Approximation Algorithms for Capacitated Stochastic Inventory Systems with Setup Costs

Cong Shi*, Huanan Zhang*, Xiuli Chao*, Retsef Levi†

* Industrial and Operations Engineering, University of Michigan, Ann Arbor, MI 48109

{shicong, zhanghn, xchao}@umich.edu

† Sloan School of Management, MIT, Cambridge, MA 02139

retsef@mit.edu

Abstract

We develop the first approximation algorithm with worst-case performance guarantee for capacitated stochastic periodic-review inventory systems with setup costs. The structure of the optimal control policy for such systems is extremely complicated, and indeed, only some partial characterization is available. Thus, finding provably near-optimal control policies has been an open challenge. In this paper we construct computationally efficient approximate optimal policies for these systems whose demands can be nonstationary and/or correlated over time, and show that these policies have a worst-case performance guarantee of 4. We demonstrate through extensive numerical studies that the policies empirically perform well, and they are significantly better than the theoretical worst-case guarantees. We also extend the analyses and results to the case with batch ordering constraints, where the order size has to be an integer multiple of a base load.

Key words: inventory, setup cost, capacity, approximation algorithms, bounds, randomized cost-balancing, worst-case performance guarantees

1 Introduction

In this paper we study capacitated stochastic periodic-review inventory systems with setup costs. The demand process may be nonstationary (time-dependent) and correlated over time, capturing demand seasonalities and forecast updates.

These systems are fundamental but notoriously hard to analyze in both theory and computation. If the ordering capacity in each period is infinity, it is well-known that state-dependent (s, S) type of policies are optimal for inventory systems with setup costs under independent demand processes. This structure for optimal policies also holds true for exogenous Markov-modulated demands (e.g., Cheng and Sethi [7]) and models with advance demand information (e.g., Gallego and Özer [9]). One might expect that some form of modified (s, S) policies is optimal for the capacitated case, but all studies have rejected the conjecture. In fact, even when the demands in different periods are independent and identically distributed, the structure of the optimal control policies is very

1

complicated and only some partial characterization is available in the literature. Thus, the design and computation of a provably near-optimal control policy have been an open challenge.

It should be noted that computing the optimal control policy using dynamic programming may not be possible due to the *curse-of-dimensionality*, i.e., the need to keep track of a state variable of large dimension. For example, the demand process for our model may be nonstationary, driven by the state-of-economy or state-of-the-world (e.g., the Markov modulated demand process), or it may be a forecast-related demand process such as the Martingale Model of Forecast Evolution (MMFE, see for example, Heath and Jackson [14]) in which the updated forecast (as well as the realization of the supply capacity in the next period) is the original forecast plus a random error with mean zero (see e.g. Lu et al. [22]). In these scenarios, the demand structure leads to a multi-dimensional stochastic dynamic program and computing the optimal policies is usually intractable.

1.1   Main results and contributions of this paper

The major results and contributions of this paper are summarized as follows. We also point out the major distinction of our proposed algorithms from previous work, in particular, Levi et al. [20] and Levi and Shi [21].

Algorithms and their worst-case analysis.   We develop the first approximation algorithms for capacitated stochastic periodic-review inventory systems with setup costs under a correlated, nonstationary and evolving stochastic demand structure. The policy proposed will be referred to as a *randomized 1/2-balancing policy (R/2)*. We show that the proposed policies admit a constant worst-case performance guarantee of 4, regardless of any specific demand instance or input parameters. Note that this constant worst-case performance guarantee does not scale with the system size or the length of the planning horizon or the input parameters. Since the structure of optimal policies for these systems is not well understood, the proposed inventory control policies provide valuable insights into how various cost components should be balanced.

As mentioned in our literature review below, Levi and Shi [21] developed a 3-approximation algorithm for the uncapacitated model with setup costs using an exact randomized balancing (i.e., exactly balance the marginal holding cost, the forced backlogging cost and the setup cost), and Levi et al. [20] proposed a concept of *forced backlogging cost accounting* for the capacitated model without setup costs. However, exact balancing is not achievable in the presence of both capacity constraints and the setup cost. The main source of difficulty lies in the fact that the policy may not be able to order a specific quantity that makes the marginal holding cost equal to the setup cost, since this particular quantity may exceed the ordering capacity; in such cases, the policy has to truncate an order at the capacity level. The approach employed in Levi and Shi [21] fails to work in this case. Instead of exact balancing, our proposed $R/2$ policy *almost* balances the marginal holding or forced backlogging cost with *half* of the setup cost. We provide a unified and much simpler analysis of Levi and Shi [21] in the well-behaved cases, and a novel analysis in the ill-behaved cases.

We also extend our results to capacitated model with setup cost under *batch order constraints*. With the batch order constraint, each order quantity has to be an integer multiple of a pre-specified base load, e.g., a truck-load. We refer interested readers to Veinott [25], Chao and Zhou [3], Chen

[4] and Huh and Janakiraman [16] for details concerning batch orders on the case with infinite ordering capacity. We propose a *modified randomized 1/2-balancing policy* and show that the worst-case performance guarantee of the proposed policy is still 4.

Empirical performance. We show how these policies can be parameterized to create a broader class of policies. We demonstrate through extensive computational studies that the proposed algorithms perform well in an empirical study (around $5\% - 15\%$ from the optimal cost), which is significantly better than the theoretical worst-case performance guarantees. The proposed inventory control policies are computationally efficient with a computational complexity of $O(T^2)$ where T is the length of the planning horizon, which is very efficient compared to the dynamic programming approach that suffers from the well-known curse of dimensionality.

1.2 Literature review

Stochastic periodic-review inventory systems have attracted the attention of many researchers over the years. The dominant paradigm in the existing literature has been to formulate and analyze these problems using dynamic programming. For many uncapacitated inventory systems with setup costs, it can be shown that some form of *(s,S) policies* are optimal (see, e.g., Scarf [24], Veinott [27]). Cheng and Sethi [7] have extended the optimality proof to exogenous Markov-modulated demands that capture cycles and seasonality to some extent. Gallego and Özer [9] have established their optimality for models under advance demand information, a demand model that allows correlation and forecast updates. Myopic policies seem to perform well for some scenarios in uncapacitated systems and are even optimal in some specific settings (see Veinott [26], Ignall and Veinott [17] and Iida and Zipkin [18]). However, capacitated problems are inherently harder, structurally and computationally, compared to their uncapacitated counterparts. The capacity constraint makes future costs heavily dependent on current decisions. Chen and Lambrecht [6] demonstrated that the optimal policy for capacitated inventory systems with setup costs exhibits an $X - Y$ band structure, with $X < Y$. That is, if the inventory level is below X, order the full capacity, and if the inventory level is over Y, order nothing; if the inventory level is between X and Y, however, the ordering policy is complicated and not known. Gallego and Scheller-Wolf [10] and Chen [5] provided some further refinements to this policy, but again, the optimal control policy remains complicated and can only be partially characterized when the inventory level at the beginning of a period is in the middle range. For example, in Gallego and Scheller-Wolf [10], it was shown that the region between the $X - Y$ bands can be further divided into two subregions. In one of them, it is optimal to either order nothing or to bring the inventory level to at least some specified level, that is, there exists a lower bound for the optimal order up-to level in this range; in the other subregion, the parameters of the solution dictate which one of the two cases hold: In the first case it is optimal to order, again to at least some specified level (thus only a lower bound is shown to exist), and in the second, the optimal policy is to either order the full capacity or order nothing. Özer and Wei [23] studied capacitated inventory systems with advance demand information. They established the optimality of a state-dependent modified base-stock policy for inventory systems with zero fixed ordering cost and for the systems with positive fixed costs, they restricted the ordering to the class of all-or-nothing policies and characterized the optimal policies within that class.

For models with infinite ordering capacities and independent and identically distributed de-

mands, Federgruen and Zipkin [8] proposed an algorithm to compute the optimal (s, S) policy in an infinite horizon model. Bollapragada and Morton [2] proposed a simple, myopic heuristic for computing the policies where the demands in different periods are assumed to have the same form of distribution function but with different means, and the coefficient of variation of the demands is assumed to be stationary. Gavirneni [11] designed a simple heuristic to compute (s, S) policies for nonstationary and capacitated model. Guan and Miller [13] proposed an exact and polynomial-time algorithm for the uncapacitated stochastic periodic-review inventory system without backlogging if the stochastic programming scenario tree is polynomially representable. Guan and Miller [12] extended these algorithms to allow for backlogging. Huang and Küçükyavuz [15] considered similar problems but with random lead times. These models allow for stochastic and correlated demands. The main limitation comes from the fact that the number of nodes in the stochastic programming scenario tree (the size of input) is likely to be exponentially large in the size of the planning horizon. Atali and Özer [1] proposed a close-to-optimal heuristic to manage a multi-item two-stage production system subject to Markov-modulated demands and production quantity requirements. All of the existing heuristics and algorithms, either lack any performance guarantees or can only be applied under restrictive assumptions on the demand processes or the input size, and to the best of our knowledge, no efficient computational policies have been reported for capacitated models that admit worst-case performance guarantees.

Our work is closely related to the recent literature on approximation algorithms in stochastic periodic-review inventory systems, first started by Levi et al. [19]. Levi et al. [19] introduced the concept of marginal cost accounting that associates the full planning horizon cost with each decision a particular policy makes. They proposed a dual-balancing policy that admits a worst-case performance guarantee of 2 for the uncapacitated model without setup costs. Subsequently, Levi et al. [20] introduced the *forced marginal backlogging cost-accounting scheme* to analyze the capacitated models without setup costs, and Levi and Shi [21] proposed the *randomized cost-balancing policy* to solve uncapacitated stochastic lot-sizing problems with setup costs. It is worthy to note that the systems studied in these papers all have nice simple structures for their optimal control policies. However and as discussed above, the structure of the optimal control policies for capacitated stochastic inventory models with setup costs is complicated and has not been fully characterized; and designing an approximation algorithm for the capacitated stochastic inventory models with setup costs remained a challenging task.

1.3 Structure of this paper

The remainder of the paper is organized as follows. In Section 2, we present the mathematical model for the capacitated stochastic periodic-review inventory system with setup cost. Section 3 reviews the marginal cost accounting scheme proposed by [20]. More specifically, we present the marginal holding cost accounting scheme in Section 3.1 and the forced backlogging cost accounting scheme in Section 3.2. In Section 4, we propose a novel randomized 1/2-balancing policy and discuss the key ideas. Then we show that the policy has a worst-case performance guarantee of 4 in Section 5. In Section 6 we extend our results to systems with batch order constraints. Finally, Section 7 is devoted to the numerical studies for our proposed policies. The parameterized policies are computationally efficient and perform well under a correlated demand structure with advance demand information (see, e.g., Gallego and Özer [9] and Özer and Wei [23]).

2   Capacitated Periodic-Review Inventory System with Setup Costs

In this section, we provide the mathematical formulation of the capacitated periodic-review inventory system with setup cost. Our model allows for nonstationary and generally correlated demand structure. The ordering *capacity* in each period is denoted by u. The planning horizon is T periods which can be either finite or infinity, and we index the period by $t = 1, \ldots, T$.

Demand structure. The demands $D_1, \ldots, D_T$ over the planning horizon T are random. At the beginning of each period s, we are given what we call an *information set* denoted by f_s. The information set f_s contains all of the information that is available at the beginning of time period s. More specifically, the information set f_s consists of the realized demands $d_1, \ldots, d_{s-1}$ over the interval $[1, s)$, and possibly some exogenous information denoted by $(w_1, \ldots, w_s)$. The information set f_s in period s is one specific realization in the set of all possible realizations of the random vector $F_s = (D_1, \ldots, D_{s-1}, W_1, \ldots, W_s)$. The set of all possible realizations is denoted by $\mathcal{F}_s$. With the information set f_s, the conditional joint distribution of the future demands $(D_s, \ldots, D_T)$ is known. The only assumption on the demands is that for each $s = 1, \ldots, T$, and each $f_s \in \mathcal{F}_s$, the conditional expectation $\mathbb{E}[D_t \mid f_s]$ is well-defined and finite for each period $t \geq s$. In particular, we allow for non-stationarity and correlation between the demands in different periods.

Cost structure. In each period t, $t = 1, \ldots, T$, four types of costs are incurred, a per-unit ordering cost c_t for ordering any number of units at the beginning of period t, a per-unit holding cost h_t for holding excess inventory from period t to $t+1$, a per-unit backlogging penalty b_t that is incurred for each unsatisfied unit of demand at the end of period t, and a *setup cost* K that is incurred in each period with strictly positive ordering quantity. Unsatisfied units of demand are usually called *backorders*. Each unit of unsatisfied demand incurs a per-unit backlogging penalty cost b_t in each period until it is satisfied. In addition, we consider a model with a lead time of $L \geq 0$ periods between the time an order is placed and the time at which it actually arrives. We remark that the analysis and results remain true when the setup cost K_t depends on period t as long as $K_t \geq K_{t+1}$ is satisfied for all t. We assume without loss of generality that the discount factor $\alpha = 1$, and that $c_t = 0$ and $h_t, b_t \geq 0$, for each t (see the discussion in [19]).

System dynamics. The goal is to coordinate a sequence of orders that minimizes the overall expected setup cost, holding cost and backlogging cost. More specifically, in each period t, $t = 1, \ldots, T$, we place an order of $Q_t \in [0, u]$ units. Given a feasible policy PL, the dynamics of the system are described using the following notation. Let NI_t denote the *net inventory* at the end of period t. Thus, NI_t^+ and NI_t^- are inventory on hand and backlog quantities in period t, respectively, where for any real number x, we let $x^+ = \max\{x, 0\}$. Since there is a lead time of L periods, one also considers the *inventory position* of the system, which is the sum of all outstanding orders plus the current net inventory. Let X_t be the *inventory position* at the beginning of period t before the order in period t is placed, i.e., $X_t = NI_{t-1} + \sum_{j=t-L}^{t-1} Q_j$, and $Q_j \in [0, u]$ denotes the number of units ordered in period j. Similarly, let Y_t be the inventory position *after* the order in period t is placed, i.e., $Y_t = X_t + Q_t$. Note that for every possible policy PL, once the information set $f_t \in \mathcal{F}_t$ is known and order Q_t is placed, the values ni_{t-1}, x_t and y_t are known, where these are the realizations of NI_{t-1}, X_t and Y_t, respectively. At the end of each period t, the costs incurred are holding cost $h_t NI_t^+$ and backlogging cost $b_t NI_t^-$. In addition, if the order quantity $Q_t > 0$, then the fixed ordering cost K is incurred. Thus, the total cost of a feasible policy PL is

$$\mathscr{C}(PL) = \sum_{t=1}^{T} \left(h_t NI_t^{+,PL} + b_t NI_t^{-,PL} + K \cdot \mathbb{1}(Q_t^{PL} > 0) \right), \tag{1}$$

where $\mathbb{1}(A)$ is the indicator function taking value 1 if statement "A" is true and 0 otherwise. The objective is to find the optimal ordering decisions Q_t^{PL}, based on information f_t, $t = 1, \ldots, T$, that minimizes the total cost (1).

3 Marginal Cost Accounting Scheme

The cost accounting scheme described in (1) above decomposes the cost by periods. Following Levi et al. [19] and [20], we next describe an alternative cost accounting scheme that is called *marginal cost accounting scheme*. The main idea underlying this approach is to decompose the cost by decisions. That is, the decision in period t is associated with all costs that, after that decision is made, become unaffected by any future decisions, and are only affected by future demands. This may include costs in all subsequent periods.

3.1 Marginal holding cost accounting

Let $D_{[s,t]}$ denote the cumulative demand over the interval $[s, t]$, i.e., $D_{[s,t]} = \sum_{j=s}^{t} D_j$. We first focus on the holding costs and assume, without loss of generality, that units in inventory are consumed on a *first-ordered first-consumed* basis. This implies that the overall holding cost of the q_s units ordered in period s (i.e., the holding cost they incur over the entire horizon $[s, T]$) is a function only of future demands, and is unaffected by any future decision. Specifically, based on the assumption that inventory is consumed on a first-ordered first-consumed basis, the q_s units on order will be used to satisfy demand only when the x_s units presently in the system have been completely consumed. Among these q_s units, the number of those still remaining in inventory at the end of period j (where $j \geq s + L$) is precisely $\left(q_s - (D_{[s,j]} - x_s)^+ \right)^+$. Thus, the total marginal holding cost associated with the decision to order q_s units in period s is, recall that the discount factor $\alpha = 1$, defined to be $\sum_{j=s+L}^{T} h_j \left(q_s - (D_{[s,j]} - x_s)^+ \right)^+$. Note that at the time the order q_s is placed, the inventory position x_s is already known and indeed the marginal holding cost is just a function of future demands. In addition, once the order in period s is determined, the backlogging cost a lead time ahead in period $s + L$, i.e., $b_{s+L} \left(D_{[s,s+L]} - (x_s + q_s) \right)^+$, is also affected only by the future demands. This leads to a marginal cost accounting scheme.

For each feasible policy PL, let H_s^{PL} be the holding cost incurred by the Q_s^{PL} units ordered in period s, for $s = 1, \ldots, T - L$, over the interval $[s, T]$. Then,

$$H_s^{PL} = H_s^{PL}(Q_s^{PL}) = \sum_{j=s+L}^{T} h_j \left(Q_s^{PL} - (D_{[s,j]} - X_s)^+ \right)^+. \tag{2}$$

It is readily verified that when we sum up the marginal holding costs of all unit ordered, we would obtain the total holding cost for all the periods. That is,

$$\sum_{t=1}^{T} h_t NI_t^{+PL} = H_{(-\infty,0]} + \sum_{t=1}^{T-L} H_t^{PL}, \tag{3}$$

where $H_{(-\infty,0]}$ denotes the total holding cost incurred by units ordered before period 1, which is independent of the ordering decisions during the planning horizon $[1, T]$.

3.2 Forced backlogging cost accounting

In capacitated models, it is no longer true that a mistake of ordering too little in the current period can always be fixed by decisions made in the future periods. Levi et al. [20] proposed a new backlogging cost accounting that associates with decision of how many units to order in period t what is called *forced backlogging cost* resulting from this decision in future periods.

Consider some period s. Suppose that x_s is the inventory position at the beginning of period s and that the number of units ordered in period is $q_s \leq u$. Let $\bar{q}_s$ be the resulting unused *slack capacity* in period t, i.e., $\bar{q}_s = u - q_s \geq 0$. Focus now on some future period $t \geq s + L$ when this order arrives and becomes available. Suppose that for some realization of the demands, we have that

$$d_{[s,t]} - (x_s + q_s + (t - s - L)u) > 0. \tag{4}$$

This implies that there exists a shortage in period t, and moreover, even if in each period after period s and until period $t - L$ the orders placed were up to the maximum available capacity, this part of the shortage in period t would still exist and incur the corresponding backlogging cost. The actual shortage may be even higher than (4) and is equal to

$$d_{[s,t]} - (x_s + q_s + \sum_{j \in (s, t-L]} q_j) > 0,$$

(recall that $q_j \leq u$ for each period j). In other words, given our decision in period s, this part of the shortage could not be avoided by any decision made over the interval $(s, t - L]$ (clearly, any order placed after period $t - L$ will not be available by time t). We conclude that, if more units had been ordered in period s, then at least some of the shortage in period t could have been avoided. More precisely, the maximum number of units of shortage that could have been avoided by ordering more units in period s is equal to

$$\min\left\{\bar{q}_s, \left[d_{[s,t]} - (x_s + q_s + (t - s - L)u)\right]^+\right\}.$$

The intuition is that by ordering more units in period s, we could have averted part of the shortage in period t, but clearly not more than the unused slack capacity $\bar{q}_s$, since we could not have ordered in period s more than additional $\bar{q}_s$ units. In this case, we would say that this part of the backlogging cost in period t was forced by the decision in period s. Denote $W_{s,t}$ as the backlogging cost in period t associated with decision made in period s. Then we can write

$$W_{st} \;=\; b_t \min\left\{\left(D_{[s,t]} - (X_s + Q_s + (t - s - L)u)\right)^+, (u - Q_s)\right\}.$$

This is significantly different from the traditional backlogging cost accounting, in which this cost would be associated with period $t - L$. Since the decision at period s could affect all succeeding period's backlogging cost, then the forced backlogging costs that are incurred by any feasible policy PL in a period s is given by

$$\bar{\Pi}_s^{PL} = \sum_{t=s+L}^{T} W_{st}^{PL}. \tag{5}$$

7

It is, again, readily verified that the summation of forced backlogging cost in all periods is equal to the total backlogging cost. That is,

$$\sum_{t=1}^{T} b_t N I_t^{-PL} = \bar{\Pi}_{(-\infty,0]} + \sum_{t=1}^{T-L} \bar{\Pi}_t^{PL}, \tag{6}$$

where $\bar{\Pi}_{(-\infty,0]}$ denotes all the forced backlogging costs of the ordering decisions made before period 1, which is independent of the policy used.

3.3 Total cost of any feasible policy

Let $\mathscr{C}(PL)$ be the total cost incurred by using the control policy PL. By (3) and (6), we can rewrite $\mathscr{C}(PL)$ as

$$
\begin{aligned}
\mathscr{C}(PL) &= \sum_{t=1}^{T} \Big(h_t N I_t^{+PL} + b_t N I_t^{-PL} + K \cdot \mathbb{1}(Q_t^{PL} > 0) \Big) \\
&= \sum_{t=1}^{T-L} \Big(K \cdot \mathbb{1}(Q_t^{PL} > 0) + H_t^{PL} + \bar{\Pi}_t^{PL} \Big) + H_{(-\infty,0]} + \bar{\Pi}_{(-\infty,0]}.
\end{aligned}
$$

Since $H_{(-\infty,0]}$ and $\bar{\Pi}_{(-\infty,0]}$ are constants that are not affected by the policy used, we will ignore them in the subsequent analysis and write the effective cost of a policy PL as

$$\mathcal{C}(PL) = \sum_{t=1}^{T-L} \Big(K \cdot \mathbb{1}(Q_t^{PL} > 0) + H_t^{PL} + \bar{\Pi}_t^{PL} \Big).$$

Clearly, to compare the performances of different policies, it suffices to compare their corresponding effective costs.

4 The Randomized 1/2-Balancing (R/2) Policy

In this section, we propose a policy called *randomized 1/2-balancing policy* ($R/2$, or R-half policy) which aims to strike a balance between three types of costs, namely, the marginal holding cost, the forced backlogging cost, and the setup cost.

There are two sources of difficulties in designing cost-balancing algorithms for capacitated stochastic periodic review inventory systems with setup costs. The first one is that we are unable to perfectly balance the three types of costs mentioned above. For instance, we may not be able to order the quantity that brings the marginal holding cost up to the setup cost K, since the particular quantity can exceed the capacity constraint u. In these cases, the balancing policy has to place a truncated order at the full capacity. This creates difficulties in analyzing the performance bounds of the policy since it is not so clear which cost component of the optimal policy can 'pay' for a constant fraction of the cost incurred by the balancing policy. The second source of difficulty is the need to balance the nonlinear setup cost against the forced backlogging cost that may have large spikes because of the variability of the demands. Thus, the balancing policy needs to employ

a randomized decision rules to make the expected setup cost incurred in each period a continuous function, rather than an indicator function K if an order is placed and 0 otherwise. However, the randomized decision rules also introduce uncertainties in the relationships between the ending inventory position of the optimal policy and the balancing policy. In some periods, it is not a-priori clear how to use the cost of the optimal policy to 'pay' for that of a balancing policy.

To describe the new policy, we modify the definition of the information set f_t to also include the randomized decisions of the randomized balancing policy up to period $t-1$. Thus, given the information set f_t, the inventory position x_t at the beginning of period t is known. However, the order quantity in period t is still *unknown* because the policy randomizes among various order quantities.

4.1 Computing auxiliary balancing quantities and costs

At the beginning of each period t with the realized information set f_s, we can efficiently compute the following auxiliary ordering quantities and costs, since the marginal holding cost $H(\cdot)$ and the forced backlogging cost $\bar{\Pi}(\cdot)$ are given in (2) and (5) in closed forms. First, compute the *balancing quantity* $\hat{q}_t$ and the *balancing cost* θ_t such that

$$\theta_t \triangleq \mathbb{E}[H_t^{R/2}(\hat{q}_t) \mid f_t] = \mathbb{E}[\bar{\Pi}_t^{R/2}(\hat{q}_t) \mid f_t].$$

The balancing quantity perfectly balances the conditional expected marginal holding cost against the conditional expected forced backlogging cost associated with the order $\hat{q}_t$. Since $\bar{\Pi}_t^{R/2}(u) = 0$, it follows that $\hat{q}_t \leq u$. Note that $H_t(\cdot)$ is convex and increasing on $[0, \infty)$ and $\bar{\Pi}(\cdot)$ is convex and decreasing to 0. Thus, $\hat{q}_t$ always exists and can be computed efficiently via bi-section search. Then, compute the *holding-cost-$K/2$ quantity* $\tilde{q}_t$ that solves

$$\mathbb{E}[H_t^{R/2}(\tilde{q}_t) \mid f_t] = \frac{K}{2}.$$

The holding-cost-$K/2$ quantity makes the conditional expected marginal holding cost equal to $K/2$, and it is well-defined since $H_t(\cdot)$ is convex and increasing on $[0, \infty)$. A caveat is that ordering $\tilde{q}_t$ may not be feasible due to the capacity constraint u in each period t. More specifically, if $\mathbb{E}[H_t^{R/2}(u) \mid f_t] < K/2$, then the quantity $\tilde{q}_t$ exceeds the capacity u and therefore cannot be ordered in full amount. It is natural to consider the order quantity $\min\{\tilde{q}_t, u\}$ which truncates the holding-cost-$K/2$ quantity at u. Thirdly, we compute the conditional expected forced backlogging cost ϕ_t if one orders the minimum of $\tilde{q}_t$ and the capacity u in period t. That is, $\phi_t \triangleq \mathbb{E}[\bar{\Pi}_t^{R/2}(\min\{\tilde{q}_t, u\}) \mid f_t]$. And finally, we compute the conditional expected forced backlogging cost ψ_t resulting from not ordering anything in period t. That is, $\psi_t \triangleq \mathbb{E}[\bar{\Pi}_t^{R/2}(0) \mid f_t]$.

4.2 Description of the $R/2$ policy

Using the quantities computed above, we propose the following procedure for a randomized policy for period t.

(i) If the balancing cost exceeds $K/2$, i.e., $\theta_t \geq K/2$, then the $R/2$ policy orders the balancing quantity $\hat{q}_t$ with probability $p_t = 1$;

(ii) if the balancing cost is less than $K/2$, i.e., $\theta_t < K/2$, then the $R/2$ policy orders the truncated holding-cost-$K/2$ quantity $\min\{\tilde{q}_t, u\}$ with probability p_t and order nothing with probability $1 - p_t$. The probability p_t is computed by the following equation,

$$p_t \frac{K}{2} = p_t \phi_t + (1 - p_t)\psi_t. \tag{7}$$

It follows from (7) that

$$p_t = \frac{\psi_t}{K/2 - \phi_t + \psi_t}.$$

We argue that $0 \leq p_t < 1$. This is because, in (ii) it holds that

$$\mathbb{E}[H_t^{R/2}(\hat{q}_t) \mid f_t] = \theta_t < K/2 = \mathbb{E}[H_t^{R/2}(\tilde{q}_t) \mid f_t],$$

hence we must have $\tilde{q}_t > \hat{q}_t$. In addition, $u \geq \hat{q}_t$ by the construction of $\hat{q}_t$. Thus, $\hat{q}_t \leq \min\{\tilde{q}_t, u\}$, which implies that $\phi_t \leq \theta_t < K/2$.

In summary, we denote the order quantity of the $R/2$ policy by $q_t^{R/2}$. Then the $R/2$ policy orders

$$q_t^{R/2} = \begin{cases} \hat{q}_t, & \text{with probability } p_t = 1 \text{ in case (i)}, \\ \min\{\tilde{q}_t, u\}, & \text{with probability } p_t \text{ in case (ii)}, \\ 0, & \text{with probability } 1 - p_t \text{ in case (ii)}, \end{cases}$$

where p_t in case (ii) is given by (7). The $R/2$ policy is depicted in Figure 1. This concludes the description of the $R/2$ policy.

Note that p_t is a-priori random and is realized with the information set $f_t \in \mathcal{F}_t$. Following the convention we use P_t to denote this a-priori random probability. Similarly, we use $Q_t^{R/2}$ to represent the random a-priori ordering quantity in period t.

4.3 Key ideas of the $R/2$ policy

In the next section, we shall show that the $R/2$ policy described above has an expected worst-case performance guarantee of 4. Here we first provide the intuition and keys ideas underlying this policy.

When the balancing cost θ_t exceeds $K/2$, we have $K \leq 2\theta_t$, implying that the setup cost K is smaller than the the total expected marginal holding and forced backlogging costs in period t. The setup cost in this case is a less dominant factor. Moreover, if the $R/2$ policy does not place an order, the conditional expected forced backlogging cost is potentially very large. Thus, it is worthwhile to order the balancing quantity $q_t^{R/2} = \hat{q}_t$ with probability 1. When the balancing cost θ_t is below $K/2$, the setup cost K becomes more dominant, and therefore it is not advisable to order with probability 1. It is natural to attempt to perfectly balance the three types of the costs, namely, marginal holding, forced backlogging and setup costs. Due to the ordering capacity constraint u, the optimal balancing ratio is no longer $1 : 1 : 1$ for each type of the costs. Intuitively, we want to increase our frequencies of ordering, keeping the sum of the marginal holding and forced backlogging equal to the setup costs. In particular, since we order the truncated holding-cost-$K/2$
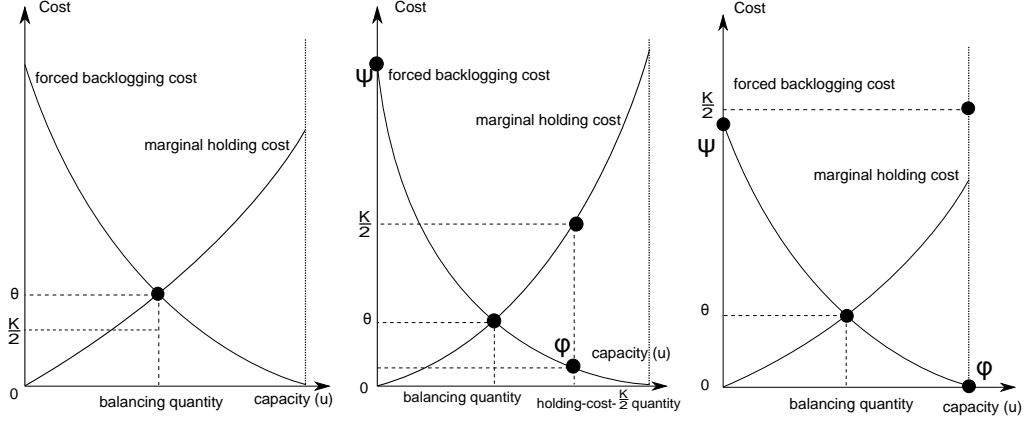
Figure 1: A graphical depiction of how the $R/2$ policy orders in the following three scenarios: (1) when the balancing cost exceeds $K/2$, the policy orders the balancing quantity; (2) when the balancing cost is below $K/2$ and the holding-cost-$K/2$ quantity is below the full capacity, order the holding-cost-$K/2$ quantity with probability p_t and nothing with probability $1 - p_t$; (3) when the balancing cost is below $K/2$ and the holding-cost-$K/2$ quantity exceeds the full capacity, order the full capacity with probability p_t and nothing with probability $1 - p_t$. Note that p_t is computed from equation (7).

quantity $\min\{\tilde{q}_t, u\}$ with probability p_t and nothing with probability $1 - p_t$, the conditional expected marginal holding cost in this case is

$$
\begin{aligned}
\mathbb{E}[H_t^{R/2}(q_t^{R/2}) \mid f_t] &= p_t \cdot \mathbb{E}[H_t^{R/2}(\min\{\tilde{q}_t, u\}) \mid f_t] + (1 - p_t) \cdot \mathbb{E}[H_t^{R/2}(0) \mid f_t] \\
&\leq p_t \cdot \mathbb{E}[H_t^{R/2}(\tilde{q}_t) \mid f_t] + (1 - p_t) \cdot \mathbb{E}[H_t^{R/2}(0) \mid f_t] = p_t K/2.
\end{aligned}
$$

By the construction of the ordering probability p_t in (7), the conditional expected forced backlogging cost is

$$
\begin{aligned}
\mathbb{E}[\bar{\Pi}_t^{R/2}(q_t^{R/2}) \mid f_t] &= p_t \cdot \mathbb{E}[\bar{\Pi}_t^{R/2}(\min\{\tilde{q}_t, u\}) \mid f_t] + (1 - p_t) \cdot \mathbb{E}[\bar{\Pi}_t^{R/2}(0) \mid f_t] \\
&= p_t \phi_t + (1 - p_t)\psi_t = p_t K/2.
\end{aligned}
$$

Finally, since p_t is the ordering probability, the expected setup cost is $p_t K$, which is twice of $p_t K/2$. It follows that this randomized decision rule *almost* balances in a parameterized way, up to the capacity constraint, the three types of costs associated with the period. The balancing ratio is $1 : 1 : 2$ for the marginal holding, the forced backlogging and the setup costs.

Remark In a way, the balancing randomized $R/2$ policy we employed for marginal holding, the forced backlogging and setup costs is optimal in terms of achieving the best worst-case bound. Indeed, we could show that, if the balancing ratio is $a : b : c$, then the worst case bound reaches its minimum at $a : b : c = 1 : 1 : 2$. For example, if our balancing ratio is $1 : 1 : 1$, then we would obtain a worst-case bound of 6.

11

5   Worst-Case Analysis of the R/2 Policy

In this section, we provide a worst-case analysis of the randomized $1/2$-balancing policy ($R/2$) and show that the $R/2$ policy has a provable worst-case performance guarantee of 4. In Section 7, we demonstrate through extensive numerical studies that the $R/2$ policy empirically performs well, and it is significantly better than the provable worst-case performance guarantees. We formally state Theorem 1, which is the main result of this paper.

Theorem 1. *For each instance of the capacitated periodic-review stochastic inventory system with setup cost, the expected cost of the randomized $1/2$-balancing policy ($R/2$) is at most 4 times the expected cost of an optimal policy OPT, i.e.,*

$$\mathbb{E}[\mathscr{C}(R/2)] \leq 4\, \mathbb{E}[\mathscr{C}(OPT)].$$

The proof of Theorem 1 is divided into a sequence of lemmas. First, let $Z_t^{R/2}$ be a random variable defined as

$$Z_t^{R/2} \triangleq \begin{cases} \Theta_t, & \text{if } \Theta_t \geq K/2; \\ P_t \frac{K}{2}, & \text{otherwise,} \end{cases} \tag{8}$$

where $\Theta_t \triangleq \mathbb{E}[H_t^{R/2}(Q_t^{R/2}) \mid F_t] = \mathbb{E}[\bar{\Pi}_t^{R/2}(Q_t^{R/2}) \mid F_t]$ is the balancing cost and P_t is the ordering probability in period t. Note that $Z_t^{R/2}$ and P_t are random variables that are realized with the information set $f_t \in F_t$ in period t. In the following lemma we show that the expected cost of the $R/2$ policy can be upper bounded using the $Z_t^{R/2}$ variables defined in (8).

Lemma 1. *Let $\mathscr{C}(R/2)$ be the total cost incurred by the $R/2$ policy. Then we have,*

$$\mathbb{E}[\mathscr{C}(R/2)] \leq 4 \sum_{t=1}^{T-L} \mathbb{E}[Z_t^{R/2}].$$

Proof. We first show that $Z_t^{R/2} \geq \mathbb{E}[H_t^{R/2}(Q_t^{R/2}) \mid F_t]$, $Z_t^{R/2} = \mathbb{E}[\bar{\Pi}_t^{R/2}(Q_t^{R/2}) \mid F_t]$ and $Z_t^{R/2} \geq P_t K/2$ with probability 1. Given any information set f_t, we know the inventory level x_t and all the quantities $\theta_t, \psi_t, \phi_t, p_t$ defined above are also known deterministically. We split the analysis into two cases.

First, if $\theta_t \geq K/2$, then $q_t^{R/2} = \hat{q}_t$ with probability $P_t = p_t = 1$ implying $z_t^{R/2} = \theta_t \geq K/2$. In addition, we have

$$z_t^{R/2} = \mathbb{E}[H_t^{R/2}(\hat{q}_t) \mid f_t] = \mathbb{E}[\bar{\Pi}_t^{R/2}(\hat{q}_t) \mid f_t],$$

and the claim follows.

Second, if $\theta_t < K/2$, then $q_t^{R/2} = \min\{\tilde{q}_t, u\}$ with probability p_t and $q_t^{R/2} = 0$ with $1 - p_t$. Thus, by the construction of the probability p_t, we have $z_t^{R/2} = p_t K/2 = \mathbb{E}[\bar{\Pi}_t^{R/2}(q_t^{R/2}) \mid f_t]$ and

$$z_t^{R/2} = p_t K/2 = \mathbb{E}[H_t^{R/2}(\tilde{q}_t) \mid f_t] \geq \mathbb{E}[H_t^{R/2}(q_t^{R/2}) \mid f_t],$$

hence the claim again follows.

Applying the above results, we obtain

$$
\begin{aligned}
\mathbb{E}[\mathscr{C}(R/2)] &= \sum_{t=1}^{T-L} \mathbb{E}[H_t^{R/2}(Q_t^{R/2}) + \bar{\Pi}_t^{R/2}(Q_t^{R/2}) + K \cdot \mathbb{1}(Q_t^{R/2} > 0)] \\
&= \sum_{t=1}^{T-L} \mathbb{E}\left[\mathbb{E}[H_t^{R/2}(Q_t^{R/2}) + \bar{\Pi}_t^{R/2}(Q_t^{R/2}) + K \cdot \mathbb{1}(Q_t^{R/2} > 0) \mid F_t]\right] \\
&\leq \sum_{t=1}^{T-L} \mathbb{E}[2Z_t^{R/2} + P_t K] \leq 4 \sum_{t=1}^{T-L} \mathbb{E}[Z_t^{R/2}].
\end{aligned}
$$

This completes the proof of the lemma. $\qquad \square$

To complete the worst-case analysis, we need to show that the expected cost of an optimal policy denoted by OPT is at least $\sum_{t=1}^{T-L} \mathbb{E}[Z_t^{R/2}]$. This will be done by amortizing the cost of OPT against the cost of the $R/2$ policy. In the subsequent analysis, we decompose the set of periods $\{1, 2, \ldots T - L\}$ into the following random partition of six sets:

$$
\mathscr{T}_{1H} = \left\{ t : \Theta_t \geq \frac{K}{2} \text{ and } Y_t^{OPT} \geq Y_t^{R/2} \right\}; \tag{9}
$$

$$
\mathscr{T}_{1\Pi} = \left\{ t : \Theta_t \geq \frac{K}{2} \text{ and } Y_t^{OPT} < Y_t^{R/2} \right\}; \tag{10}
$$

$$
\mathscr{T}_{2H} = \left\{ t : \Theta_t < \frac{K}{2} \text{ and } Y_t^{OPT} \geq X_t^{R/2} + Q_t^{R/2} \text{ and } Q_t^{R/2} = \tilde{Q}_t \leq u \right\}; \tag{11}
$$

$$
\mathscr{T}_{2S} = \left\{ t : \Theta_t < \frac{K}{2} \text{ and } Y_t^{OPT} \geq X_t^{R/2} + Q_t^{R/2} \text{ and } Q_t^{R/2} = u < \tilde{Q}_t \right\}; \tag{12}
$$

$$
\mathscr{T}_{2\Pi} = \left\{ t : \Theta_t < \frac{K}{2} \text{ and } X_t^{R/2} \geq Y_t^{OPT} \right\}; \tag{13}
$$

$$
\mathscr{T}_{2M} = \left\{ t : \Theta_t < \frac{K}{2} \text{ and } X_t^{R/2} < Y_t^{OPT} < X_t^{R/2} + Q_t^{R/2} \text{ and } Q_t^{R/2} = \min\{\tilde{Q}_t, u\} \right\}. \tag{14}
$$

Note that the sets $(9) - (14)$ are disjoint and their union is the complete set of periods. It is also straightforward to check that conditioning on f_t, it is already known which part of the partition period t belongs. We first analyze the sets $\mathscr{T}_{1H}$, $\mathscr{T}_{2H}$, $\mathscr{T}_{1\Pi}$ and $\mathscr{T}_{2\Pi}$ since we can identify the cost components of the optimal policy larger than those in the $R/2$ policy. This gives rise to Lemma 2 below.

Lemma 2. *The total holding and backlogging costs incurred by* OPT, *denoted by* H^{OPT} *and* Π^{OPT} *respectively, satisfy*

$$
\mathbb{E}[H^{OPT}] \geq \mathbb{E}\left[\sum_t Z_t^{R/2} \cdot \mathbb{1}(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H})\right], \tag{15}
$$

$$
\mathbb{E}[\Pi^{OPT}] \geq \mathbb{E}\left[\sum_t Z_t^{R/2} \cdot \mathbb{1}(t \in \mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi})\right]. \tag{16}
$$

Proof. Note that in each period $t \in (\mathscr{T}_{1H} \bigcup \mathscr{T}_{2H})$, we have $Y_t^{OPT} \geq Y_t^{R/2}$ with probability one. By the argument of Lemma 4.2 in [19], since the inventory level of the optimal policy is higher than that of the $R/2$ policy, the optimal policy must have ordered $Q_t^{R/2}$ no later than the $R/2$ policy. Thus, the total holding cost associated with $Q_t^{R/2}$ in the optimal policy must exceed that of the $R/2$ policy. It remains to check that $\mathbb{E}[H_t^{R/2}(Q_t^{R/2}) \mid F_t] = Z_t^{R/2}$ in the two sets $\mathscr{T}_{1H}$ and $\mathscr{T}_{2H}$.

For $t \in \mathscr{T}_{1H}$, since the $R/2$ policy orders the balancing quantity, i.e., $Q_t^{R/2} = \hat{Q}_t$,

$$\mathbb{E}[H_t^{R/2}(Q_t^{R/2}) \mid F_t] = \mathbb{E}[H_t^{R/2}(\hat{Q}_t) \mid F_t] = \Theta_t = Z_t^{R/2}.$$

Now for $t \in \mathscr{T}_{2H}$, since $Q_t^{R/2} = \tilde{Q}_t \leq u$ by the construction of $\mathscr{T}_{2H}$,

$$\mathbb{E}[H_t^{R/2}(Q_t^{R/2}) \mid F_t] = \mathbb{E}[H_t^{R/2}(\min\{\tilde{Q}_t, u\}) \mid F_t] = \mathbb{E}[H_t^{R/2}(\tilde{Q}_t) \mid F_t] = P_t \frac{K}{2} = Z_t^{R/2}.$$

Note that $\tilde{Q}_t$ is the holding-cost-$K/2$ quantity; in the set $\mathscr{T}_{2H}$, the $R/2$ policy can order this quantity in full amount since it is below the capacity u. Thus, we conclude that

$$
\begin{aligned}
\mathbb{E}[H^{OPT}] &\geq \mathbb{E}\left[\sum_t H_t^{R/2}(Q_t^{R/2}) \cdot \mathbb{1}(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H})\right] \\
&= \mathbb{E}\left[\mathbb{E}\left[\sum_t H_t^{R/2}(Q_t^{R/2}) \cdot \mathbb{1}(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H}) \mid F_t\right]\right] \\
&= \mathbb{E}\left[\sum_t Z_t^{R/2} \cdot \mathbb{1}(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H})\right].
\end{aligned}
$$

Note that in each period $t \in \mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi}$, we have $Y_t^{OPT} < Y_t^{R/2}$ with probability one. By the argument of Lemma 2 in [20], we have $\mathbb{E}[\Pi^{OPT}] \geq \mathbb{E}\left[\sum_t \bar{\Pi}_t^{R/2} \mathbb{1}(t \in \mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi})\right]$. Since $\mathbb{E}[\bar{\Pi}_t^{R/2}(Q_t^{R/2}) \mid F_t] = Z_t^{R/2}$ is automatic by the construction of the $R/2$ policy, we have

$$
\begin{aligned}
\mathbb{E}[\Pi^{OPT}] &\geq \mathbb{E}\left[\sum_t \bar{\Pi}_t^{R/2}(Q_t^{R/2}) \cdot \mathbb{1}(t \in \mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi})\right] \\
&= \mathbb{E}\left[\mathbb{E}\left[\sum_t \bar{\Pi}_t^{R/2}(Q_t^{R/2}) \cdot \mathbb{1}(t \in \mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi}) \mid F_t\right]\right] \\
&= \mathbb{E}\left[\sum_t Z_t^{R/2} \cdot \mathbb{1}(t \in \mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi})\right].
\end{aligned}
$$

This completes the proof of Lemma 2. $\qquad\square$

In each period $t \in \mathscr{T}_{2S}$, the $R/2$ policy can no longer order the holding-cost-$K/2$ quantity in full amount due to the capacity constraint u, i.e., $Q_t^{R/2} = u < \tilde{Q}_t$. Then we have

$$\mathbb{E}[H_t^{R/2}(Q_t^{R/2}) \mid F_t] = \mathbb{E}[H_t^{R/2}(u) \mid F_t] < \mathbb{E}[H_t^{R/2}(\tilde{Q}_t) \mid F_t] = P_t \frac{K}{2} = Z_t^{R/2}.$$

Thus, we can no longer argue that the holding cost of the optimal policy is greater than

$$\mathbb{E}\left[\sum_t Z_t^{R/2} \cdot \mathbb{1}(t \in \mathscr{T}_{2S})\right],$$

even though the ending inventory of the optimal policy is higher than that of the $R/2$ policy. However, we show in Lemma 3 that half of the total setup costs incurred by the optimal policy can be used to pay $\mathbb{E}\left[\sum_t Z_t^{R/2} \cdot \mathbb{1}(t \in \mathscr{T}_{2S})\right]$ incurred by the $R/2$ policy.

Lemma 3. *Half of the total setup costs incurred by the optimal policy is lower bounded by*

$$\frac{1}{2}\,\mathbb{E}\left[\sum_t K \cdot \mathbb{1}(Q_t^{OPT} > 0)\right] \;\geq\; \mathbb{E}\left[\sum_t Z_t^{R/2} \cdot \mathbb{1}(t \in \mathscr{T}_{2S})\right]. \tag{17}$$

Proof. In each period $t \in \mathscr{T}_{2S}$, we have $Q_t^{R/2} = u < \tilde{Q}_t$, the $R/2$ policy will order the capacity u with probability p_t incurring strictly less than $p_t K/2$ expected marginal holding cost. Since $Z_t^{R/2} = P_t K/2$ in each period $t \in \mathscr{T}_{2S}$, we shall show that half of the setup costs incurred by the optimal policy is greater than $\mathbb{E}\left[\sum_t P_t \frac{K}{2} \cdot \mathbb{1}(t \in \mathscr{T}_{2S})\right]$.

Fix a period $t \in \mathscr{T}_{2S}$. First we claim that the number of orders placed by the optimal policy over the interval $[1, t]$ is at least the number of orders in which the $R/2$ policy orders the full capacity u over $[1, t]$. We prove the claim by contradiction. Suppose otherwise, the number of orders placed by the optimal policy over the interval $[1, t]$ is m and the number of orders in which the $R/2$ policy orders exactly the capacity u over $[1, t]$ is n and $m < n$. The maximum inventory position of the optimal policy in period t is

$$X_1^{OPT} + mu - \sum_{s=1}^{t-1} D_s \geq Y_t^{OPT},$$

whereas the minimum inventory position of the $R/2$ policy in period t is

$$X_1^{R/2} + nu - \sum_{s=1}^{t-1} D_s \leq Y_t^{R/2} = X_t^{R/2} + u.$$

Since $m < n$ and both policies start with the same inventory position in period 1, i.e., $X_1^{OPT} = X_1^{R/2}$, this implies that $Y_t^{OPT} < X_t^{R/2} + u$ which contradicts to the fact that $t \in \mathscr{T}_{2S}$. The claim thus holds true.

Thus, by letting A be the event that the $R/2$ policy orders exactly the capacity u, we have

$$\begin{aligned}
2\,\mathbb{E}\left[\sum_t Z_t^{R/2} \cdot \mathbb{1}(t \in \mathscr{T}_{2S})\right] &= \mathbb{E}\left[\sum_t P_t K \cdot \mathbb{1}(t \in \mathscr{T}_{2S})\right] \\
&= \mathbb{E}\left[\sum_t K \cdot \mathbb{1}(A \cap t \in \mathscr{T}_{2S})\right] \\
&\leq \mathbb{E}\left[\sum_t K \cdot \mathbb{1}(Q_t^{OPT} > 0)\right].
\end{aligned}$$

15

The last inequality holds true because of our previous claim that the number of orders placed by the optimal policy is not smaller than the total number of *full capacity* orders placed by the $R/2$ policy within $[1, t]$ where $t \in \mathscr{T}_{2S}$. This completes the proof of Lemma 3. $\qquad\square$

In each period $t \in \mathscr{T}_{2M}$, the $R/2$ policy orders the truncated holding-cost-$K/2$ quantity $Q_t^{R/2} = \min\{\tilde{Q}_t, u\}$ with probability p_t and nothing with probability $1 - p_t$. The randomized decision rule introduces uncertainties in the relation between the inventory positions after ordering of the $R/2$ policy and the optimal policy. Thus, we cannot argue that the holding cost or the backlogging cost of the optimal policy is greater than $\mathbb{E}\left[\sum_t Z_t^{R/2} \cdot \mathbb{1}(t \in \mathscr{T}_{2M})\right]$. We resort to the setup costs incurred by the optimal policy again, and show in Lemma 4 that half of the total setup costs incurred by the optimal policy is sufficient to pay $\mathbb{E}\left[\sum_t Z_t^{R/2} \cdot \mathbb{1}(t \in \mathscr{T}_{2S})\right]$ incurred by the $R/2$ policy.

Lemma 4. *Half of the total setup cost incurred by the optimal policy is lower bounded by*

$$\frac{1}{2} \mathbb{E}\left[\sum_t K \cdot \mathbb{1}(Q_t^{OPT} > 0)\right] \;\geq\; \mathbb{E}\left[\sum_t Z_t^{R/2} \cdot \mathbb{1}(t \in \mathscr{T}_{2M})\right]. \tag{18}$$

Proof. Consider an arbitrary sample path with $f_T \in \mathcal{F}_T$. We denote the period in which the optimal policy makes the nth order by t_n. Then we can partition the planning horizon $\{1, \ldots, T\} = [0, t_1) \cup [t_1, t_2) \cdots [t_{N-1}, t_N) \cup [t_N, t_{N+1})$, where $t_{N+1} = T + 1$ and N is the total number of orders that the optimal policy have placed through T.

First we claim that there does not exist a period $s \in [0, t_1)$ such that $s \in \mathscr{T}_{2M}$. Since the $R/2$ policy and optimal policy have the same initial inventory $X_1^{OPT} = X_1^{R/2}$ and face the same demands, if the optimal policy has not placed any orders, we must have $X_s^{R/2} \geq X_s^{OPT} = Y_s^{OPT}$, which implies that s does not belong to the set $\mathscr{T}_{2M}$. Therefore the claim is proved.

Next we claim that the $R/2$ policy will at most make one order in each set of periods $\mathscr{T}_{2M} \cap [t_i, t_{i+1})$ where $1 \leq i \leq N$. In each period $t \in \mathscr{T}_{2M} \cap [t_i, t_{i+1})$, we have $X_t^{R/2} < Y_t^{OPT} < X_t^{R/2} + Q_t^{R/2}$. By the construction of the $R/2$ policy, we will order $Q_t^{R/2}$ with probability p_t and nothing otherwise. Now let A be the event in which the $R/2$ policy places an order and define the stopping time

$$k = \inf\left\{m \geq t_i : A \cap m \in \mathscr{T}_{2M}\right\}.$$

If $k \geq t_{i+1}$, the claim holds since the $R/2$ policy does not place any orders within $\mathscr{T}_{2M} \cap [t_i, t_{i+1})$. Now suppose that $k < t_{i+1}$. It suffices to show that $\mathscr{T}_{2M} \cap [k+1, t_{i+1}) = \emptyset$. Since $k \in \mathscr{T}_{2M}$ and the $R/2$ policy places an order, then we must have $Y_k^{OPT} < X_k^{R/2} + Q_k^{R/2} = Y_k^{R/2}$. In addition, we know that $Y_m^{OPT} = X_m^{OPT}$ for all $m \in (t_i, t_{i+1})$ since the optimal policy does not place any orders in the set (t_i, t_{i+1}). Then for each period $j \in [k+1, t_{i+1})$, by the dynamics of the model, we have

$$Y_j^{OPT} = X_j^{OPT} = Y_k^{OPT} - \sum_{m=k}^{j-1} D_m < Y_k^{R/2} - \sum_{m=k}^{j-1} D_m \leq Y_k^{R/2} + \sum_{m=k}^{j-1} Q_m - \sum_{m=k}^{j-1} D_m = X_j^{R/2},$$

which implies that j does not belong to the set $\mathscr{T}_{2M}$. Thus the second claim is also proved.

16

Then we can conclude with probability 1, it holds that

$$\sum_{t=1}^{T-L} K \cdot \mathbb{1}(A \cap t \in \mathscr{T}_{2M}) \leq NK.$$

Thus, by the above inequality, we have

$$
\begin{aligned}
2\,\mathbb{E}\left[\sum_{t=1}^{T-L} Z_t^{R/2} \cdot \mathbb{1}(t \in \mathscr{T}_{2M})\right] &= \mathbb{E}\left[\sum_{t=1}^{T-L} P_t K \cdot \mathbb{1}(t \in \mathscr{T}_{2M})\right] = \mathbb{E}\left[\sum_{t=1}^{T-L} K \cdot \mathbb{1}(A \cap t \in \mathscr{T}_{2M})\right] \\
&\leq \mathbb{E}\left[NK\right] = \mathbb{E}\left[\sum_{t} K \cdot \mathbb{1}(Q_t^{OPT} > 0)\right].
\end{aligned}
$$

This completes the proof of Lemma 4. $\qquad\square$

Summing up inequalities (15), (16), (18) and (17), we then obtain

$$\mathbb{E}[\mathscr{C}(OPT)] \geq \sum_{t=1}^{T-L} \mathbb{E}[Z_t^{R/2}]. \tag{19}$$

Hence, by (19) and Lemma 1, we have established Theorem 1, i.e., the $R/2$ policy has an expected worst-case performance guarantee of 4.

Before closing this section, we provide some intuitions why our proposed policy has a worst case performance guarantee of 4 but not 3 (in which Levi and Shi [21] were able to prove for the uncapacitated stochastic problems). It can be readily observed that the set $\mathscr{T}_{2S}$ defined in (12) can be merged into the set $\mathscr{T}_{2H}$ defined in (11) in models with infinite ordering capacities. This follows from the fact that the policy can always order up to holding-cost-K quantity $\tilde{Q}_t$. Following the arguments in Lemma 2, we can show that the holding cost incurred by OPT can cover our balancing cost in $\mathscr{T}_{2H} \bigcup \mathscr{T}_{2S}$. Together with the analysis of the remaining partitions, this leads to the 3-approximation algorithm in Levi and Shi [21] for the uncapacitated lot-sizing problem. With ordering capacity constraints, the holding cost incurred by OPT can no longer cover our balancing cost in $\mathscr{T}_{2S}$. Instead, we have shown that the setup costs incurred by OPT can be used to cover this gap. Since analyzing the problematic set $\mathscr{T}_{2M}$ requires the use of setup costs incurred by OPT once again, we have in fact used the setup costs incurred by OPT twice. If our balancing ratio is $1:1:1$ (marginal holding, the forced backlogging and setup costs), then we would obtain a worst-case bound of 6, which is not optimal in terms of achieving the tightest worst-case bound. As we discussed earlier, the worst case bound reaches its minimum at $1:1:2$, which yields a 4-approximation algorithm.

6 Extensions to Batch Ordering System

In this section we extend our results to stochastic periodic-review inventory models with setup cost under *batch ordering constraints*. The batch constraint specifies that, every order quantity has to be an integer multiple of a pre-specified base batch size, say q_0, which can be, for example, a box, a pallet, a truckload, etc. The case of batch ordering with infinite order capacity has been studied in

the literature, see, for example Veinott [25], and Chen [4], among others. Without loss of generality, we can assume the capacity is also an integer multiple of batch size, since the excess quantity which is less than a base batch size q_0 can never be used under batch order constraints. More specifically, let $u = mq_0$ where m is a given positive integer, and a feasible policy can only order quantities of iq_0 for some integer i taking value from $\{0, 1, \ldots, m\}$.

6.1   Modified randomized 1/2-balancing ($MR/2$) policy

Since the ordering quantity can only be an integer multiple of the base quantity q_0, the marginal holding cost function $H_t(Q_t)$ and the forced backlogging cost function $\bar{\Pi}_t(Q_t)$ are defined only at $Q_t = iq_0$ where $i = 0, 1, \ldots, m$. For other non-negative integer value Q_t, we can extend the two functions $H_t(Q_t)$ and $\bar{\Pi}_t(Q_t)$ by interpolating piecewise linear extensions of these batch quantities. More specifically, for any integer value Q_t, there exists a scalar $\lambda_t \in [0, 1)$ such that

$$Q_t = (1 - \lambda_t)Q_t^{lower} + \lambda_t Q_t^{upper},$$

where

$$Q_t^{lower} = \lfloor Q_t/q_0 \rfloor q_0, \qquad Q_t^{upper} = \lfloor Q_t/q_0 + 1 \rfloor q_0,$$

and the floor function $\lfloor a \rfloor$ is the largest integer less than or equal to a. The corresponding marginal holding cost and forced backlogging cost are defined, using linear interpolation, as

$$\begin{aligned}
\mathbb{E}[H_t(Q_t) \mid F_t] &= (1 - \lambda_t) \cdot \mathbb{E}[H_t(Q_t^{lower}) \mid F_t] + \lambda_t \cdot \mathbb{E}H_t(Q_t^{upper}) \mid F_t], \\
\mathbb{E}[\bar{\Pi}_t(Q_t) \mid F_t] &= (1 - \lambda_t) \cdot \mathbb{E}[\bar{\Pi}_t(Q_t^{lower}) \mid F_t] + \lambda_t \cdot \mathbb{E}\bar{\Pi}_t(Q_t^{upper}) \mid F_t].
\end{aligned}$$

It is clear that these extended cost functions $H_t(Q_t)$ and $\bar{\Pi}_t(Q_t)$ preserve the properties of convexity and monotonicity.
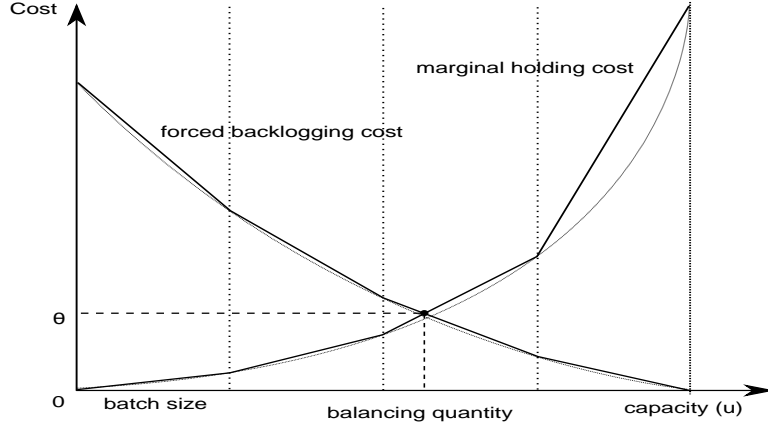

Figure 2: A graphical depiction of the extended cost functions by linear interpolation.

We now propose a *modified randomized 1/2-balancing policy* ($MR/2$). At the beginning of each period t with the realized information set f_t, we compute the auxiliary order quantities and cost functions discussed in Section 4.1. Note that these auxiliary functions are defined only on integer

18

multiples of the base batch size q_0. Thus, to properly balance the cost functions, we need to define their corresponding lower and upper quantities.

First, compute the balancing quantity $\hat{q}_t$ and the balancing cost θ_t such that $\theta_t \triangleq \mathbb{E}[\bar{\Pi}_t(\hat{q}_t) \mid f_t] = \mathbb{E}[\bar{\Pi}_t(\hat{q}_t) \mid f_t]$. Then, there exists a scalar $\hat{\lambda}_t \in [0, 1)$ such that

$$\hat{q}_t = (1 - \hat{\lambda}_t)\hat{q}_t^{lower} + \hat{\lambda}_t \hat{q}_t^{upper}$$

where $\hat{q}_t^{lower} = \lfloor \hat{q}_t/q_0 \rfloor q_0$ and $\hat{q}_t^{upper} = \lfloor \hat{q}_t/q_0 + 1 \rfloor q_0$.

Next, compute the holding-cost-$K/2$ quantity $\tilde{q}_t$ such that $\mathbb{E}[H_t(\tilde{q}_t) \mid f_t] = K/2$. There exists another scalar $\tilde{\lambda}_t \in [0, 1)$ such that

$$\tilde{q}_t = (1 - \tilde{\lambda}_t)\tilde{q}_t^{lower} + \tilde{\lambda}_t \tilde{q}_t^{upper},$$

where $\tilde{q}_t^{lower} = \lfloor \tilde{q}_t/q_0 \rfloor q_0$ and $\tilde{q}_t^{upper} = \lfloor \tilde{q}_t/q_0 + 1 \rfloor q_0$.

Third, compute the resulting forced backlogging cost $\phi_t \triangleq \mathbb{E}[\bar{\Pi}_t(\min\{\hat{q}_t, u\}) \mid f_t]$ if one orders the minimum of the holding-cost-$K/2$ quantity and the capacity u.

Finally, compute the forced backlogging cost $\psi_t \triangleq \mathbb{E}[\bar{\Pi}_t(0) \mid f_t]$ if one orders nothing.

The modified randomized 1/2-balancing (MR/2) order policy we propose for the case with batch ordering constraint is described as follows:

(i) If the balancing cost $\theta_t \geq K/2$, then the $MR/2$ policy orders $\hat{q}_t^{lower}$ with probability $1 - \hat{\lambda}_t$ and $\hat{q}_t^{upper}$ with probability $\hat{\lambda}_t$.

(ii) If the balancing cost $\theta_t < K/2$, then compute the ordering probability p_t from $p_t K/2 = p_t \phi_t + (1 - p_t)\psi_t$ similar to (7). The $MR/2$ policy orders $\min\{\tilde{q}_t^{lower}, u\}$ with probability $p_t(1 - \tilde{\lambda}_t)$, order $\min\{\tilde{q}_t^{upper}, u\}$ with probability $p_t \tilde{\lambda}_t$, and order nothing with probability $1 - p_t$.

To summarize, if we denote the order quantity of the $MR/2$ policy by $q_t^{MR/2}$, then it is given by

$$q_t^{MR/2} = \begin{cases} \hat{q}_t^{lower}, & \text{with probability } 1 - \hat{\lambda}_t \text{ in case (i);} \\ \hat{q}_t^{upper}, & \text{with probability } \hat{\lambda}_t \text{ in case (ii);} \\ \min\{\tilde{q}_t^{lower}, u\}, & \text{with probability } p_t(1 - \tilde{\lambda}_t) \text{ in case (ii);} \\ \min\{\tilde{q}_t^{upper}, u\}, & \text{with probability } p_t \tilde{\lambda}_t \text{ in case (ii);} \\ 0, & \text{with probability } 1 - p_t \text{ in case (ii).} \end{cases}$$

It is clear that the modified randomized 1/2-balancing policy balances the three types of costs in a similar manner as the original randomized 1/2-balancing policy without the batch order constraints.

6.2 Worst-case analysis

To conduct the performance analysis, we define all the sets similarly to (9) to (14) as follows.

$$
\mathscr{T}_{1H} = \left\{ t : \Theta_t \geq \frac{K}{2} \text{ and } Y_t^{OPT} \geq X_t^{MR/2} + \hat{q}_t^{upper} \right\}, \tag{20}
$$

$$
\mathscr{T}_{1\Pi} = \left\{ t : \Theta_t \geq \frac{K}{2} \text{ and } Y_t^{OPT} \leq X_t^{MR/2} + \hat{q}_t^{lower} \text{ and } \hat{q}_t^{lower} < \hat{q}_t^{upper} \right\}, \tag{21}
$$

$$
\mathscr{T}_{2H} = \left\{ t : \Theta_t < \frac{K}{2} \text{ and } Y_t^{OPT} \geq X_t^{MR/2} + \min\{\tilde{q}_t^{upper}, u\} \text{ and } \tilde{q}_t^{upper} \leq u \right\}, \tag{22}
$$

$$
\mathscr{T}_{2S} = \left\{ t : \Theta_t < \frac{K}{2} \text{ and } Y_t^{OPT} \geq X_t^{MR/2} + \min\{\tilde{q}_t^{lower}, u\} \text{ and } u \leq \tilde{q}_t^{lower} < \tilde{q}_t^{upper} \right\}, \tag{23}
$$

$$
\mathscr{T}_{2\Pi} = \left\{ t : \Theta_t < \frac{K}{2} \text{ and } X_t^{MR/2} \geq Y_t^{OPT} \right\}, \tag{24}
$$

$$
\mathscr{T}_{2M} = \left\{ t : \Theta_t < \frac{K}{2} \text{ and } X_t^{MR/2} < Y_t^{OPT} \leq X_t^{MR/2} + \min\{\tilde{q}_t^{lower}, u\} \right\}. \tag{25}
$$

To ensure that the union of the sets from (20) to (25) is the complete set of all periods, we show by Lemma 5 that, conditional on the same demand realization, the base batch load q_0 must divide the absolute difference between inventory levels of the optimal policy and the $MR/2$ policy. Thus, it is impossible to have

$$
Y_t^{OPT} \in (X_t^{MR/2} + \min\{\tilde{q}_t^{lower}, u\}, X_t^{MR/2} + \min\{\tilde{q}_t^{upper}, u\})
$$

when the balancing cost $\Theta_t \geq K/2$. Similarly, since u is an integer multiple of q_0, it is also impossible for this to happen when the balancing cost $\Theta_t < K/2$.

Lemma 5. *For any realization, the base batch size q_0 must divide $Y_t^{OPT} - Y_t^{MR/2}$, the difference between the inventory positions of the optimal policy and the $MR/2$ policy for each period $t = 1, \ldots, T - L$.*

Proof. Consider an arbitrary period s and suppose the optimal policy ordered $m_s^{OPT} q_0$ while the $MR/2$ policy ordered $m_s^{MR/2} q_0$, where m_s^{OPT} and $m_s^{R/2}$ are nonnegative integers. Suppose that the starting inventory positions at the beginning of period 1 are the same for both policies, i.e., $X_1^{OPT} = X_1^{MR/2}$. Then for an arbitrary period $t = 1, \ldots, T - L$, we have

$$
\begin{aligned}
&Y_t^{OPT} - Y_t^{MR/2} \\
&= \left(X_1^{OPT} + \left(\sum_{s=1}^{t} m_s^{OPT} \right) q_0 - D_{[1,t)} \right) - \left(X_1^{MR/2} + \left(\sum_{s=1}^{t} m_s^{MR/2} \right) q_0 - D_{[1,t)} \right) \\
&= \left(\sum_{s=1}^{t} (m_s^{OPT} - m_s^{MR/2}) \right) q_0.
\end{aligned}
$$

Thus, the base batch load q_0 must divide $Y_t^{OPT} - Y_t^{MR/2}$. $\square$

By Lemma 5, we have constructed the disjoint sets (20) - (25) and their union is a complete set. It can be readily verified that the Lemmas 1, 2, 3 and 4 continue to hold if we replace all the sets (9) – (14) with (20) – (25). We formally state the result for capacitated stochastic inventory problem with setup cost under batch order constraints.

Theorem 2. *For each instance of the capacitated stochastic periodic-review inventory problem with setup cost under batch ordering constraints, the expected cost of the modified randomized 1/2-balancing policy (MR/2) is at most four times the expected cost of an optimal policy OPT, i.e.,*

$$\mathbb{E}[\mathscr{C}(MR/2)] \leq 4\mathbb{E}[\mathscr{C}(OPT)].$$

Remark Consider a special case of the batch ordering system with the base batch order size equal to the capacity, i.e., we restrict ourselves to *all-or-nothing ordering policies*. (Özer and Wei [23] established the optimality of a state-dependent threshold policy with this class of policies.) In this special case, we can conveniently transform the original unit ordering cost $c_t q_t$ plus the setup cost $K \cdot \mathbb{1}(q_t = u)$ into an equivalent modified unit ordering cost $c'_t q_t$ where $c'_t = c_t + K/u$, since q_t can only take values 0 or u. Then the model will be reduced to the one studied in Levi et al. [20] where a *dual-balancing* policy yields a 2-approximation. It should be noted that this simple transformation fails to work for any more-than-two point ordering policies, since the setup cost exhibits a concave ordering cost structure.

7   Numerical Experiments

In this section, we conduct a numerical study on the performance of the R/2 policy developed in Section 5. As noted by Levi and Shi [21], the randomized cost-balancing policy can be parameterized to obtain general classes of policies, respectively, and the worst-case analysis discussed above can then be viewed as choosing parameter values that perform well against any possible instance. In contrast, one can try to find the 'best' parameter values, for each given instance. This gives rise to policies that have at least the same worst-case performance guarantees, but are likely to work better empirically, since we refined the parameters according to the specific instance being solved. Using simulation based optimization, we have implemented this approach and tested the empirical performance of the resulting policies. The policies were tested using the demand model of *advance demand information* proposed by Gallego and Özer [9], and Özer and Wei [23]. To the best of our knowledge, these are the only papers that reported optimal computational costs (by brute force dynamic programming) for the capacitated stochastic periodic-review inventory system with setup cost and dependent demand structures.

Parameterized policies.   We describe a class of parameterized policies involving parameters β, γ and η, where β controls the holding-cost-$\beta K/2$ quantity, γ controls the ratio of the marginal holding cost to the forced backlogging cost and η controls the level of the forced backlogging cost resulting from not ordering. Specifically, the parameterized policy first computes several quantities.

1) The balancing quantity $\hat{q}_t$ that solves $\mathbb{E}[H_t^{R/2}(\hat{q}_t) \mid f_t] = \gamma \mathbb{E}[\bar{\Pi}_t^{R/2}(\hat{q}_t) \mid f_t] := \theta_t$.

2) The holding-cost-$\beta\frac{K}{2}$ quantity $\tilde{q}_t$ that solves $\mathbb{E}[H_t^{R/2}(\tilde{q}_t) \mid f_t] = \beta\frac{K}{2}$.

3) The resulting conditional expected forced backlogging cost if one orders $\min\{\tilde{q}_t, u\}$ units in period t, denoted by ϕ_t. That is, $\phi_t = \mathbb{E}[\bar{\Pi}_t^{R/2}(\min\{\tilde{q}_t, u\}) \mid f_t]$.

4) The conditional expected forced backlogging cost resulting from not ordering in period t, denoted by ψ_t. That is, $\psi_t = \eta\mathbb{E}[\bar{\Pi}_t^{R/2}(0) \mid f_t]$.

Based on the above quantities computed, the following randomized rule is employed to obtain the ordering quantity for each period t.

a) If $\theta_t \geq \beta K/2$, the $R/2$ policy orders $q_t^{R/2} = \hat{q}_t$ with probability $p_t = 1$ in period t.

b) If $\theta_t < \beta K/2$, the $R/2$ policy orders $q_t^{R/2} = \min\{\tilde{q}_t, u\}$ with probability p_t and order nothing with probability $1 - p_t$ in period t, where p_t solves $0 \leq p_t = \psi_t/(\beta\frac{K}{2} - \phi_t + \psi_t) < 1$.

If we denote the order quantity of the $R/2$ policy by $q_t^{R/2}$, then the $R/2$ policy orders

$$
q_t^{R/2} = \begin{cases} \hat{q}_t, & \text{with probability } p_t = 1 \text{ in case (a)}, \\ \min\{\tilde{q}_t, u\}, & \text{with probability } p_t \text{ in case (b)}, \\ 0, & \text{with probability } 1 - p_t \text{ in case (b)}. \end{cases}
$$

End-of-horizon rule. To prevent the policy from over-ordering too much near the end of horizon, we also incorporate the following end-of-horizon rule. In period t, we estimate the total expected cumulative backlogging cost (assuming no orders are placed) over the interval $[t, T]$. If the amount is less than K, the policy does not place an order in period t.

Algorithmic complexity. We describe the procedures of finding the optimal parameters for a specific instance of the problem. First, assume that there exists a positive constant U such that the optimal parameters β^*, γ^*, η^* are upper bounded by U. In addition, we discretize U with some step-size Δ, i.e., $\beta, \gamma, \eta \in [0, U]$ can only take values as integer multiples of Δ. Then we conduct an exhaustive search on a cube of $U \times U \times U$ for the parameters β, γ and η. In our numerical studies, $U = 100$ and $\Delta = 1$ are chosen to be the upper bound and the resolution for discretization, respectively. The algorithm runs on every point on this cube, simulates the cost of each parameterized policy and returns the best possible $(\beta^*, \gamma^*, \eta^*)$ that minimize the cost. Secondly, assume that there exists a positive constant $\hat{U}$ that serves as an upper bound on the balancing and hold-cost-$K/2$ quantities. For each $t = 1, \ldots, T$, the complexity for evaluating marginal holding cost is $O(T)$ and the complexity for carrying out bisection search is $O(\log \hat{U})$. The algorithm runs in time $O(T^2 \log \hat{U})$ for each set of parameters (β, γ, η). Hence, the algorithm that returns both the optimal parameters and the lowest cost runs in $O(U^3 \Delta^{-3} T^2 \log \hat{U}) \approx O(T^2)$ since $U^3 \Delta^{-3} \log \hat{U}$ is some positive constant. For all tested instances with $T = 10$, the average CPU time per test instance on a Pentium 3.0GHz PC is 170s. In contrast, the dynamic programming algorithm takes 1800s on average per test instance.

Design of experiments and numerical results. When the demand process is correlated over time, the computation of exact optimal solution becomes impossible for reasonable problem sizes.

Thus, in order to compare with the optimal minimum cost, we consider a planning horizon of 10 periods. The cost parameters selected for each demand class are as follows: we normalize $h = 1$ and then vary other parameters, $c \in \{0, 2, 5\}$, $b \in \{5, 10, 15\}$, $K \in \{10, 50\}$ and $u \in \{3, 6, 9\}$.

For the demand process with *advance demand information* (ADI), we adopt a model studied in Gallego and Özer [9] and Özer and Wei [23]. That is, we assume that customers could place orders 2 periods ahead. Thus for each period t, a demand vector $(D_{t,t}, D_{t,t+1}, D_{t,t+2})$ is received, where $D_{t,s}$ is the order placed in period t for period $s \geq t$. The total demand for period t is $D_t = D_{t-2,t} + D_{t-1,t} + D_{t,t}$. We tested the cases for which each entry $D_{t,t+i}$ follows a Poisson distribution with mean λ_i. Note that the actual demands over periods are correlated due to the presence of advance demand information. The performance error of an approximation policy P is defined by

$$err = \left(\frac{\mathscr{C}(P)}{\mathscr{C}(OPT)} - 1 \right) \times 100.$$

In words, the performance error of an approximation policy is the percentage of total cost increase of this policy over the planning horizon with respect to the optimal minimum total cost.

To report all the numerical results for the ADI demand model, we group the instances as follows: The purchasing cost are L ($c = 0$), M ($c = 5$), and H ($c = 10$); the setup costs are L ($K = 10$), and H ($K = 50$); the capacities are L ($u = 3$), M ($u = 6$), and H ($u = 9$). For each triplet (c, K, u), three values of shortage cost $b \in \{5, 10, 15\}$ are used in our tests.

The performance errors for the $R/2$ policy are reported in Table 1. As observed, the numerical results show that the $R/2$ policies perform on average 9% of the error from the optimal cost, which is significantly better than the theoretical worst-case performance guarantees.

Table 1: Performance of the $R/2$ policy for ADI demand structures (err).
$h = 1$, $c \in \{0, 2, 5\}$, $b \in \{5, 10, 15\}$, $K \in \{10, 50\}$, $u \in \{3, 6, 9\}$ and $T = 10$.

	c	L				M				H				All
	K	L		H		L		H		L		H		mean
demand	u	mean	max	mean	max	mean	max	mean	max	mean	max	mean	max	**mean**
	L	10.63	15.10	15.61	20.12	8.75	12.60	14.73	18.53	8.10	11.80	13.20	16.38	11.84
(6,0,0)	M	10.40	15.01	15.40	19.94	8.49	12.00	14.38	18.46	7.90	11.68	13.15	16.12	11.62
	H	10.10	14.95	14.96	19.69	8.10	12.10	14.20	18.31	7.78	11.50	13.08	16.10	11.37
	L	8.70	12.55	13.76	18.60	7.32	11.99	13.28	17.46	7.10	10.31	12.76	16.01	10.49
(3,3,0)	M	8.50	12.41	13.60	18.22	7.30	11.52	13.13	17.19	6.88	10.21	12.58	15.92	10.33
	H	8.44	12.34	13.40	18.13	7.21	11.40	13.07	17.11	6.80	10.13	12.50	15.90	10.23
	L	8.33	12.00	13.23	18.04	6.51	10.68	12.42	16.56	6.22	9.56	10.98	14.27	9.62
(0,6,0)	M	7.60	11.60	13.11	17.98	6.33	10.34	12.32	16.46	6.12	9.21	10.77	14.17	9.38
	H	7.54	11.48	13.01	17.50	6.30	10.15	12.28	16.42	6.03	9.08	10.67	14.12	9.31
	L	6.70	11.08	11.58	15.89	5.45	9.98	11.40	15.64	4.33	6.55	8.23	11.50	7.95
(0,3,3)	M	6.22	9.69	11.40	15.80	5.30	9.50	11.28	15.46	4.17	6.32	8.15	11.28	7.75
	H	6.10	9.51	11.19	15.61	5.19	9.40	11.19	15.33	4.12	6.21	8.08	11.12	7.65
	L	5.45	7.85	10.30	14.60	4.83	6.85	9.19	13.38	2.88	3.54	6.88	9.58	6.59
(0,0,6)	M	5.35	7.69	10.28	14.51	4.45	6.55	8.88	13.10	2.78	3.21	6.78	9.45	6.42
	H	5.27	7.49	10.10	14.29	4.31	6.50	8.37	13.00	2.70	3.18	6.70	9.43	6.24

Acknowledgment

References

[1] A. Atali and Ö. Özer, *Stochastic multi-item inventory systems with markov-modulated demands and production quantity requirements*, Probability in the Engineering and Information Sciences **26** (2012), 263–293.

[2] S. Bollapragada and T. E. Morton, *A simple heuristic for computing nonstationary (s,S) policies*, Operations Research **47** (1999), 576–584.

[3] X. Chao and S. X. Zhou, *Optimal policy for a multiechelon inventory system with batch ordering and fixed replenishment intervals*, Operations Research **57** (2009), no. 2, 377.

[4] F. Chen, *Optimal policies for multi-echelon inventory problems with batch ordering*, Operations Research **48** (2000), 376–389.

[5] S. Chen, *The infinite horizon periodic review problem with setup costs and capacity constraints: A partial characterization of the optimal policy*, Operational Research **52** (2004), no. 3, 409–421.

[6] S. Chen and M. Lambrecht, *X-Y band and modified (s, S) policy*, Operations Research **44** (1996), 1013–1019.

[7] F. Cheng and S. Sethi, *Optimality of (s,S) policies in inventory models with markovian demand*, Operational Research **45** (1997), no. 6, 931–939.

[8] A. Federgruen and P. H. Zipkin, *An efficient algorithm for computing optimal (s, S) policies*, Operations Research **32** (1984), no. 6, 1268–1285.

[9] G. Gallego and Ö. Özer, *Integrating replenishment decisions with advance demand information*, Management Science **47** (2001), 1344–1360.

[10] G. Gallego and A. Scheller-Wolf, *Capacitated inventory problems with fixed order costs: Some optimal policy structure*, European Journal of Operational Research **126** (2000), 603–613.

[11] S. Gavirneni, *An efficient heuristic for inventory control when the customer is using a (s, S) policy*, Operations Research Letters **28** (2001), 187–192.

[12] Y. Guan and A. J. Miller, *A polynomial time algorithm for the stochastic uncapacitated lot-sizing problem with backlogging*, IPCO 2008, volume 5035 of Lecture Notes in Computer Science (2008), 450–462.

[13] ———, *Polynomial time algorithms for stochastic uncapacitated lot-sizing problems*, Operations Research **56** (2008), no. 5, 1172–1183.

[14] D. C. Heath and P. L. Jackson, *Modeling the evolution of demand forecasts with application to safety stock analysis in production/distribution system*, IIE Transactions **26** (1994), no. 3, 17–30.

[15] K. Huang and S. Küçükyavuz, *On stochastic lot-sizing problems with random lead times*, Operations Research Letters **36** (2008), no. 3, 303–308.

[16] W. T. Huh and G. Janakiraman, *Technical note - on optimal policies for inventory systems with batch ordering*, Operations Research **60** (2012), no. 4, 797–802.

[17] E. Ignall and A. F. Veinott, *Optimality of myopic inventory policies for several substitute products*, Management Science **15** (1969), 284–304.

[18] T. Iida and P. H. Zipkin, *Approximate solutions of a dynamic forecast-inventory model*, Manufacturing Service Operations Management **8** (2006), 407–425.

[19] R. Levi, M. Pál, R. O. Roundy, and D. B. Shmoys, *Approximation algorithms for stochastic inventory control models*, Mathematics of Operations Research **32** (2007), 284–302.

[20] R. Levi, R. O. Roundy, D. B. Shmoys, and V. A. Truong, *Approximation algorithms for capacitated stochastic inventory models*, Operations Research **56** (2008), 1184–1199.

[21] R. Levi and C. Shi, *Approximation algorithms for the stochastic lot-sizing problem with order lead times*, Operations Research **61** (2013), no. 3, 593–602.

[22] X. Lu, J. S. Song, and A. C. Regan, *Inventory planning with forecast updates: approximate solutions and cost error bounds*, Operations Research **54** (2006), no. 6, 1079–1097.

[23] Ö. Özer and W. Wei, *Inventory control with limited capacity and advance demand information*, Operations Research **52** (2004), 988–1000.

[24] H. Scarf, *The optimality of (S,s) policies in the dynamic inventory problem*, Mathematical Methods in Social Sciences, Stanford Univ. Press, Stanford, CA, 1960.

[25] A. F. Veinott, *The optimal inventory policy for batch ordering*, Operations Research **13** (1965), 424–432.

[26] ———, *Optimal policy for a multi-product, dynamic, non-stationary inventory problem*, Management Science **12** (1965), 206–222.

[27] ———, *On the optimality of (s,S) inventory policies: new conditions and a new proof*, SIAM J. Appl. Math **14** (1966), 1067–1083.