Inventory control with partial batch ordering

Citation for published version (APA):
Alp, O., Huh, W. T., & Tan, T. (2009). *Inventory control with partial batch ordering*. (BETA publicatie : working papers; Vol. 283). Technische Universiteit Eindhoven.

Document status and date:
Published: 01/01/2009

Document Version:
Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

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.

Download date: 21. Sep. 2023

Inventory Control with Partial Batch Ordering

Osman Alp[*] Woonghee Tim Huh[†] Tarkan Tan[‡]

July 4, 2009

Abstract

In an infinite-horizon, periodic-review, single-item production/inventory system with random demand and back-ordering, we study the feature of batch ordering, where a separate fixed cost is associated for each batch ordered. Contrary to majority of the literature on this topic, we do not restrict the order quantities to be integer multiples of the 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, and we discuss several properties of the optimal solution. Based on the analysis of a single-period problem and a multi-period lower-bound problem, we study two heuristic policies for the original partial batch ordering problem, both of which perform very well computationally for a wide range of problem parameters. Finally, we compare the performance of the optimal policy to the performance of the best full-batch-size ordering policy to quantify the value of partial ordering flexibility.

1. Introduction and Related Literature

In most production environments, multiple units of items are procured, processed, manufactured or shipped together in batches. The batch ordering nature often arises from capacitated facilities, such as industrial ovens, containers, trucks and ships, that are utilized to produce or supply the item. In other cases, it is due to economic or technological requirements at the powerful supplier or the manufacturer, manifested through case pack sizes, minimum order quantities, etc.

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

[†]huh@ieor.columbia.edu. Department of Industrial Engineering and Operations Research, Columbia University, New York, NY 10027, USA. Research partially supported by NSF grant DMS-0732169.

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

In a batch ordering environment, system costs are heavily influenced by the number of batches used for ordering. A separate fixed cost is charged for ordering (or producing, shipping, etc) each batch, as well as a variable cost for each unit ordered. While this kind of a cost structure is easy to describe and frequently observed in practice, it does not lend itself to straightforward analytical tractability. As a result, the treatment of batch ordering in the academic literature has been somewhat simplified. Typical approaches for modeling the procurement cost are either (i) to ignore the batch cost altogether by considering only the linear cost, (ii) to include a single fixed cost for ordering, i.e., to assume an infinite batch size, (iii) to cap the total order quantity, or (iv) to impose additional requirement such as full batch sizes (also referred to as "full truck loads" or "full containers"). The first three variants have been analyzed well. When the ordering cost is linear, the base-stock policy is optimal. When the fixed cost of ordering is associated with each period, the optimal policy is an (s, S) policy (Scarf, 1960; Veinott, 1966). When the capacity is imposed on the ordering quantity in each period without explicit modeling of the fixed cost, the modified base-stock policy is optimal (Federgruen and Zipkin, 2000).

In the literature, the phrase "batch ordering" typically refers to the problem with the full batch size ordering restriction. This problem is introduced by Veinott (1965), who assumes stochastic demand and requires that orders must be multiples of the full batch size. In a single-stage setting, he 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.[1] 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). 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.

In an environment where the fixed cost of ordering a batch is dominant, ordering full batches is reasonable and commonly used. However, when this fixed cost is moderate, rounding up or down the order quantity to a multiple of the full batch size may not provide an optimal trade-off between service level, inventory holding cost and the batch ordering

[1]Her conditions are mildly reminiscent of the K-convexity of Scarf (1960) that is used to prove the optimality of (s, S) policies with the single fixed cost for ordering (the case of (ii) above).

cost. By allowing the flexibility of a partial batch, the overall cost can be reduced, and we expect that such a saving to be high when the full batch size is large. In this paper, we consider an infinite-horizon, periodic-review, single-item inventory system with random demand and batch ordering, where a separate fixed cost is associated for each batch ordered. We do not restrict the order quantities to be integer multiples of the batch size and allow the possibility of partial batches, in which case the fixed cost for ordering the batch is still fully charged.

The partial batch ordering flexibility, a key feature of this paper, has received relatively little attention in the literature. To our knowledge, there is no paper that analyzes the optimal policies with partial batch ordering and stochastic demand in a single-stage or serial system. While a single-stage problem with partial batch ordering is already a difficult problem that has not been analyzed, we are aware of only two papers, Cachon (2001) and Tanrikulu et al. (2009), that contain this feature as a part of a more complex multi-item joint-replenishment context where heuristic methods are proposed and evaluated.

Our main contributions can be summarized as follows. (1) We build a model that particularly takes the batch ordering cost structure into account and allows partial batch ordering. Even though a multi-item version of this problem has been studied by Cachon (2001) and Tanrikulu et al. (2009), we focus on a simpler yet still challenging problem in order to derive analytic results and develop intuition. (2) To facilitate our analysis, we introduce a novel alternative cost accounting scheme. Here, instead of charging the per-batch cost to every batch ordered, we impose an appropriate penalty cost for any batch that is less than full. While this scheme is equivalent to the original cost structure, it allows develop several properties of the optimal solution. (3) We characterize the optimal solution for the single-period problem. 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 our problem. (4) We examine two policies that can be used to solve the problem in a heuristic manner. These policies are designed for the partial-ordering case, 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. We compare them to the case of full-batch-ordering policy to quantify the value of partial-ordering flexibility. (5) Finally, we build managerial insights for the inventory systems with batch ordering costs.

We mention several 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). Both Lippman (1969) and Alp et al. (2003) consider a general version of our problem under *deterministic* demand settings and show some optimality properties.

The rest of the paper is organized as follows. We present our model and an alternative cost accounting scheme in Section 2. We analyze the problem, show optimality results, and propose heuristic policies in Section 3. The performance of these policies is reported and managerial insights are presented in Section 4. We conclude the paper in Section 5.

2. Model

2.1 Description

In this section, we describe the details of our model, and then present a dynamic programming formulation. 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 of batch ordering, all of which are exogenously determined and non-negative. We ignore unit ordering costs without loss of generality in the long run. We assume full availability of the ordered quantities, and that the lead times can be neglected. The batch size is assumed to be fixed. In contrast to most existing models on batch ordering, we do not restrict the order quantities to be integer multiples of the batch size, and we allow the possibility of fractional number of batches. However, the batch ordering cost is a function of the number of batches, regardless of whether all the batches are full or not. Note that this function is neither convex nor concave in the number of units ordered.

Let $t \in \{1, 2, \ldots\}$ 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 based on the beginning inventory level x_t, and incurs the ordering cost of $\tilde{c}(q_t)$ given by

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

where $K \geq 0$ represents the ordering cost per batch, $Q > 0$ denotes the fixed batch size, 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 required to order q units). Note that $\tilde{c}(q)$ is a right-continuous step function where the 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 $(D_1, D_2, \ldots)$ are independent and identically distributed, and we denote the common distribution by D. For simplicity of exposition, we assume that demands are discrete with integer supports, but our model and analysis can easily be generalized to the case of the continuous demand. An appropriate linear overage or underage cost is charged, where the per-unit per-period overage and underage costs are denoted by h and b, respectively. (4) The excess demand, if any, is backlogged, and therefore the beginning inventory level in the next period is given by $x_{t+1} = y_t - D_t$.

The expected overage and underage cost in each period t depends only on y_t, and can be written as $\mathcal{L}(y_t)$, where

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

Note that $\mathcal{L}$ is a convex function. The expected cost incurred in period t can be written as

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

The objective is to minimize the long-run average cost, i.e., $\limsup_{T\to\infty} \sum_{t=1}^{T} \tilde{C}_t(x_t, y_t)/T$. Let $\tilde{C}^*$ denote the optimal long-run average cost.

2.2 Equivalent Alternate Cost Formulation

We introduce an alternative cost accounting scheme that would be easier to work with in our analysis. Since every unit procured is eventually sold, the long-run average number of ordered units is $E[D]$. Thus, the long-run average ordering cost would have been $E[D] \cdot K/Q$ if each order is an integer multiple of the batch quantity Q. However, in general, the average ordering cost would be higher, due to the fact that the batch cost K cannot be divided among Q units if the order quantity is less than the batch size Q. Thus, based on this observation, we take $E[D] \cdot K/Q$ as a baseline of the batch ordering cost, and we introduce an alternative but equivalent cost structure. Define

$$c(q) \quad = \quad K \cdot (\lceil q/Q \rceil - q/Q) \ . \tag{2}$$

If q is a multiple of Q, then the above expression is zero. Otherwise, we interpret $c(q)$ as the cost of ordering a less-than-full batch. Let

$$C_t(x_t, y_t) \quad = \quad c(y_t - x_t) + \mathcal{L}(y_t) \ . \tag{3}$$

Then, minimizing the long-run average cost in terms of C_t is equivalent to minimizing the long-run average of the original cost $\tilde{C}_t$ (up to an additive constant), i.e.,

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

Therefore, for the remainder of this paper, we adopt the objective of minimizing the right-hand-side expression of the above equation. We let C^* denote the optimal long-run average cost in terms of the C_t. Clearly, $C^* = \tilde{C}^* - E[D] \cdot K/Q$.

One of the common approaches to specify and solve inventory problems is to use dynamic programming. The finite-horizon T-period problem can be formulated as follows:

$$\begin{aligned} f_t(x_t) &= \min_{y_t \geq x_t} c(y_t - x_t) + \mathcal{L}(y_t) + E_{D_t}\left[f_{t+1}(y_t - D_t)\right] \qquad \text{for } t = 1, \ldots, T, \text{ and} \\ f_{T+1}(x_{T+1}) &= v \cdot x_{T+1} \ , \end{aligned}$$

where v is the per-unit salvage value. Under mild technical conditions, the average of the T-period cost converges to the long-run average cost as $T \to \infty$. Dynamic programming is a commonly-used tool in the inventory theory literature, and its analysis typically takes advantage of special structures that are preserved through dynamic programming recursions, such as convexity or its generalizations (see, for example, Chen, 2000; and Gallego and Toktay, 2004). However, convexity does not hold in our problem since the ordering cost function $c(\cdot)$ is not convex, and thus the optimal solution for the dynamic program is likely to be quite complex in general.

3. Analysis

Since the optimal policy for the multiple-period dynamic programming formulation is difficult to analyze, we first focus on the single-period problem, for which we show that the optimality of a threshold-type policy (Section 3.1). Then, in a multiple-period setting, we consider a variant of the original problem, and show how we compute the optimal solution for this problem (Section 3.2). This problem provides a lower bound on the optimal cost. Inspired by the optimal policies for these related problems, we propose heuristic policies for the original problem and present conditions under which these policies are indeed optimal (Section 3.3). The performance of these policies is reported later in Section 4.

3.1 The Single-Period Problem and the Myopic Policy

We consider the ordering problem in the last period $t = T$. For the sake of simplicity, we assume no salvage value, i.e., $v = 0$. (Otherwise, the salvage value can be incorporated by modifying the $\mathcal{L}$ function.) Then, the last-period problem, which we refer to as the single-period problem, can be written as:

$$\min_{y_T \geq x_T} \; c(y_T - x_T) + \mathcal{L}(y_T) \; , \tag{4}$$

where x_T is the starting inventory level in the last period. We let $y_T(x_T)$ denote any inventory policy for the above problem, $y_T^*(x_T)$ being the optimal one; this notation makes it explicit that the action y_T depends on the current state x_T.

Since the batch size is Q, the modular arithmetic based on Q plays an important role in our analysis. For any integer z, we define

$$[z] \;\; = \;\; z \mod Q$$

such that $[z] \in \{0, 1, \ldots, Q-1\}$ i.e., $z = [z] + kQ$ for some integer k. Mathematically, $[\cdot]$ defines a set of equivalence classes.

Recall that $\mathcal{L}$ given in (1) is a convex function, satisfying $\mathcal{L}(y) \to \infty$ as $y \to \infty$ or $y \to -\infty$. Let θ^* be the minimizer of $\mathcal{L}$. (If $\mathcal{L}$ has multiple minimizers, then fix θ^* at the largest minimizer, to avoid ambiguity in the definition.) We note that θ^* 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 θ^*, i.e.,

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

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

$$\begin{aligned} \mathcal{S}^L \;\; &= \;\; \{\theta^* - Q + 1, \; \theta^* - Q + 2, \; \ldots, \; \theta^*\} \qquad \text{and} \\ \mathcal{S}^R \;\; &= \;\; \{\theta^*, \; \theta^* + 1, \; \ldots, \; \theta^* + Q - 1\} \; . \end{aligned}$$

Note that the set $\{\theta^{[z]} \mid 0 \le z < Q\}$ 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

$$\mathcal{L}(y') \ \le \ \mathcal{L}(y'') \qquad \text{for any} \quad y' \in \mathcal{Y} \quad \text{and} \quad y'' \notin \mathcal{Y} \ .$$

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 $[y] = [z]$.

We establish basic properties of the optimal policy for the single-period problem.

Proposition 1. *For the single-period problem,*

(a) For any $x_T \ge \theta^*$, it is optimal not to order, i.e., $y_T^*(x_T) = x_T$.

(b) For any $x_T < \theta^*$, $y_T^*(x_T) \in \{\theta^{[x_T]},\ \theta^{[x_T]}+1,\ \ldots,\ \theta^{[x_T]}+Q\}$.

(c) For any pair of $x'_T, x''_T < \theta^*$ satisfying $[x'_T] = [x''_T]$, $y_T^*(x'_T) = y_T^*(x''_T)$.

Proof. Part (a) follows from the fact that $c(\cdot)$ is a nonnegative function and that $\mathcal{L}(\theta^*) \le \mathcal{L}(x_T) \le \mathcal{L}(y)$ for any $\theta^* \le x_T \le y$.

For part (b), $C(x_T, y_T)$ is bounded below as follows:

$$c(y_T(x_T) - x_T) + \mathcal{L}(y_T(x_T)) \ \ge \ \mathcal{L}(y_T(x_T)) \quad \ge \ \mathcal{L}(\theta^{[x_T]}) \ ,$$

where the first inequality follows from the nonnegativity of c, and the second inequality follows from the convexity of $\mathcal{L}(\cdot)$ and the fact that $y_T(x_T) < \theta^{[x_T]} \le \theta^*$. Thus, selecting $\theta^{[x_T]}$ as the after-ordering inventory level would be at least as good as selecting $y_T(x_T)$. Similarly, if $y_T(x_T) > \theta^{[x_T]} + Q$, then a similar argument shows

$$c(y_T(x_T) - x_T) + \mathcal{L}(y_T(x_T)) \ \ge \ \mathcal{L}(y_T(x_T)) \quad \ge \ \mathcal{L}(\theta^{[x_T]} + Q) \ .$$

Thus, selecting $\theta^{[x_T]} + Q$ would also be at least as good as selecting $y_T(x_T)$, completing the proof of part (b).

Finally, part (c) follows from the fact that $c(y - x'_T) = c(y - x''_T)$ whenever $x'_T, x''_T < y$. $\square$

Proposition 1 gives a partial characterization of the optimal policy – that it is optimal not to order if the beginning inventory level exceeds θ^*, and otherwise the after-ordering inventory level must belong to a subset of $\mathcal{S}^L \cup \mathcal{S}^R = \{\theta^* - Q + 1,\ \theta^* - Q + 2,\ \ldots,\ \theta^* + Q - 1\}$. Furthermore, parts (a) and (c) imply that it suffices to specify the optimal policy within a subset of the state space for the dynamic programming formulation, namely the set $\mathcal{S}^L = \{\theta^* - Q + 1, \ldots, \theta^*\}$. (If $x_T \geq \theta^*$, then it is optimal not to order; if $x_T \leq \theta^* - Q$, then the order-up-to level is the same as that of x'_T where $x'_T \in \mathcal{S}^L$ and $[x_T] = [x'_T]$.) Furthermore, for any $x_T \in \mathcal{S}^L$, part (b) shows that it is adequate to consider a restricted action space given by $y_T(x_T) \leq Q + x_T$, i.e., ordering at most one batch.

Based on the above discussion, the single-period problem of (4) can be written as follows: For any $x_T \in \mathcal{S}^L$,

$$\min_{x_T \leq y_T \leq x_T + Q} \mathcal{L}(y_T) + \psi(y_T) \tag{5}$$

where

$$\psi(y_T) = \begin{cases} (K/Q) \cdot (x_T + Q - y_T) & \text{if } y_T > x_T, \\ 0 & \text{if } y_T = x_T. \end{cases}$$

While $\mathcal{L}(y_T)$ is convex, $\psi(y_T)$ is concave in y_T. Thus, the objective function in (5) is neither convex or concave. Despite this property, it turns out that the above problem can easily be solved. Note that expression (5) depends on y_T only through $\mathcal{L}(y_T)$ and $(K/Q) \cdot y_T$. It is useful to define

$$\tilde{\theta} \;=\; \arg\min_{y \geq 0} \quad \mathcal{L}(y) - (K/Q) \cdot y\ . \tag{6}$$

This minimization problem given in (6) is a convex function minimization problem. We note that while $\tilde{\theta}$ is bounded below by θ^* (i.e., $\theta^* \leq \tilde{\theta}$), it may or may not belong to the set $\mathcal{Y}$ or $\mathcal{S}^R$.

If the optimal solution to (5) is at a boundary x_T or $x_T + Q$ then it belongs to $\mathcal{Y}$ and is given by $y^{[x_T]}$ since x_T belongs to $\mathcal{S}^L$ and $\mathcal{Y}$ is the set of consecutive integers corresponding to the Q smallest values of $\mathcal{L}$, the optimal boundary solution belongs to $\mathcal{Y}$ and is given by $y^{[x_T]}$.

Therefore, the optimal solution to (5) for given $x_T \in \mathcal{S}^L$ is either the boundary solution $y^{[x_T]}$ or an interior solution $\tilde{\theta}$. Proposition 1 below shows which of these two solutions is indeed optimal, and that this decision depends on whether the value of $y^{[x_T]} \in \mathcal{Y}$ exceeds thresholds or not. Before we state this proposition, we define $\underline{\theta}$ as follows:

$$\underline{\theta} = \begin{cases} \min\{\theta : \mathcal{L}(\theta) \leq \mathcal{L}(\tilde{\theta}) + \frac{K}{Q} \cdot (\theta + Q - \tilde{\theta})\} & \text{if } \tilde{\theta} \leq \max \mathcal{Y}, \\ -\infty & \text{if } \tilde{\theta} > \max \mathcal{Y}. \end{cases} \tag{7}$$

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 $\underline{\theta} = \tilde{\theta}$). We note that $\underline{\theta}$ is bounded above by $\tilde{\theta}$ (i.e., $\underline{\theta} \leq \tilde{\theta}$), and it may or may not belong to the set $\mathcal{Y}$.

Lemma 1. *For the single-period problem, the following policy is optimal for any* $x_T \in \mathcal{S}^L = \{\theta^* - Q + 1, \ \ldots, \ \theta^*\}$*:*

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

Proof. For any $x_T \in \mathcal{S}^L$, we choose the optimal policy by comparing the cost associated with $y^{[x_T]}$ and $\tilde{\theta}$. Let $\mathrm{Cost}(y|[x_T])$ denote the expected cost associated with the decision of ordering up to y. Then,

$$\mathrm{Cost}(y^{[x_T]}|[x_T]) \ = \ \mathcal{L}(y^{[x_T]}) \ .$$

From the definition of $\tilde{\theta}$, it follows that $\tilde{\theta} \geq \theta^*$. Thus, the following cases are exhaustive.

- Case $\tilde{\theta} > \max \mathcal{Y}$. The definition of $\mathcal{Y}$ implies $\mathcal{L}(y^{[x_T]}) \leq \mathcal{L}(\tilde{\theta})$. Thus, $y_T(x_T) = y^{[x_T]}$ for any value of $y^{[x_T]}$, which corresponds to the first case of $y_T(x_T)$ since $y^{[x_T]} \leq \max \mathcal{Y} < \tilde{\theta}$.

- Case $\tilde{\theta} \in \mathcal{Y}$ and $y^{[x_T]} > \tilde{\theta}$. In this case, $x_T < \theta^*$ and $y^{[x_T]} > \tilde{\theta} \geq \theta^*$. From the definition of $\tilde{\theta}$,

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

  where the last inequality implies $Cost(\tilde{\theta}|[x_T]) \leq Cost(y^{[x_T]}|[x_T])$.

- Case $\tilde{\theta} \in \mathcal{Y}$ and $\underline{\theta} \leq y^{[x_T]} \leq \tilde{\theta}$. Due to the definition of $\underline{\theta}$, we have

$$\mathcal{L}(\underline{\theta}) \ \leq \ \mathcal{L}(\tilde{\theta}) + \frac{K}{Q}(\underline{\theta} + Q - \tilde{\theta}) \tag{8}$$

  For $\underline{\theta} \leq y^{[x_T]} \leq \theta^*$, we have

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

since $\mathcal{L}(y^{[x_T]})$ is a decreasing function and $\frac{K}{Q}(y^{[x_T]} + Q - \tilde{\theta})$ is an increasing function in the given range of $y^{[x_T]}$. Combining the last inequality with (8), we have

$$\mathcal{L}(y^{[x_T]}) \leq \mathcal{L}(\tilde{\theta}) + \frac{K}{Q}(y^{[x_T]} + Q - \tilde{\theta}).$$

For $\theta^* < y^{[x_T]} \leq \tilde{\theta}$, we have

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

since $\mathcal{L}(y^{[x_T]})$ and $\frac{K}{Q}(y^{[x_T]} + Q - \tilde{\theta})$ are both increasing functions in the given range of $y^{[x_T]}$. Hence $Cost(y^{[x_T]}|[x_T]) \leq Cost(\tilde{\theta}|[x_T])$.

- Case $y^{[x_T]} < \underline{\theta}$. Due to the definition of $\underline{\theta}$,

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

which implies that $y^{[x_T]} = \tilde{\theta}$.

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

Proposition 1 shows that the optimal policy on $\mathcal{S}^L$ can be characterized by two thresholds $\underline{\theta}$ and $\tilde{\theta}$, where the value of $\underline{\theta}$ belongs to the set $\mathcal{Y}$ but $\tilde{\theta}$ may not. If $y^{[x_T]}$ falls in the interval $[\underline{\theta}, \tilde{\theta}]$, it is optimal to order up to $y^{[x_T]}$, in which case, there is no partial batch. Otherwise, if $y^{[x_T]} \in \mathcal{Y} \setminus [\underline{\theta}, \tilde{\theta}]$, then it is optimal to order up to $\tilde{\theta}$, in which case the batch is partial. Thus, the optimal policy has a nice threshold structure characterized by two parameters.

We are now ready to state the optimal policy for any value of x_T.

Theorem 1. *For the single-period problem, the following policy is optimal:*

$$y_T(x_T) \quad = \quad \begin{cases} x_T & \text{if } x_T > \theta^* \\ y^{[x_T]} & \text{if } x_T \leq \theta^*, \text{ and } y^{[x_T]} \in [\underline{\theta}, \tilde{\theta}] \\ \tilde{\theta} & \text{if } x_T \leq \theta^* \text{ and } y^{[x_T]} \notin [\underline{\theta}, \tilde{\theta}]. \end{cases}$$

Proof. If $x_T > \theta^*$, then $\mathcal{L}(x_T) \leq \mathcal{L}(y)$ for any $y \geq x_T$. Thus, it is optimal not to order. Note that Lemma 1 has shown the required result if $x_T \in \mathcal{S}^L$. For $x_T \leq \theta^* - Q$, the optimal decision at x_T is the same as the optimal decision at $x_T + Q$ (by Proposition 1(c)). Thus, the result of Lemma 1 easily generalizes to any $x_T \leq \theta^* - Q$, and we complete the proof. $\square$

We discuss the properties of the solution given in Theorem 1. The first case ($x_T > \theta^*$) corresponds to the case where there is too much inventory initially, and thus no additional units are ordered. In the remaining two cases, the after-ordering inventory level is either the full-batch solution, $y^{[x_T]}$, or a "partial batch" solution, $\tilde{\theta}$ (which may coincide with the full batch solution, $y^{[x_T]}$). If $y^{[x_T]} \geq \tilde{\theta} + 1$ then changing the after ordering inventory level from $y^{[x_T]}$ to $\tilde{\theta}$ costs $\frac{K}{Q}$. However, if $\tilde{\theta} = y^{[x_T]} + 1$ then changing the after ordering inventory level from $y^{[x_T]}$ to $\tilde{\theta}$ costs $\frac{K}{Q}(Q-1)$ which is larger than $\frac{K}{Q}$ for $Q > 2$, thus, in this case, it is not attractive to increase the inventory level, and the after-ordering inventory level remains at $y^{[x_T]}$. This explains the asymmetry of the optimal action around $\tilde{\theta}$.

The following result is a corollary of Theorem 1, and shows structural relations between $\underline{\theta}, \tilde{\theta}$, and K. In particular, when K is sufficiently small, then it is optimal to order up to $\tilde{\theta}$, which itself converges to θ^*. 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 $y^{[x_T]}$ for any x_T. This is exactly the interval-based modified base-stock policy for the batch-ordering problem with full batches only.

Corollary 1. *$\tilde{\theta}$ is non-decreasing and $\underline{\theta}$ 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_T, the order-up-to-$\tilde{\theta}$ policy is optimal if $K \leq \underline{K}$ and the order-up-to-$y^{[x_T]}$ policy is optimal if $K \geq \overline{K}$.*

Proof. 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 (6), $\tilde{\theta}_K$ can easily be shown to be non-decreasing in K.

Now we prove that $\underline{\theta}_K$ is non-increasing in K. If $\tilde{\theta}_K > \max \mathcal{Y}$ then $\underline{\theta}_K = -\infty$ from (7). It remains to prove that $\underline{\theta}_K$ is non-increasing in K if $\tilde{\theta}_K \in \mathcal{Y}$ in the interval $[\min \mathcal{Y}, \tilde{\theta}_K]$. For any $K \geq 0$, define

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

which is a linear function of z. Let K_1 and K_2 be real numbers such that $K_1 \leq K_2$ and $\tilde{\theta}_{K_1}, \tilde{\theta}_{K_2} \in \mathcal{Y}$. By the earlier result, $\tilde{\theta}_{K_1} \leq \tilde{\theta}_{K_2}$. We claim that

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

From the definition of $\underline{\theta}_{K_1}$ and $\underline{\theta}_{K_2}$ given in (7), the above claim implies that if both $\underline{\theta}_{K_1}$ and $\underline{\theta}_{K_2}$ belong to $\mathcal{Y}$, then $\underline{\theta}_{K_1} \geq \underline{\theta}_{K_2}$, as required.

To prove the claim, let $\Delta = \tilde{\theta}_{K_2} - \tilde{\theta}_{K_1} \geq 0$. Note that

$$\begin{aligned} l_{K_1}(\tilde{\theta}_{K_2}) - K_1 &= \mathcal{L}(\tilde{\theta}_{K_1}) + \frac{K_1}{Q} \cdot (\tilde{\theta}_{K_2} + Q - \tilde{\theta}_{K_1}) - K_1 \\ &= \mathcal{L}(\tilde{\theta}_{K_1}) + \frac{K_1}{Q} \cdot (\tilde{\theta}_{K_2} - \tilde{\theta}_{K_1}) \\ &\leq \mathcal{L}(\tilde{\theta}_{K_2}) \\ &= l_{K_2}(\tilde{\theta}_{K_2}) - K_2 \ , \end{aligned} \tag{9}$$

where the inequality follows from the convexity of $\mathcal{L}$ and the definition of $\tilde{\theta}_{K_1}$ by (6). 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 (9) 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 (6) and (7), it is easy to see that both of these quantities converge to θ^* as $K \downarrow 0$. Thus, for sufficiently small K, we obtain that $\tilde{\theta}$ becomes θ^* and the set $[\underline{\theta}, \tilde{\theta}]$ becomes to a singleton set $\{\theta^*\}$; then, the optimal policy given in Theorem 1 becomes

$$y_T(x_T) = \begin{cases} x_T & \text{if } x_T > \theta^* \\ \theta^* & \text{if } x_T \leq \theta^* \ , \end{cases}$$

which is the order-up-to-θ^* policy. Now, for sufficiently large K, it follows that $\tilde{\theta}_K$ exceeds $\max \mathcal{Y}$, in which case, $\underline{\theta}_K = -\infty$ and the optimal policy in Theorem 1 becomes

$$y_T(x_T) = \begin{cases} x_T & \text{if } x_T > \theta^* \\ y^{[x_T]} & \text{if } x_T \leq \theta^* \ , \end{cases}$$

which is the order-up-to-$y^{[x_T]}$ policy. $\square$

3.2 Relaxed Problem and the Reduced MDP Approach

The multiple-period problem is difficult to analyze because of the lack of convex 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 introduced in Section 3.3 (this policy performs very well as presented in Section 4).

For this relaxation, we no longer impose the constraint $y_t \geq x_t$, and allow the possibility that inventory can be scrapped. The cost of scrapping inventory is also given by the same c function defined in (2) and (3), even for the negative q values. For example, if $q \in (-(n+1)Q, -nQ)$ for some nonnegative integer n, then

$$c(q) \;=\; K \cdot (\lceil q/Q \rceil - q/Q) \;=\; K \cdot (-n - q/Q) \;=\; K \cdot (|q| - nQ)/Q \;.$$

We refer to this problem as the *multiple-period relaxed problem*. The following lemma partially characterizes the optimal policy.

Lemma 2. *For the multiple-period 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.*

Proof. 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 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 $\hat{y}_t \in \mathcal{Y}$, and also obtain $c(\hat{y}_t - \hat{x}_t) = c(y_t - x_t)$, where x_t and $\hat{x}_t$ denote the before-ordering inventory levels under the original policy and under the new policy, respectively. Thus, for any $T \geq 1$,

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

Thus, we conclude that the new policy is optimal. □

Lemma 2 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 $[x_t] = [\hat{x}_t]$. (This is an obvious extension of Proposition 1(c) to the relaxed problem.) Thus, the optimal action depends on x_t only through the modular arithmetic class to which it belongs, i.e., $[x_t]$.

Therefore, we can define a reduced Markov Decision Process (MDP) with the state space indexed by $\mathcal{Y}$ (corresponding to the inventory level before ordering) and the action space $\mathcal{Y}$ (corresponding to the inventory level after ordering). 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

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

For the transition probability, we define $p^{[i]} = P\{D \equiv [i]\}$, which is the probability that the demand belongs to the set $\{i + kQ \mid k = 0, 1, 2, \ldots\}$. Clearly, $p^{[0]} + p^{[1]} + \cdots + p^{[Q-1]} = 1$. 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 we consider the long-run average-cost criterion.)

Theorem 2. *For the 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,*

$$\hat{C} \;\leq\; C^* \;=\; \tilde{C}^* - E[D] \cdot K/Q \ .$$

Proof. From Lemma 2 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. □

Section 3.3 includes a discussion identifying conditions under which the lower bound stated by Theorem 2, $\hat{C}$, is tight. In the numerical experiments that we conduct in Section 4, we have observed that the average and the maximum gap between the optimal solution and the lower bound are 0.01% and 1.13%, respectively. We conclude this section with the following observation.

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

Proof. Since $p^{[0]} = p^{[1]} = \cdots = p^{[Q-1]}$, 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 inventory policy is first to order up to the interval $\mathcal{S}^L$ using full batches only, and then order up to the quantity specified by Theorem 1. □

3.3 The Multiple-Period Problem and Heuristic Policies

Since the original multiple-period problem presented in Section 2.2 is difficult to analyze, we have considered the single-period problem and the relaxed version of the original problem in Section 3.1 and Section 3.2, respectively. We now return to the original problem. We first establish the bounds on the optimal ordering policy, and propose two types of heuristic methods based on our earlier discussion.

The following proposition establishes the upper and lower bounds on the optimal action $y^*(x)$, analogous to Proposition 1.

Proposition 3. *There exists an optimal policy $y^*(x)$ for the original multiple-period problem satisfying the following properties:*

(a) For any $x \geq \theta^*$, $y^*(x) = x$.

(b) For any $x < \theta^*$, $y^*(x) \in \{\theta^{[x]},\ \theta^{[x]} + 1,\ \ldots,\ \theta^{[x]} + Q\}$.

(c) For any pair of $x', x'' < \theta^*$ satisfying $[x'] = [x'']$, $y^*(x') = y^*(x'')$.

Proof. (a) We show that there exists an optimal policy such that, for any sample path of demand, $y_t = x_t$ whenever $x_t \geq \theta^*$. Suppose that $\{(x'_t, y'_t)|t = 1, 2, \ldots\}$ denote a sequence of before-ordering and after-ordering inventory levels in a system by any policy. We define an alternate policy such that the sequence of inventory levels in a system managed by this policy, denoted by $\{(x''_t, y''_t)|t = 1, 2, \ldots\}$, satisfies

$$y''_t \;=\; \begin{cases} x''_t & \text{if } x''_t \geq \theta^* \\ \min\{y'_t, y^{[y'_t]}\} & \text{if } x''_t < \theta^*. \end{cases}$$

Then, an inductive argument shows the following results for each $t \geq 1$: (i) $y''_t \leq y'_t$; (ii) if $x''_t \geq \theta^*$, then $\theta^* \leq y''_t \leq y'_t$; (iii) if $x''_t < \theta^*$, then $[y''_t] = [y'_t]$ and either $y''_t = y'_t$ or $y''_t \in \mathcal{Y}$. Therefore, we can easily establish

$$\mathcal{L}(y''_t) \quad \leq \quad \mathcal{L}(y'_t) \qquad \text{for each } t \geq 1 \ . \tag{10}$$

Furthermore, it can be shown that

$$\sum_{t=1}^{T} c(y''_t - [x''_t]) \quad \leq \quad \sum_{t=1}^{T} c(y'_t - [x'_t]) \qquad \text{for any } T \geq 1. \tag{11}$$

It follows that $\sum_{t=1}^{T} C_t([x''_t], y''_t) \leq \sum_{t=1}^{T} C_t([x'_t], y'_t)$ holds for any $T \geq 1$. Therefore, we conclude that the alternate policy is also optimal.

(b) The existence of the optimal policy with the property $y^*(x) \leq \theta^{[x]} + Q$ whenever $x < \theta^*$ follows directly from the construction given in part (a). Now, suppose that $\{(x'_t, y'_t) | t = 1, 2, \ldots\}$ denote a sequence of before-ordering and after-ordering inventory levels in a system by an optimal policy such that $y'_t = x'_t$ if $x'_t \geq \theta^*$. Suppose, by way of contradiction, that the property $y'_t \leq \theta^{[x'_t]} + Q$ whenever $x'_t < \theta^*$ does not hold. Then, we construct an alternate policy such that

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

In this alternate policy, one does not order if the current inventory is at least θ^*. Suppose $x''_t < \theta^*$. If $y'_t \leq \theta^{[x''_t]}$, then the full-batch solution of the alternate system, $\theta^{[x''_t]}$ is closer to θ^* than the solution of the original system. Otherwise, we order up to the smaller of y'_t and $\theta^{[x''_t]} + Q$. This ensures $\mathcal{L}(y''_t) \leq \mathcal{L}(y'_t)$. If $y''_t \neq \theta^{[y'_t]}$, the $c(\cdot)$ cost of the alternate system is zero; if $y''_t = \theta^{[y'_t]}$, then it can be argued that the cost of not ordering full batches, $c(y''_t - [x''_t])$, in the alternate system in the current period, can be accounted by the cost of not ordering full batches by the original system in the current or previous periods. Thus, it can be shown that both (10) and (11) hold, and we conclude that the alternate policy is also optimal.

(c) This result follows from part (b) and the fact that, for any pair of $x', x'' < \theta^*$ satisfying $[x'] = [x'']$, $C_t([x'], y) = C_t([x''], y)$ holds for any $y \geq \max\{x', x''\}$. $\square$

We now propose two types of policies that are based on our discussion in Sections 3.1 and 3.2. Under these policies, the inventory level after ordering belongs to $\mathcal{Y}$, an interval of length Q. They differ regarding which element of Q is chosen as a function of the starting inventory level in each period.

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:

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

We interpret this policy as a generalization of the base-stock policy. If the initial inventory exceeds θ^*, then we do not order. Otherwise, we consider θ_U as a target order-up-to level. If the partial batch 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 batch cost K for a small number of units. We refer to this policy as the IB policy. Note that the IB policy is a base-stock policy if $\theta_L = \theta_U$.

The optimal choice of (θ_L, θ_U) can be attained by solving a two-dimensional optimization problem. A special 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 IB policy is equivalent to the single-product version of the (Q, S) policy in Cachon (2001). In his paper, the motivation of this heuristic policy is to impose a minimum load percentage in each batch order. In our paper, the motivation of this policy comes from the fact that it is optimal for the myopic problem. Note that we reach the same policy structure using a different approach, which strengthens the attractiveness of this class of policy.

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. 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:

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

To implement this policy, we first need to solve the relaxed MDP problem, which has Q states, each with Q possible actions. We now state the following theorem that relates heuristic policies with the optimal policy.

Theorem 3. *Suppose that $D \geq Q$ holds with probability 1. Then, the following statements hold.*

(a) The RMB policy is optimal.

(b) If $p^{[i]} = 1/Q$ holds for each $i \in \{0, 1, \ldots, Q-1\}$, then the solution to the myopic problem is optimal.

Proof. For part (a), we assume that $x_1 \leq \min \mathcal{Y}$. This assumption is without loss of generality since $E[D] > 0$. 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 2. We conclude that the RMB policy is optimal.

Part (b) follows directly from part (a) and Proposition 2. □

4. Numerical Study

The purpose of this section is twofold. We first test in Section 4.1 the performance of the proposed heuristics for the partial batch ordering problem. Then, in Section 4.2 we investigate the full batch size ordering policies that are frequently adapted in practice and studied in the literature. In particular, we compare the optimal full batch ordering policy with the optimal partial batch ordering policy and build managerial insights as to when it is safer to operate with full batch sizes and when not.

In our numerical experiments, we consider all possible combinations of the following problem parameters: Gamma demand distribution with a mean of 25 units per period, with

$$\begin{aligned} \text{Coeff. of variation } (CV) &\in \{0.2, 0.5, 1.0, 1.5\} \\ h &= 1 \\ b &\in \{2, 5, 10, 50\} \\ K &\in \{2, 5, 10, 50, 100, 200\} \\ Q &\in \{5, 10, 25, 50, 100, 200\} \; , \end{aligned}$$

resulting in $4 \times 4 \times 6 \times 6 = 576$ problem instances. We evaluate the performance of various heuristic policies and report it using the relative error with respect to the optimal policy, which we obtain by solving a dynamic program. We denote the relative cost by Δ, where a subscript is used to denote the type of a heuristic policy.

4.1 Performance of the Heuristics

We first analyze the performance of IB policy (IBP). Recall that IBP has two parameters, one is an order-up-to point, and the other is a threshold for switching from order-up-to to full batch ordering. We have searched for the parameter space to find the best values for the parameters. In our computation, the minimum, average, and maximum values of the relative error, Δ_{IBP}, over all cases are

$$\begin{aligned} \min\{\Delta_{IBP}\} &= 0.000\% \ , \\ \text{average}\{\Delta_{IBP}\} &= 0.001\% \ , \qquad \text{and} \\ \max\{\Delta_{IBP}\} &= 0.253\% \ . \end{aligned}$$

Furthermore, IBP gives the optimal solution in 571 of the total 576 problem instances.

If the optimal policy has the same two-parameter structure mentioned above, then IBP always gives the optimal policy. While this is true in many of the cases, the optimal policy does not always have a simple structure. As an example, the optimal policy with $CV = 0.2$, $b = 10$, $K = 100$, and $Q = 100$ exhibits three order-up-to points as shown is Figure 1.

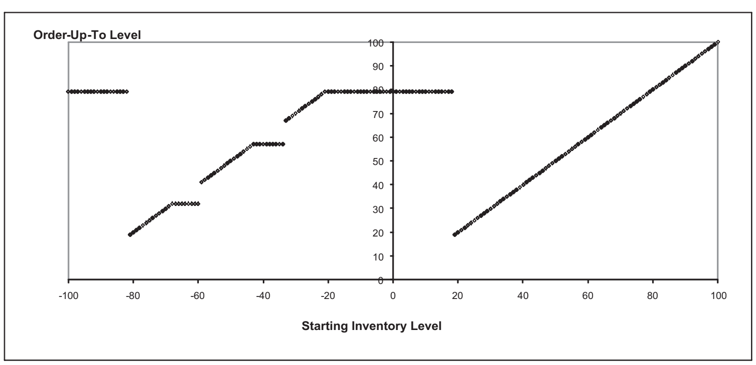

Figure 1: An Optimal Policy with Three Order-Up-To Points. $CV = 0.2$, $b = 10$, $K = 100$, and $Q = 100$

While IBP performs extremely well in general, the computational time to search for the best parameters may increase discouragingly fast in Q if an exhaustive search of the two-dimensional space is conducted (as in our experiment). However, the second heuristic method we study, RMB policy (RMBP) requires negligible amount of computational time, and gives extremely good results. The minimum, average, and maximum values of Δ_{RMBP} are

$$\begin{aligned} \min\{\Delta_{RMBP}\} &= 0.000\% \ , \\ \text{average}\{\Delta_{RMBP}\} &= 0.001\% \ , \qquad \text{and} \\ \max\{\Delta_{RMBP}\} &= 0.194\% \ . \end{aligned}$$

Unlike IBP, RMBP may detect optimal policies with multiple order-up-to points. The problem instances that RMBP failed to find the optimal solution are given in Table 1. As can be observed from this table, these problem instances (10 in total out of 576 instances) have large batch sizes and large fixed costs.

Table 1: Problem instances that RMBP fails to find the optimal solution

b	50	50	50	5	50	50	50	10	10	10
Q	50	100	200	50	50	100	200	50	100	200
K	100	100	100	50	50	50	50	50	50	50
$\Delta_{RMBP}(in\%)$	0.006	0.027	0.027	0.047	0.050	0.062	0.062	0.083	0.194	0.194

In all problem instances where demand has a CV higher than 0.2, both IBP and RMBP found the optimal solution. In all of these instances, the optimal policy has at most one order-up-to point (an optimal policy without an order-up-to point corresponds to the policy of ordering multiples of full batches only, which can be detected by both policies). We observe that as the CV decreases, the total system cost becomes more sensitive to the order-up-to points, and hence the number of distinct order-up-to points tend to increase in the optimal policy. Therefore, a "fine tuning" on the order-up-levels becomes essential if the demand is more predictable. On the other hand, when the demand is less predictable due to higher variability, such an action is unnecessary and the system prefers to operate with at most one order-up-to point. (This result is consistent with the general intuition from Proposition 2 – since the uniform distribution over Q equivalence classes signifies the least predictable demand.) To validate this, we have considered four additional lower CV values and found the optimal solution of all $4 \times 6 \times 6 = 144$ problem instances for each CV. It turns out that there are 24, 33, 35, and 54 problem instances with more than one order-up-to point when the CV of demand is 0.15, 0.10, 0.05, and 0.01, respectively. (In particular, there is a case with five different order-up-to points when CV is 0.01, which is the maximum that we obtained in our numerical tests.)

While both IBP and RMBP give close-to-optimal solutions, the policy parameters are found by solving non-trivial mathematical expressions – either searching over a two-dimensional space or solving an MDP formulation. Thus, we have tested the myopic policy (MP), which is a special member of IBP. This policy could be a preferred alternative in practice as it requires resolution of much simpler mathematical expressions. In what follows, we test the performance of the myopic policy.

We first note that the value of CV has a significant effect on the performance of MP. The minimum, average, and maximum Δ_{MP} values for different CV values over all problem instances is summarized in Table 2.

Table 2: Performance of MP for different CV values

CV	$\min\{\Delta_{MP}\}$	Average$\{\Delta_{MP}\}$	$\max\{\Delta_{MP}\}$
0.2	0%	3.34%	58.03%
0.5	0%	1.55%	20.21%
1.0	0%	0.64%	12.48%
1.5	0%	0.41%	11.44%

In Table 3, we present Δ_{MP} values for some particular problem instances averaged over all b values. Myopic policy performs very well for lower values of K. We might expect this since the myopic policy is optimal for the special case of our problem with $K = 0$. But as K increases up to some moderate values, Δ_{MP} also increases. Nevertheless, under considerably large K values, the optimal policy tends to behave like a full batch ordering policy, which is easier to be detected by the myopic policy due to high penalty of partial batch ordering, and hence Δ_{MP} starts to decrease. The myopic policy also performs in a similar manner for different batch sizes. For small and very large batch sizes, MP performs very well, because for small Q values, the penalty of partial batch ordering is high, and for considerably high Q values, the problem approaches to an uncapacitated problem.

Table 3: Average Δ_{MP} Values (in %) for CV = 0.2

	K					
Q	2	5	10	50	100	200
5	0	0	0	0	0	0
10	0	0	0	0	0	0
25	0.01	0.07	0.09	0.12	0	0
50	0	0.07	0.72	2.11	0.02	0
100	0	0	0.05	23.25	4.45	0.12
200	0	0	0	23.41	46.84	18.71

4.2 Full Batch Ordering Policy

In this section, we test the performance of the full batch ordering (FBO) policy. We study this 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.

Table 4: Δ_{FBO} Values (in %) Averaged Over All b and CV Values

	Q						
K	5	10	25	50	100	200	Overall
2	0.1	2.16	15.76	37.62	56.62	67.95	30.03
5	0	0.73	8.79	32.25	51.41	62.82	26.00
10	0	0	4.44	26.28	45.09	56.33	22.02
50	0	0	0.54	1.02	13.99	25.23	6.80
100	0	0	0	0.11	4.56	15.26	3.32
200	0	0	0	0	0.15	8.36	1.42
Overall	0.02	0.48	4.92	16.21	28.64	39.33	14.93

Table 4 shows that FBO coincides with the optimal policy when Q is low and K is high, as expected. For the other extreme of high Q and low K, the additional cost incurred by imposing FBO policy is as high as 67.95%. This means that adapting FBO – as is done often in practice – is not rational for relatively small fixed costs or relatively high full batch sizes.

The performance of FBO with respect to CV and b is more complicated. First of all, rather interestingly, we observe that FBO performs much better as CV increases, as shown in Table 5. This is because the optimal policy tends to order in full batches for high CV values;

Table 5: Δ_{FBO} Values (in %) Averaged Over All K and Q Values

b	CV				Overall
	1	2	3	4	
2	87.08	35.64	17.81	12.66	38.3
5	82.04	32.21	12.9	6.96	33.53
10	78.35	28.53	10.03	4.6	30.38
50	66.7	21.85	6.43	2.4	24.35
Overall	78.54	29.56	11.79	6.66	31.64

Table 6: Batch Utilization (in %) of the Optimal Policy Averaged Over All K and Q values

b	CV				Overall
	0.2	0.5	1.0	1.5	
2	78.9	82.12	88.79	91.88	85.42
5	78.04	80.85	88.8	92.59	85.07
10	77.68	80.3	88.95	92.86	84.95
50	77.41	79.82	88.89	93.2	84.83
Overall	78.01	80.78	88.85	92.63	85.07

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 *batch utilization* values (average batch size occupied by an order relative to the full batch size), and reported them in Table 6.

Finally, in practice, full batch ordering policy is favored when the backorder costs are relatively high (Tanrikulu et al., 2009). The performance of FBO in Table 5 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, resulting in a smaller proportion of the batch-related cost in the total cost. 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.[2]

Moreover, rather surprisingly, even the best full batch ordering policy performs quite poorly under low demand variability. This can be explained by the fact that, with full batch ordering, low demand uncertainty is more likely to result in more predictable overage or underage; the single-period expected cost is smoother when demand is more uncertain.

[2] In our computation, we have found the following examples without this monotonicity property: (Q, K, CV) triplets of $(100, 50, 0.2)$, $(100, 50, 0.5)$ and $(200, 100, 0.2)$.

In particular, the average batch utilization in the optimal policy is as low as 78% when CV=0.2, compared to 93% when CV=1.5. The average penalty of adapting a full batch ordering policy is as high as 79% when CV=0.2, which is only a lower bound considering that the full batch ordering policy could be adapted in a suboptimal way in practice.

5. Conclusions

In this paper, we pose the partial batch ordering problem in an inventory system. 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 reformulating the cost structure of the problem, we identify a "relaxed" problem that can be easily solved and yields a lower bound for the optimal cost. We also characterize the optimal ordering policy for the single period version of the original problem. Both of these results aid us to identify two sets heuristic algorithms for the original problem. Our numerical results indicate that both heuristics perform extremely well.

Finally, we investigate the full batch size ordering policies that are frequently adapted in practice and studied in the literature, by comparing the best full batch ordering policy with the optimal partial batch ordering policy. We show that restricting ordering quantities to full batch sizes does not always perform well, especially when the cost per batch is high or the batch size is big. This restriction increases, on average, 31.64% of total cost in our experiments, and the optimal partial-batch policy exhibits, on average, only 85.07% of utilization. The performance of the full batch ordering policy deterioates as the demand variability decreases; in comparison, as the backorder cost increases, its performance improves, usually but not always.

This research can be extended in several ways. It may be the case that there are several different options for the possible batch sizes, 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.

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**. 211–229.

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**. 376–389.

Federgruen, A. and P. Zipkin. 2000. An Inventory Model with Limited Production Capacity and Uncertain Demands II. The Discounted-Cost Criterion. *Mathematics of Operations Research*. **48**. 376–389.

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.

Iwaniec, K. 1979. An Inventory Model with Full Load Ordering. *Management Science*. **25**(4). 374–384.

Lippman, S. 1969. Optimal Inventory Policy with Multiple Set-up Costs. *Management Science*. **16**(1). 118–138.

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. 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. 2009. Joint Replenishment Problem with Truck Cost Structures. *International Journal of Production Research*. to appear.

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.

Zheng, Y., F. Chen. 1992. Inventory policies with quantized ordering. *Naval Research Logistics*. **39**. 285–305.