, S)	0.15	(s, S)	0.04	(s, S) 0.00
$K = 10$	FBO	0.00	FBO	0.00	FBO	0.00	FBO	0.28	(s, S)	0.10	(s, S) 0.01
$K = 50$	FBO	0.00	FBO	0.00	FBO	0.00	FBO	0.00	FBO	0.87	(s, S) 0.11
$K = 100$	FBO	0.00	FBO	0.00	FBO	0.00	FBO	0.00	FBO	0.15	(s, S) 0.35
$K = 200$	FBO	0.00	FBO	0.00	FBO	0.00	FBO	0.00	FBO	0.00	(s, S) 1.12

(b) CV=0.05 and $b = 100$											
(Best Policy, Δ)	$Q = 5$		$Q = 10$		$Q = 25$		$Q = 50$		$Q = 100$		$Q = 200$
$K = 2$	(s, S)	0.00	OM	0.00	(s, S)	0.00	OM	0.00	OM	0.00	OM 0.00
$K = 5$	OM	3.11	OM	0.00	OM	2.61	OM	0.00	OM	0.00	OM 0.00
$K = 10$	FBO	0.00	FBO	3.52	OM	9.69	OM	0.00	OM	0.00	OM 0.00
$K = 50$	FBO	0.00	FBO	0.00	FBO	91.28	(s, S)	42.92	(s, S)	0.00	(s, S) 0.00
$K = 100$	FBO	0.00	FBO	0.00	FBO	62.71	FBO	34.63	(s, S)	0.00	(s, S) 0.00
$K = 200$	FBO	0.00	FBO	0.00	FBO	36.49	FBO	25.68	(s, S)	6.99	(s, S) 0.00

batch utilization is a dominating factor, (s, S) policy is not sensible.

(iii) OM POLICY. OM policy performs consistently well when Q is high and K is low, but ceases to perform well as K increases. See Table 5. We also note that OM policy, a myopic policy based on the original single-period cost performs much worse than the myopic policy, based on the alternate formulation.

(iv) BEST PERFORMANCE. As the traditional policies discussed above perform well for particular problem parameters, we now address the question whether it suffices to consider the best performing traditional policy to attain the best practice. In Table 8(a), we present which traditional policy performs the best for $CV=1.5$, $b = 100$, and all K, Q combinations, together with the corresponding relative cost error values. While it appears that one can infer a general pattern, it is nonetheless difficult to identify *a priori* which policy will perform the best (e.g., compare the cases with $CV=0.05$ and 1.5 in Table 8), except for extreme cases, making the decision of selecting the best traditional policy non-trivial. Furthermore, even the best policy (among traditional policies) does not always perform well (up to 91% relative cost error in our test bed).

In closing Section 4, the performance of RMB policy is the closest to the optimal policy, with a negligible performance gap. Unlike IMP and MP, RMB policy is capable of being optimal even if the optimal policy has more than one order-up-to level. We find this policy elegant since its framework is based on the modularity of Q , the batch size. In terms of the solution times, RMB policy yields almost instantaneous solutions (less than a second) in all problem instances we have tested on a 2.93 GHz CPU computer. Furthermore, since this policy has no parameter, it does not require searching for the parameter space.

5 Conclusions

In this paper, we pose the multiple set-up costs problem in an inventory system where partial batch ordering is allowed. We build an infinite horizon dynamic programming model to find the optimal solution. Characterizing the solution of this problem is difficult because even for more restricted versions of this problem, total expected cost functions are not convex and the optimal ordering policies do not have an easily identifiable simple structure. After a cost transformation, we develop a computationally efficient optimal solution method. We also develop several properties of the

optimal solution. Our results aid us to identify two sets of heuristic algorithms. Our numerical results indicate that both heuristics perform extremely well.

Finally, we investigate the performances of three “traditional” policies that are frequently adopted in practice and studied in the literature, discussed the parameter ranges where these policies perform well. We observe that there are a number of prospective problems with using traditional policies. It is difficult to know *a priori* which policy should be applied, except for “trivial” cases. Also, even the best policy can result in substantial errors. In comparison, the maximum relative cost error terms obtained by the proposed heuristic algorithms IB policy and RMB policy are 2.92% and 0.35%, respectively. Furthermore, the solution times of policies such as (s, S) are also prohibitive when the problem needs to be solved on a frequent basis.

This research can be extended in several ways. It may be the case that there are several different options for the possible base quantities, such as different truck capacities. Different batch options might also involve different fixed costs. The alternative cost accounting scheme that we proposed in this paper can also be extended to more general settings, such as multiple items, multiple stages, Markov modulated demand case, etc. A similar accounting scheme could also be applied to other problem environments where there is a fixed cost of utilizing capacitated facilities, such as the stochastic lot sizing problem. Note that the classical capacitated inventory problem with fixed costs is a special case of our problem where there is only one batch. The optimal policy of this problem has not yet been fully characterized in the literature (see for example, Shaoxiang and Lambrecht, 1996; and Gallego and Scheller-Wolf, 2000). The ideas presented in our paper could be a basis for the analysis of that problem.

The methodological approach we have taken in the analysis exploits the discreteness of the units and the infinite-horizon average cost criterion, and it does not immediately extend to the case where either of these assumptions fails. However, we believe that the cost accounting scheme of this paper can be useful even in the continuous-unit setting or the finite-horizon setting to motivate good heuristics and to generate interesting analytical properties.

References

- Alp, O., N. K. Erkip and R. Gullu. 2003. Optimal Lot-Sizing/Vehicle-Dispatching Policies under Stochastic Lead Times and Stepwise Fixed Costs. *Operations Research*. **51**(1). 160–166.

- Cachon, G. 2001. Managing a retailer's shelf space, inventory and transportation. *Manufacturing and Service Operations Management*. **3**(3). 211–229.
- Caliskan-Demirag, O., Y. Chen and Y. Yang. 2012. Ordering policies for periodic-review inventory systems with quantity-dependent fixed costs. *Operations Research*. **60**(4). 785–796.
- Chao, X. and S.X. Zhou. 2009. Optimal Policy for a Multiechelon Inventory System with Batch Ordering and Fixed Replenishment Intervals. *Operations Research*. **57**(2). 377–390.
- Chen, F. 2000. Optimal policies for multi-echelon inventory problems with batch ordering. *Operations Research*. **48**(3). 376–389.
- Federgruen, A. and P. Zipkin. 1986a. An Inventory Model with Limited Production Capacity and Uncertain Demands I. The Discounted-Cost Criterion. *Mathematics of Operations Research*. **11**(2). 193–207.
- Federgruen, A. and P. Zipkin. 1986b. An Inventory Model with Limited Production Capacity and Uncertain Demands II. The Average-Cost Criterion. *Mathematics of Operations Research*. **11**(2). 208–215.
- Gallego, G. and Scheller-Wolf A. 2000. Capacitated inventory problems with fixed order costs: some optimal policy structure. *European Journal Operations Research* **11**(2), 208–215.
- Gallego, G. and L. B. Toktay. 2004. All-or-Nothing Ordering under a Capacity Constraint, *Operations Research* **52**(6). 1001–1002.
- Hammond, J. H. and M. Kelly. 1990. Merloni Elettrodomestici SpA: The Transit Point Experiment. Harvard Business School Case.
- Heyman, D. P. and M. J. Sobel. 1984. *Stochastic Models in Operations Research: Volume II*. Dover Publications, Inc.: Mineola, New York.
- Huh, W. T. and G. Janakiraman. 2012. On Optimal Policies for Inventory Systems with Batch Ordering. *Operations Research*. **60**(4). 797–802.
- Iwaniec, K. 1979. An Inventory Model with Full Load Ordering. *Management Science*. **25**(4). 374–384.
- Lippman, S. 1969a. Optimal Inventory Policy with Multiple Set-up Costs. *Management Science*. **16**(1). 118–138.
- Lippman, S. 1969b. Optimal Inventory Policy with Subadditive Ordering Costs and Stochastic Demands. *SIAM Journal on Applied Mathematics*. **17**(3). 543–559.
- Lippman, M. L 1971. Economic Order Quantities and Multiple Set-up Costs. *Management Science*. **18**(1). 39–47.
- Puterman, M. L. 1994. *Markov Decision Process: Discrete Stochastic Dynamic Programming*. Wiley & Sons, Inc.: Hoboken, New Jersey.
- Scarf, H. 1960. The Optimality of (S, s) Policies in the Dynamic Inventory Problem. Chapter 13 in K. J. Arrow, S. Karlin, and P. Suppes (Eds.), *Mathematical Methods in Social Sciences*, Stanford University Press, Stanford, CA.
- Shaoxiang, C. 2004. The infinite horizon periodic review problem with setup costs and capacity constraints: A partial characterization of the optimal policy. *Operations Research*. **52**(3). 409–421.
- Shaoxiang, C. and M. Lambrecht. 1996. X - Y Band and Modified (s, S) Policy. *Operations Research*. **44**(6). 1013–1019.
- Tanrikulu M. M., A. Sen, O. Alp. 2010. A joint replenishment policy with individual control and constant size orders. *International Journal of Production Research*. **48**(14). 4253–4271.
- Veinott, A. F., Jr. 1965. The Optimal Inventory Policy for Batch Ordering, *Operations Research*, **13**(3), 424–432.

- Veinott, A. F., Jr. 1966. On the optimality of (s, S) inventory policies: New conditions and a new proof. *SIAM Journal of Applied Mathematics*. **14**(5), 1067–1083.
- Zhang, L., S. Çetinkaya. 2011. Stochastic dynamic inventory problem under explicit inbound transportation cost capacity. *Working Paper*. Texas A&M University, College Station, TX.
- Zheng, Y., F. Chen. 1992. Inventory policies with quantized ordering. *Naval Research Logistics*. **39**(3). 285–305.

A Online Appendix

A.1 Proof of Theorem 1

First, we claim the existence of bounds l and u such that there is an optimal policy satisfying $l \leq y_t \leq u$ for each t . It is straightforward to see that this claim implies an existence of an optimal policy satisfying $\lim_{T \rightarrow \infty} \sum_{t=1}^T (y_t - x_t)/T = E[D]$ and every unit procured is eventually sold.

To prove the claim, let $l = -Q$ and $u = y^\circ + Q$ where y° is the largest minimizer of $\mathcal{L}$. Fix an optimal policy and a sample path of y_t 's. We define an alternative policy as a base-stock policy where the order-up-to level is y_t if $y_t \in [l, u]$, $y_t + Q$ if $y_t < l$, and $y_t - Q$ if $y_t > u$. Then, since $\mathcal{L}$ is decreasing in $(-\infty, 0]$ and increasing in (y°, ∞) , it can be shown that the alternative policy performs as well as the original policy and the after-ordering inventory level is within the interval $[l, u]$, completing the proof of the claim.

Now suppose $\lim_{T \rightarrow \infty} \sum_{t=1}^T (y_t - x_t)/T = E[D]$. By adding and subtracting the term $\frac{K}{Q} \sum_{t=1}^T (y_t - x_t)$ to the function $\tilde{C}$, the limit of expression (2) as $T \rightarrow \infty$ can be written as

$$\begin{aligned} \lim_{T \rightarrow \infty} & \left[\frac{1}{T} \cdot \frac{K}{Q} \cdot \sum_{t=1}^T \{y_t - x_t\} + \frac{K}{T} \cdot \sum_{t=1}^T \left\{ \left\lceil \frac{y_t - x_t}{Q} \right\rceil - \frac{y_t - x_t}{Q} \right\} + \frac{1}{T} \sum_{t=1}^T \mathcal{L}(y_t) - \frac{1}{T} \cdot v \cdot x_{T+1} \right] \\ & = \lim_{T \rightarrow \infty} \left[\frac{1}{T} \cdot \frac{K}{Q} \cdot \sum_{t=1}^T \{y_t - x_t\} + \frac{1}{T} \cdot \sum_{t=1}^T c(y_t - x_t) - \frac{1}{T} \cdot v \cdot x_{T+1} \right]. \end{aligned}$$

Since $\lim_{T \rightarrow \infty} \sum_{t=1}^T (y_t - x_t)/T = E[D]$, the first term in the above limit is equal to $E[D] \cdot K/Q$. The last term in the limit is zero since v is a constant. This completes the proof of the first part of the theorem, and the second part directly follows.

A.2 Proof of Theorem 2

Proposition 9. *For the single-period problem,*

- (a) *For any $x \geq y^\circ$, it is optimal not to order, i.e., $y^*(x) = x$.*
- (b) *For any $x < y^\circ$, there exists an optimal policy satisfying $y^*(x) \in \{\theta^{[x]}, \theta^{[x]} + 1, \dots, \theta^{[x]} + Q\}$.*
- (c) *For any pair of $x', x'' < y^\circ$ satisfying $[x'] = [x'']$, there exists an optimal policy satisfying $y^*(x') = y^*(x'')$.*

Proof. Part (a) follows from the fact that $c(\cdot)$ is a nonnegative function and that $\mathcal{L}(y^\circ) \leq \mathcal{L}(x) \leq \mathcal{L}(y)$ for any $y^\circ \leq x \leq y$. More specifically, for any $y \geq x \geq y^\circ$,

$$C(x, y) = c(y - x) + \mathcal{L}(y) \geq 0 + \mathcal{L}(x) = c(0) + \mathcal{L}(x) = C(x, x),$$

where the first and the last equalities follow from the definition of $C(\cdot)$ in (5), the inequality and the second equality follow from the definition of $c(\cdot)$ in (4). Thus, it is optimal not to order if $x \geq y^\circ$.

For part (b), recall from the definition of $\theta^{[x]}$ that $\theta^{[x]} \leq y^\circ < \theta^{[x]} + Q$. We first consider the case of $y < \theta^{[x]}$. In this case,

$$C(x, y) = c(y - x) + \mathcal{L}(y) \geq \mathcal{L}(y) \geq \mathcal{L}(\theta^{[x]}) = C(x, \theta^{[x]}) ,$$

where the first inequality follows from the nonnegativity of $c(\cdot)$, and the second inequality follows from the convexity of $\mathcal{L}(\cdot)$ and the fact that $y < \theta^{[x]} \leq y^\circ$, where y° is the minimizer of $\mathcal{L}(\cdot)$. Thus, selecting $\theta^{[x]}$ as the after-ordering inventory level would be at least as good as selecting y , and it is optimal to select $y(x) \geq \theta^{[x]}$. Now, consider the case of $y > \theta^{[x]} + Q$. A similar argument shows that

$$\begin{aligned} C(x, y) &= c(y - x) + \mathcal{L}(y) \geq \mathcal{L}(y) \\ &\geq \mathcal{L}(\theta^{[x]} + Q) = c((\theta^{[x]} + Q) - x) + \mathcal{L}(\theta^{[x]} + Q) = C(\theta^{[x]} + Q, y) , \end{aligned}$$

where the second equality follows from the fact that the difference between x and $\theta^{[x]}$ is a nonnegative integer multiple of Q . Thus, selecting $\theta^{[x]} + Q$ would also be at least as good as selecting y , and it is optimal to select $y(x) \leq \theta^{[x]} + Q$. Therefore, we complete the proof of part (b).

Finally, for part (c), note that $[x'] = [x'']$ implies $c(y - x') = c(y - x'')$ whenever $x', x'' < y$, and also $\theta^{[x']} = \theta^{[x']}$. Thus, for any $y \in \{\theta^{[x']}, \theta^{[x']} + 1, \dots, \theta^{[x']} + Q\}$, we have

$$C(x', y) = c(y - x') + \mathcal{L}(y) = c(y - x'') + \mathcal{L}(y) = C(x'', y) .$$

Thus, from part (b), we conclude that $y^*(x') = y^*(x'')$. □

We now characterize the optimal solution to (8) given $x \in \mathcal{S}^L$.

Lemma 10. *For the single-period problem, the following policy is optimal for any $x \in \mathcal{S}^L = \{y^\circ - Q + 1, \dots, y^\circ\}$:*

$$y(x) = \begin{cases} y^{[x]} & \text{for } \underline{\theta} \leq y^{[x]} \leq \tilde{\theta} \\ \tilde{\theta} & \text{for } y^{[x]} < \underline{\theta} \text{ or } y^{[x]} > \tilde{\theta} . \end{cases}$$

Proof. For any $x \in \mathcal{S}^L$, we choose the optimal policy by comparing the cost associated with $y^{[x]}$ and $\tilde{\theta}$. Recall, from (6), the expected cost that we want to minimize is $C(x, y) = c(y - x) + \mathcal{L}(y)$, where $c(y - x)$ is given in (7) for $y \in \{x, x + 1, \dots, x + Q\}$. From the earlier discussion, the optimal

order-up-to level $y(x)$ is either a boundary solution $y^{[x]}$ or an interior solution $\tilde{\theta}$. Since $y^{[x]} \in \mathcal{Y}$ is either x or $x + Q$, we have

$$C(x, y^{[x]}) = \mathcal{L}(y^{[x]}) .$$

We compare $C(x, y^{[x]})$ to $C(x, \tilde{\theta})$ by considering the following cases. (These cases are exhaustive since $\tilde{\theta} \geq y^\circ \in \mathcal{Y}$ follows from the definition of $\tilde{\theta}$ in (9).)

- Case $\tilde{\theta} > \max \mathcal{Y}$. Since $y^{[x]} \in \mathcal{Y}$ and $\tilde{\theta} \notin \mathcal{Y}$, the definition of $\mathcal{Y}$ shows $\mathcal{L}(y^{[x]}) \leq \mathcal{L}(\tilde{\theta})$, implying $C(x, y^{[x]}) \leq C(x, \tilde{\theta})$. Thus, it is better to order up to $y^{[x]}$ than up to $\tilde{\theta}$, implying $y(x) = y^{[x]}$ for any value of $y^{[x]}$. This corresponds to the first case of $y(x)$ in the statement of the lemma since $-\infty = \underline{\theta} \leq y^{[x]} \leq \max \mathcal{Y} < \tilde{\theta}$. (See the definition of $\underline{\theta}$ in (10).)
- Case $\tilde{\theta} \in \mathcal{Y}$ and $y^{[x]} > \tilde{\theta}$. First, we claim $y^{[x]} > \tilde{\theta} \geq y^\circ > x$. To prove this claim, the first inequality follows from the assumption of this case, and the second inequality follows from the definition of $\tilde{\theta}$. For the last inequality, $x \in \mathcal{S}^L \in \{y^\circ - Q + 1, \dots, y^\circ\}$ implies $x \leq y^\circ$. Thus, we have $y^{[x]} > \tilde{\theta} \geq y^\circ \geq x$. If the last inequality were to hold with equality (i.e., $y^\circ = x$), then it would have implied $x = y^{[x]}$ since $y^\circ \in \mathcal{Y}$. Thus, we would have obtained $y^{[x]} > \tilde{\theta} \geq y^\circ = x = y^{[x]}$, which is a contradiction. This completes the proof of the claim.

By the above claim, we have $y^{[x]} > x$, which implies $y^{[x]} = x + Q$ since $x \in \mathcal{S}^L$. From (9),

$$\begin{aligned} \mathcal{L}(\tilde{\theta}) - \frac{K}{Q}\tilde{\theta} &\leq \mathcal{L}(y^{[x]}) - \frac{K}{Q}y^{[x]} \\ \mathcal{L}(\tilde{\theta}) + \frac{K}{Q}(y^{[x]} - \tilde{\theta}) &\leq \mathcal{L}(y^{[x]}) \end{aligned}$$

where the last inequality implies $C(x, \tilde{\theta}) \leq C(x, y^{[x]})$. Hence, we conclude that $y(x) = \tilde{\theta}$ is optimal.

- Case $\tilde{\theta} \in \mathcal{Y}$ and $\underline{\theta} \leq y^{[x]} \leq \tilde{\theta}$. If $y^{[x]} = \tilde{\theta}$, then the two cases in the statement of the lemma collapses into one case. Thus, we proceed by assuming otherwise, i.e., $y^{[x]} < \tilde{\theta}$. We consider two sub-cases separately. If $y^{[x]} \geq y^\circ$, then by the convexity of $\mathcal{L}(\cdot)$ and $y^\circ \leq y^{[x]} < \tilde{\theta}$ imply

$$C(x, y^{[x]}) = \mathcal{L}(y^{[x]}) \leq \mathcal{L}(\tilde{\theta}) \leq C(x, \tilde{\theta})$$

If $y^{[x]} < y^\circ$, then $\underline{\theta} \leq y^{[x]} < y^\circ$. We have

$$\mathcal{L}(y^{[x]}) \leq \mathcal{L}(\underline{\theta}) \leq \mathcal{L}(\tilde{\theta}) + \frac{K}{Q} \cdot (\underline{\theta} + Q - \tilde{\theta}) \leq \mathcal{L}(\tilde{\theta}) + \frac{K}{Q} \cdot (y^{[x]} + Q - \tilde{\theta}) ,$$

where the first inequality follows since $\mathcal{L}(\cdot)$ is decreasing between $\underline{\theta}$ and y° , the second inequality follows from the definition of $\underline{\theta}$ in (10), and the last inequality follows from the assumption

of this case. Note that the leftmost expression of the above inequality is $C(x, y^{[x]})$ and the rightmost expression is $C(x, \tilde{\theta})$. Thus, it is optimal to order up to $y^{[x]}$, i.e., $y(x) = y^{[x]}$, which is the first case in the statement of the lemma.

- Case $y^{[x]} < \underline{\theta}$. From the definition of $\underline{\theta}$ in (10), it follows

$$\mathcal{L}(y^{[x]}) > \mathcal{L}(\tilde{\theta}) + \frac{K}{Q}(y^{[x]} + Q - \tilde{\theta}) ,$$

which implies that it is optimal to order up to $\tilde{\theta}$, i.e., $y(x) = \tilde{\theta}$.

Combining the above cases, we complete the proof. $\square$

Lemma 10 considers the case of $x \in \mathcal{S}^L$, where the optimal ordering quantity is at most the base quantity Q , requiring at most one set-up. (This follows from the definition $\mathcal{S}^L$ and Proposition 9(b).) The optimal policy on $\mathcal{S}^L$ can be characterized by two thresholds $\underline{\theta}$ and $\tilde{\theta}$. If $y^{[x]}$ falls between $\underline{\theta}$ and $\tilde{\theta}$, it is optimal to order up to $y^{[x]}$, in which case, the order quantity is either zero or the base quantity Q . Otherwise, it is optimal to order up to $\tilde{\theta}$, in which case the order quantity is less than Q . Lemma 10 shows that the optimal policy has a nice threshold structure characterized by two parameters.

Now, we are ready to prove Theorem 2. If $x > y^\circ$, then $\mathcal{L}(x) \leq \mathcal{L}(y)$ for any $y \geq x$. Thus, it is optimal not to order. Note that Lemma 10 has shown the required result if $x \in \mathcal{S}^L$. For $x \leq y^\circ - Q$, the optimal decision at x is the same as the optimal decision at $x + Q$ (by Proposition 9(c)). Thus, the result of Lemma 10 easily generalizes to any $x \leq y^\circ - Q$, and we complete the proof.

A.3 Proof of Corollary 3

For this proof, we use the notation $\underline{\theta}_K$ and $\tilde{\theta}_K$ to denote their dependency on K explicitly. From its definition in (9), $\tilde{\theta}_K$ can easily be shown to be non-decreasing in K .

Now, we prove that $\underline{\theta}_K$ is non-increasing in K . Let K_1 and K_2 be real numbers such that $K_1 \leq K_2$, and we want to show $\underline{\theta}_{K_1} \geq \underline{\theta}_{K_2}$. Since we have shown that $\tilde{\theta}_K$ is non-decreasing in K , it follows that $\tilde{\theta}_{K_1} \leq \tilde{\theta}_{K_2}$. If $\tilde{\theta}_{K_2} > \max \mathcal{Y}$, then $\underline{\theta}_{K_2} = -\infty$ follows from the definition of $\underline{\theta}_{K_2}$ in (10), and thus the required result holds. We proceed by assuming $\tilde{\theta}_{K_1} \leq \tilde{\theta}_{K_2} \leq \max \mathcal{Y}$. Since y^* is a lower bound for any $\tilde{\theta}$ (from the definition in 9), it follows that $\tilde{\theta}_{K_1}, \tilde{\theta}_{K_2} \in \mathcal{Y}$. For any $K \geq 0$, define

$$l_K(z) = \mathcal{L}(\tilde{\theta}_K) + \frac{K}{Q} \cdot (z + Q - \tilde{\theta}_K) ,$$

which is a linear function of z with the slope K/Q . We claim that

$$l_{K_1}(z) \leq l_{K_2}(z) \quad \text{for any } z \in \{\min \mathcal{Y}, \tilde{\theta}_{K_2}\}.$$

From the definition of $\underline{\theta}_{K_1}$ and $\underline{\theta}_{K_2}$ given in (10), the above claim implies that $\underline{\theta}_{K_1} \geq \underline{\theta}_{K_2}$, as required.

To prove the claim, recall $\tilde{\theta}_{K_1} \leq \tilde{\theta}_{K_2}$, and note that

$$\begin{aligned} l_{K_1}(\tilde{\theta}_{K_2}) - K_1 &= \mathcal{L}(\tilde{\theta}_{K_1}) + \frac{K_1}{Q} \cdot (\tilde{\theta}_{K_2} + Q - \tilde{\theta}_{K_1}) - K_1 = \mathcal{L}(\tilde{\theta}_{K_1}) + \frac{K_1}{Q} \cdot (\tilde{\theta}_{K_2} - \tilde{\theta}_{K_1}) \\ &\leq \mathcal{L}(\tilde{\theta}_{K_2}) = l_{K_2}(\tilde{\theta}_{K_2}) - K_2, \end{aligned} \tag{11}$$

where the inequality follows from constructing a supporting hyperplane of the convex function $\mathcal{L}$ at $\tilde{\theta}_{K_1}$ using the fact that K_1/Q is a subgradient of $\mathcal{L}$ at this point (see the definition of $\tilde{\theta}_{K_1}$ in (9)). Therefore, for any z satisfying $\min \mathcal{Y} \leq z \leq \tilde{\theta}_{K_2} \leq \max \mathcal{Y}$,

$$\begin{aligned} l_{K_2}(z) - l_{K_1}(z) &= \left[l_{K_2}(\tilde{\theta}_{K_2}) - \frac{K_2}{Q} \cdot (\tilde{\theta}_{K_2} - z) \right] - \left[l_{K_1}(\tilde{\theta}_{K_2}) - \frac{K_1}{Q} \cdot (\tilde{\theta}_{K_2} - z) \right] \\ &= \left[l_{K_2}(\tilde{\theta}_{K_2}) - l_{K_1}(\tilde{\theta}_{K_2}) \right] - (K_2 - K_1) \cdot \frac{\tilde{\theta}_{K_2} - z}{Q} \\ &\geq [K_2 - K_1] - (K_2 - K_1) \cdot \frac{\tilde{\theta}_{K_2} - z}{Q} \\ &\geq 0, \end{aligned}$$

where the first inequality follows from (11) and the second inequality follows the fact that both $\tilde{\theta}_{K_2}$ and z belong to the interval $[\min \mathcal{Y}, \max \mathcal{Y}]$ where $|\mathcal{Y}| = Q$. Thus, we complete the proof of the claim.

Now, from the definitions of $\tilde{\theta}_K$ and $\underline{\theta}_K$ in (9) and (10), it is easy to see that both of these quantities converge to y° as $K \downarrow 0$. Thus, for sufficiently small K , we obtain that $\tilde{\theta}$ becomes y° and the set $\{\underline{\theta}, \underline{\theta} + 1, \dots, \tilde{\theta}\}$ becomes a singleton set $\{y^\circ\}$; then, the optimal policy given in Theorem 2 becomes

$$y(x) = \begin{cases} x & \text{if } x > y^\circ \\ y^\circ & \text{if } x \leq y^\circ, \end{cases}$$

which is the order-up-to- y° policy. Now, for sufficiently large K , it follows that $\tilde{\theta}_K$ exceeds $\max \mathcal{Y}$ (see (9)), and in case, we have $\underline{\theta}_K = -\infty$ (see (10)). Then, the optimal policy in Theorem 2 becomes

$$y(x) = \begin{cases} x & \text{if } x > y^\circ \\ y^{[x]} & \text{if } x \leq y^\circ, \end{cases}$$

which is the same as the order-up-to- $y^{[x]}$ policy.

A.4 Proof of Lemma 4

Suppose that $y_t(x_t)$ is any given policy for starting inventory level x_t in period t . We define a new policy $\hat{y}_t(x_t)$ such that $\hat{y}_t(x_t) \in \mathcal{Y}$ and $[\hat{y}_t(x_t)] = [y_t(x_t)]$, i.e., $\hat{y}_t(x_t)$ and $y_t(x_t)$ differ by a nonnegative integer multiple of Q . For any sample path of demand realization, let $\{y_t\}$ and $\{\hat{y}_t\}$ denote the after-ordering inventory levels under the original policy and under the new policy, respectively. Then, for each $t \geq 1$, we obtain $\mathcal{L}(\hat{y}_t) \leq \mathcal{L}(y_t)$ since $\hat{y}_t \in \mathcal{Y}$. Also, let x_t and $\hat{x}_t$ denote the before-ordering inventory levels under the original policy and under the new policy, respectively. Then, it is easy to show by induction that $[\hat{y}_t] = [y_t]$ and $[\hat{x}_t] = [x_t]$ hold for each t , and thus it follows that $c(\hat{y}_t - \hat{x}_t) = c(y_t - x_t)$ holds for each t . Therefore, for any $T \geq 1$,

$$\sum_{t=1}^T C(\hat{x}_t, \hat{y}_t) = \sum_{t=1}^T c(\hat{y}_t - \hat{x}_t) + \mathcal{L}(\hat{y}_t) \leq \sum_{t=1}^T c(y_t - x_t) + \mathcal{L}(y_t) = \sum_{t=1}^T C(x_t, y_t),$$

and we conclude that the new policy is optimal.

A.5 Proof of Theorem 5

From Lemma 4 and the construction of $\mathcal{M}$, solving the reduced problem is equivalent to solving $\mathcal{M}$. Furthermore, $\hat{C}$ is a lower bound for the optimal cost C^* since the relaxed problem does not have the $y_t \geq x_t$ constraint in each period. The above equality follows from the relationship between the original optimal cost $\tilde{C}^*$ and the equivalent alternate cost C^* (see Section 2.2).

A.6 Proof of Proposition 6

Since $p^{[0]} = p^{[1]} = \dots = p^{[Q-1]}$, the probability distribution of the equivalent class to which the ending inventory position belongs is independent of the ordering decision. Thus, the optimal action of $\mathcal{M}$ in each period is to minimize the cost of the current period. For this effect, the optimal inventory policy is first to order the quantity specified by Theorem 2.

A.7 Proof of Theorem 7

We first prove that there exists an optimal policy, not necessarily stationary, for the original infinite-horizon problem satisfying (a), (b) and (c). Suppose that $\{(x'_t, y'_t) | t = 1, 2, \dots\}$ denotes a sequence of before-ordering and after-ordering inventory levels according to any given policy, which we refer to as the ‘incumbent policy’. We construct a modified system with the following policy

$$y''_t = \begin{cases} x''_t & \text{if } x''_t \geq y^\circ \\ \theta^{[x''_t]} & \text{if } x''_t < y^\circ \text{ and } y'_t \leq \theta^{[x''_t]} \\ \theta^{[x''_t]} & \text{if } x''_t < y^\circ \text{ and } y'_t > \max \mathcal{Y} \geq \theta^{[x''_t]} \geq \min \mathcal{Y} \\ \theta^{[x''_t]} + Q & \text{if } x''_t < y^\circ \text{ and } y'_t > \max \mathcal{Y} > \min \mathcal{Y} > \theta^{[x''_t]} \\ \min\{y'_t, \theta^{[x''_t]} + Q\} & \text{if } x''_t < y^\circ \text{ and } \max \mathcal{Y} \geq y'_t > \theta^{[x''_t]}. \end{cases}$$

Note that this policy orders either nothing or (an integer multiple of) full batch, or mimics the incumbent policy as for the order-up-to level. It is straightforward to verify that the above equation defines a feasible policy for the modified system, and that it satisfies both properties stated in (a) and (b).

In this modified policy, one does not order if the current inventory is at least y° . Suppose $x_t'' < y^\circ$. If $y_t' \leq \theta^{[x_t'']}$ or $y_t' > \max \mathcal{Y}$ (i.e. branches 2-4 in the above equation), then the ordering quantity of the modified system is a nonnegative integer multiple of the base quantity Q , and the after-ordering inventory level of the modified system is $\theta^{[x_t'']}$ or $\theta^{[x_t'']} + Q$. Finally, if $\theta^{[x_t'']} < y_t' \leq \max \mathcal{Y}$, the modified system orders up to the smaller of y_t' and $\theta^{[x_t'']} + Q$. From our construction of the modified policy, it can be verified that we maintain the following invariant in each period: (i) $y_t' \leq y_t'' \leq y^\circ$, (ii) $y^\circ \leq y_t'' \leq y_t'$, or (iii) $\min \mathcal{Y} \leq y_t'' \leq \max \mathcal{Y} < y_t'$. Furthermore, we can show $\mathcal{L}(y_t'') \leq \mathcal{L}(y_t')$ from the above invariant. If (i) or (ii) holds, it follows that $\mathcal{L}(y^\circ) \leq \mathcal{L}(y_t'') \leq \mathcal{L}(y_t')$ holds by the convexity of $\mathcal{L}$. If (iii) holds, we have $y_t'' \in \mathcal{Y}$ and $y_t' \notin \mathcal{Y}$, which ensures $\mathcal{L}(y_t'') \leq \mathcal{L}(y_t')$ by the definition of $\mathcal{Y}$.

Now, we compare the $c(\cdot)$ cost of the two systems. Note that the modified policy always orders integer multiples of the base quantity Q unless this policy “couples” with the incumbent policy, i.e., $y_t'' = y_t'$. In other words, the positive $c(\cdot)$ in the modified system occurs only if $y_t'' = y_t'$. Given two systems $\{(x_t', y_t') | t = 1, 2, \dots\}$ and $\{(x_t'', y_t'') | t = 1, 2, \dots\}$, let τ_i be the period in which the order-up-to levels of both systems are the same for the i^{th} time, i.e., let $\tau_0 = 0$, and $\tau_i = \min\{t > \tau_{i-1} \mid y_t'' = y_t'\}$ for any $i \geq 1$. Then, for any i , $c(y_t'' - x_t'') = 0$ for any $t \in \{\tau_{i-1} + 1, \dots, \tau_i - 1\}$ by construction. Furthermore,

$$\sum_{t=\tau_{i-1}+1}^{\tau_i} c(y_t'' - x_t'') = c(y_{\tau_i}'' - x_{\tau_i}'') = c\left(\sum_{t=\tau_{i-1}+1}^{\tau_i} (y_t' - x_t')\right) \leq \sum_{t=\tau_{i-1}+1}^{\tau_i} c(y_t' - x_t')$$

The first equality follows since the modified policy does not order any partial batch between periods $\tau_{i-1} + 1$ and $\tau_i - 1$. The second equality follows as the total under-utilization in both systems are equal to each other in modulus Q and the order-up-to levels in both systems are equal to each other in periods τ_{i-1} and τ_i . Finally, the last inequality follows from the definition of c .

It follows that $\sum_{t=1}^T C(x_t'', y_t'') \leq \sum_{t=1}^T C(x_t', y_t')$ holds for any $T \geq 1$. Property (c) follows from (b) and the fact that, for any pair of $x', x'' < y^\circ$ satisfying $[x'] = [x'']$, $C([x'], y) = C([x''], y)$ holds for any $y \geq \max\{x', x''\}$. Hence, the modified policy also satisfies (c). Consequently, the modified policy –which satisfies (a), (b), and (c)– generates a total cost which is less than or equal to that of any other policy and therefore it is optimal.

We now proceed to show the existence of an optimal *stationary* policy. From properties (a)-(c) for an optimal policy, it suffices to consider a finite set of states and a finite set of actions. Thus, a stationary optimal policy exists by Theorem 4-14 of Heyman and Sobel (1984).

A.8 Proof of Theorem 8

For part (a), we assume that $x_1 \leq \min \mathcal{Y}$. (If this assumption does not hold, then we can easily modify the following argument by ignoring the initial transition periods until the first positive quantity is ordered by the RMB policy. This is acceptable since we consider the average cost criterion.) Then, since demand D_t is at least Q units in each period, it follows from an induction argument that, under the RMB policy, both $x_t \leq \min \mathcal{Y}$ and $y_t \in \mathcal{Y}$ hold for each period $t \geq 1$. Thus, the after ordering inventory level of the RMB policy is

$$y_t(x_t) = \max\{\hat{y}([x_t]), x_t\} = \hat{y}([x_t]) .$$

Therefore, the long-run average cost under the RMB policy is the same as the lower bound $\hat{C}$ given in Theorem 5. We conclude that the RMB policy is optimal.

Part (b) follows directly from part (a) and Proposition 6.