



MIT Open Access Articles

Approximation Algorithms for the Stochastic Lot-Sizing Problem with Order Lead Times

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation	Levi, Retsef, and Cong Shi. "Approximation Algorithms for the Stochastic Lot-Sizing Problem with Order Lead Times." <i>Operations Research</i> 61, no. 3 (June 2013): 593–602.
As Published	http://dx.doi.org/10.1287/opre.2013.1162
Publisher	Institute for Operations Research and the Management Sciences (INFORMS)
Version	Author's final manuscript
Citable link	http://hdl.handle.net/1721.1/87771
Terms of Use	Creative Commons Attribution-Noncommercial-Share Alike
Detailed Terms	http://creativecommons.org/licenses/by-nc-sa/4.0/

Approximation Algorithms for the Stochastic Lot-sizing Problem with Order Lead Times

Retsef Levi

Sloan School of Management, Massachusetts Institute of Technology, retsef@mit.edu

Cong Shi

Industrial & Operations Engineering, University of Michigan, shicong@umich.edu

We develop new algorithmic approaches to compute provably near-optimal policies for multi-period stochastic lot-sizing inventory models with positive lead times, general demand distributions and dynamic forecast updates. The policies that are developed have worst-case performance guarantees of 3 and typically perform very close to optimal in extensive computational experiments. The newly proposed algorithms employ a novel randomized decision rule and depart from the previous work of Levi et al. (2007). We believe that these new algorithmic and performance analysis techniques could be used in designing provably near-optimal randomized algorithms for other stochastic inventory control models, and more generally in other multistage stochastic control problems.

Key words: inventory, stochastic lot-sizing, approximation algorithms, randomized cost-balancing

History: Submitted on June 10, 2009, Revised June 7, 2011, July 4, 2012, December 29, 2012

1. Introduction

In this paper, we develop new provably near-optimal algorithms for stochastic inventory control models with *fixed costs*, general demand distributions and dynamic forecast updates. Fixed costs arise in many real-life scenarios, and reflect the fact that ordering, production and transportation in large quantities lead to economies of scales. Specifically, we study several general variants of the classical *stochastic lot-sizing problem*. Finding optimal policies in these settings is often computationally intractable. Instead, we develop new algorithmic approaches that yield a *3-approximation*, i.e., they have a *worst-case performance guarantee of 3*. This implies that the algorithms are guaranteed to have expected cost at most three times the optimal expected cost, regardless of the input instance.

Our contributions. The new algorithmic and performance analysis approaches that are developed in this paper depart from the previous work of Levi et al. (2007), and provide multi-fold contributions to the study of stochastic inventory control as well as more generally to the design and analysis of randomized algorithms. The paper extends the recent stream of work to develop *cost-balancing* algorithmic techniques for computationally challenging multi-period stochastic inventory control problems. This stream of work has been initiated by Levi et al. (2007) and subsequent work (Levi et al. (2005, 2008a, 2007, 2008c)), which primarily studied stochastic inventory control problems with *no* fixed costs. The conceptual idea underlying cost-balancing based algorithms is a repeated attempt to balance opposing costs, for example, in models without fixed ordering cost one seeks to balance the cost of over-ordering (*holding cost*) and the cost of under-ordering (*backlogging cost*) based on the notion of *marginal cost accounting schemes* (Levi et al. (2005, 2007, 2008c)) (see also the discussion in Section 4.1).

The existence of fixed costs adds a third *nonlinear* component to the cost, and makes the cost balancing more subtle. Levi et al. (2007) did study a very special case of the model studied in this paper, in which orders arrive instantaneously and demand in each period is known deterministically at the beginning the period *before* the ordering decision is made. They proposed the *triple-balancing policy* that aims to balance the fixed ordering cost, the holding cost and the backlogging cost over each *time interval* between consecutive orders. Their policy is a 3-approximation. However, the algorithm and the worst-case analysis can be applied effectively only to models, in which there is no lag, commonly called *lead time*, from when an order is placed until it arrives. In fact, in models with positive lead times the assumption in Levi et al. (2007) is equivalent to knowing deterministically the cumulative demand over the lead time. This is clearly a *very* restrictive assumption, since in many scenarios forecasting the demand over the lead time is the major challenge. Moreover, in Section 3, we show that if this assumption does not hold, the triple-balancing policy can perform arbitrarily worse than an optimal policy. This stands in contrast to most of the analytical work done on inventory models with backlogged demand, for which the extensions from models with no lead time to models with positive lead time are often immediate.

To address the nonlinearity induced by the fixed costs, a novel *randomized decision rule* is employed to balance the expected fixed ordering costs, holding costs and backlogging costs, in each *period*. In particular, the order quantity in each period is decided based on a carefully designed randomized rule that chooses among various possible order quantities with carefully chosen probabilities. To the best of our knowledge, this is the first randomized policy proposed for stochastic inventory control policies. Levi et al. (2007) used a straightforward randomized rule for the model with no fixed costs, but merely as a ‘rounding’ technique to address the constraint to order in integer quantities. Unlike the triple-balancing policy that balances the costs over intervals, the newly randomized policy balances the costs in each period. Like the triple-balancing policy, the randomized cost-balancing policy proposed in this paper has a worst-case guarantee of 3, but this holds under very general assumptions, i.e., general demand distributions and positive lead times . The worst-case performance analysis of the randomized policy employs several fundamental new ideas that depart from the previous work of Levi et al. (2007). Like the previous work, the analysis is based on an *amortization* of the cost incurred by the balancing policy against the cost of an optimal policy. However, all of the previous work is entirely based on *sample-path* arguments. In contrast, the analysis in this paper is based on more subtle *averaging* arguments. We believe that the new algorithmic and analysis techniques developed in this paper will turn out to be effective in the design of provably near-optimal algorithms for other stochastic inventory control problems.

Our proposed randomized policies can be parameterized to create a broader class of policies. A simulation based optimization is used to find the ‘best’ parameters for a given instance of the problem. This preserves the same worst-case guarantees. Moreover, relatively extensive computational experiments that we conducted indicate that it typically leads to near-optimal policies that perform empirically within few percentages of optimal, significantly better than the worst-case performance guarantees.

In addition, the work in this paper contributes to the body of work on *randomized algorithms*. The last two decades have witnessed a tremendous growth in the area of randomized algorithms. During this period, randomized algorithms went from being a tool in computational number theory

to finding widespread applications in other fields, such as data structures, geometric algorithms, graph algorithms, number theory, enumeration, parallel algorithms, approximation algorithms and online algorithms. Part of the reason why randomized algorithms are attractive is the fact that they are usually conceptually simple and computationally fast. Randomized decision rules have been used extensively to obtain approximation algorithms with worst-case guarantees for many deterministic NP-hard optimization problems, including several examples of deterministic inventory management problems (see for example, Teo and Bertsimas (1996), Levi et al. (2008b)). In addition, randomized decision rules are very common in the field of *online algorithms* (see Borodin and El-Yaniv (1998)), in which they are used to obtain algorithms with *competitive ratios*. However, in spite of the increasing use of randomized algorithms, there have been relatively few successful attempts to incorporate randomized decision rules to obtain algorithms for multistage stochastic control problems. Rust (1997) proposed random versions of successive approximations and multi-grid algorithms for computing approximate solutions to Markovian decision problems. Prandini et al. (1999) designed a randomized algorithm to obtain an estimate of the probability of aircraft conflict. Bouchard et al. (2005) studied a maturity randomization technique for approximating optimal control problems to price American put options. Shmoys and Talwar (2008) proposed a randomized 4-approximation algorithm of the *a priori Traveling Salesman Problem*. Shmoys and Swamy (2006) gave a fully polynomial randomized approximation scheme for solving 2-stage stochastic integer optimization problems. However, the techniques developed in this paper are different and we believe they have a promising potential to apply in other multistage stochastic optimization models.

Literature review. The dominant paradigm in most of the existing literature has been to formulate stochastic inventory control problems (including the models studied in this paper) using a dynamic programming framework. This approach turned out to be effective in characterizing the structure of optimal policies. For many of these models, it can be shown that *state-dependent* (s, S) *policies* are optimal. The ordering decision in each period is driven by two thresholds. Specifically, an order is placed if and only if the inventory level falls below the threshold s . In addition, if an order

is placed the inventory level is brought up to the threshold S . The thresholds s and S are determined based on the state of the system at the beginning of the period. Scarf (1960) and Veinott (1966) have established the optimality of (s, S) policies in models with independent demands. Cheng and Sethi (1997) have extended the optimality proof to exogenous Markov-modulated demands that capture cycles and seasonality to some extent. Gallego and Özer (2001) have shown that (s, S) policies are optimal under *advance demand information*, a demand model that allows correlation and forecast updates.

Unfortunately, the rather simple forms of these optimal policies do not usually lead to efficient algorithms for computing the optimal policies. There are very few cases, in which there are efficient algorithms to compute the optimal policies. Federgruen and Zipkin (1984) proposed an algorithm to compute the optimal stationary (s, S) policy in a model with infinite horizon and independent and identically distributed demands. Federgruen and Zheng (1991) described a simple and efficient algorithm to compute the infinite horizon optimal policy in a continuous-review system with demand that is generated by a renewal process. (In this setting, (s, S) policies are equivalent to (R, Q) policies, in which one places an order of Q units, whenever the inventory level drops below R .) For other more complex variants of the model, there are currently no known exact algorithms, but only heuristics. Bollapragada and Morton (1999) proposed a simple myopic policy, assuming that the demands in different periods have the same form of distribution function with the same coefficient of variation but with different means. Gavirneni (2001) designed an efficient heuristic to compute (s, S) policies for nonstationary and capacitated model. Song and Zipkin (1993) considered uncapacitated models with exogenous Markov-modulated Poisson demand. They developed an algorithm to compute the optimal (s, S) policy using a modified value iteration approach. However, they impose strong assumptions on the structure and the size of the state space of the underlying Markov process. Gallego and Özer (2001) and Özer and Wei (2004) considered uncapacitated and capacitated inventory models with advance demand information, respectively. They proposed backward induction algorithms to numerically solve problems with a relatively short planning horizon, and conducted computational experiments to study the impact of advance demand information

on the optimal policy. (In the computational experiments in Section 5, we have applied the newly proposed policies to the instances they considered.) Guan and Miller (2008b) proposed an exact and polynomial-time algorithm for the uncapacitated stochastic economic lot-sizing problem if the stochastic programming scenario tree is polynomially representable. Guan and Miller (2008a) extended these algorithms to allow backlogging. Huang and Küçükyavuz (2008) considered similar problems with random lead times. These models allow 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. To the best of our knowledge, all of the existing heuristics and algorithms, either lack any performance guarantees or can be applied under restrictive assumptions on the demand distributions or the input size.

Structure of this paper. The remainder of the paper is organized as follows: In Section 2, we present the model formulation. Section 3 reviews the *triple-balancing policy* proposed by Levi et al. (2007). We provide a bad example in which the triple-balancing policy fails to work under general demand assumptions. In section 4, we propose what is called *randomized cost-balancing policy* that makes use of order randomization. We show that the policy has a worst-case performance guarantee of 3 under general demand assumptions. Section 5 is devoted to the numerical experiments tested for our newly-proposed policies. The parameterized policies are computationally efficient and near-optimal under advance demand information by Gallego and Özer (2001).

2. The Periodic-Review Stochastic Lot-Sizing Inventory Control Problem

In this section, we provide the mathematical formulation of the periodic-review stochastic lot-sizing inventory control problem. We consider a finite planning horizon of T periods indexed $t = 1, \dots, T$. The demands over these periods are random variables, denoted by D_1, \dots, D_T , and the goal is to coordinate a sequence of orders over the planning horizon to satisfy these demands with minimum cost. As a general convention, from now on we will refer to a random variable and its realization using capital and lower case letters, respectively. Script font is used to denote sets.

In each period $t = 1, \dots, T$, 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 *fixed ordering* cost K that is incurred in each period with strictly positive ordering quantity. (Our model can allow nonstationary K_t , i.e., $\alpha K_{t+1} \leq K_t$ with discount factor α .) 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 t until it is satisfied. In addition, we consider a model with a lead time of L periods between the time an order is placed and the time at which it actually arrives. We assume that the lead time is a known integer L . Following the discussion in Levi et al. (2007), 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 .

At the beginning of each period s , we observe what is called 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, \dots, d_{s-1} over the interval $[1, s)$, and possibly some exogenous information denoted by (w_1, \dots, 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 $F_s = (D_1, \dots, D_{s-1}, W_1, \dots, W_s)$. The set of all possible realizations is denoted by \mathcal{F}_s . The observed information set f_s induces a given conditional joint distribution of the future demands (D_s, \dots, D_T) . For ease of notation, D_t will always denote the random demand in period t according to the conditional joint distribution in some period $s \leq t$, where it will be clear from the context to which period s it refers. The index t will be used to denote a general time period, and s will always refer to the current period. The only assumption on the demands is that for each $s = 1, \dots, T$, and each $f_s \in \mathcal{F}_s$, the conditional expectation $E[D_t | f_s]$ is well defined and finite for each period $t \geq s$. In particular, we allow non-stationary and correlation between the demands in different periods.

The goal is to find an ordering policy that minimizes the overall expected discounted fixed ordering cost, holding cost and backlogging cost. We consider only policies that are *nonanticipatory*,

i.e., at time s , the information that a feasible policy can use consists only of f_s and the current inventory level. The superscripts PL and OPT will be used to refer to a given feasible policy PL and an optimal policy, respectively.

Given a feasible policy PL , the dynamics of the system are described using the following notation. Let $D_{[s,t]}$ denote the cumulative demand over the interval $[s, t]$, i.e., $D_{[s,t]} = \sum_{j=s}^t D_j$. In addition, let NI_t denote the *net inventory* at the end of period t . Thus, $NI_t^+ = \max(NI_t, 0)$ and $NI_t^- = \max(-NI_t, 0)$ are inventory on hand and backlog quantities in period t , respectively. 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., $X_t := NI_{t-1} + \sum_{j=t-L}^{t-1} Q_j$ (for $t = 1, \dots, T$), where Q_j 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 given, 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 $h_t NI_t^+$ holding cost and $b_t NI_t^-$ backlogging cost. 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

$$\mathcal{C}(PL) = \sum_{t=1}^T (h_t NI_t^{+, PL} + b_t NI_t^{-, PL} + K \cdot \mathbf{1}(Q_t^{PL} > 0)). \quad (1)$$

3. Triple-Balancing Policy -Bad Example

In this section, we briefly discuss the triple-balancing policy proposed by Levi et al. (2007) for a special case of the stochastic lot-sizing problem. The discussion sheds light on the limitation of this policy, and motivates the newly proposed randomized cost-balancing policy discussed in section 5. Levi et al. (2007) considered a model in which in each period $t = 1, \dots, T$, conditioning on some information set $f_t \in \mathcal{F}_t$, the conditional distribution of future demands (D_t, \dots, D_T) is such that the demand D_t is known deterministically (i.e., with probability one). This implies that the order in period t is placed *after* the demand in that period is already known. The underlying

assumption here is that at the beginning of period t , our forecast for the demand in that period is sufficiently accurate, so that we can assume forecast to be given deterministically. A primary example is make-to-order systems. However, this assumption does not hold if there is a positive lead time and one considers D_{t+L} instead.

3.1. Description of the policy

First we briefly discuss the original triple-balancing policy in Levi et al. (2007), denoted by TB . This policy is based on the following two rules.

(I) When to order. At the beginning of period t , let s be the last period in which an order is placed before t . An order is placed in period t if and only if by not placing it in period t , the cumulative backlogging cost over the interval $(s, t]$ will exceed K . Once a new order is placed, s is updated to be equal to t . Observe that since, at the beginning of each period t , the conditional joint distribution of future demands is such that D_t is known deterministically, this procedure is well-defined. Notice that an optimal policy will never incur any backlogging costs in a period when an order is placed, since the cumulative backlog quantities are known prior to placing the order.

(II) How much to order. Suppose that an order is placed in period $t < T$. Focus on the holding cost incurred by the units ordered in period t over the interval $[t, T]$. The order is set to the maximum quantity q_t^{TB} , such that the conditional expected marginal holding cost incurred does not exceed K . (The exact definition of marginal holding cost is provided in Section 4.1.)

Worst-case Analysis. The analysis in Levi et al. (2007) showed that the triple-balancing policy has a worst-case performance guarantee of 3. In particular, one can show that, for each time interval between two consecutive orders of the triple-balancing policy, the expected cost incurred by an optimal policy over that interval is at least one-third of the expected cost incurred by the triple-balancing policy over the same interval. However, this is only valid under the restrictive assumptions of no lead times and period demand known at the beginning of the period.

If the period demand is not known at the beginning of the period (or there is a positive lead time), then (I) above is enforced on expectation. It turns out that this policy can perform arbitrarily

bad compared to an optimal policy and does not have a worst-case performance guarantee where the assumptions are dropped. As a result this policy may not be applicable in more general and realistic settings. The example that shows this fact is discussed in section 3.2.

3.2. A bad example

The *triple-balancing policy* can be applied in general settings and one might hope to obtain a worst-case performance guarantee in general. However, the following example shows that such guarantee fails to exist in general. Consider the following instance with infinite horizon $T = \infty$, let $h_t = h = 0$, $b_t = b = 1$, $\forall t \in \mathbb{Z}^+$, $L = 1$ and $K \in \mathbb{Z}^+$, and

$$D_t = \begin{cases} \lambda K & \text{with probability } \frac{1}{\lambda} - \frac{\epsilon}{\lambda K}, \\ 0 & \text{otherwise} \end{cases}, \quad (2)$$

where ϵ is a positive number satisfying $0 < \epsilon < K$. Moreover, the demand drops to 0 in all periods after the first positive demand. Note that the per-unit holding cost is $h = 0$, and therefore there is no penalty for holding extra units in the inventory. The optimal policy orders λK units at the beginning of period 1. The demand λK will eventually come in some period with probability 1. Thus, the optimal cost incurs fixed ordering K only. However, if no demand has arrived, the cumulative backlogging cost is 0, and the expected backlogging cost upon not ordering is $K - \epsilon$. This implies that the policy does not place any orders before the positive demand λK occurs. Thus, the policy incurs a cost of $K + \lambda K$. If we let $\lambda \rightarrow \infty$, the cost ratio goes to ∞ , indicating that the triple-balancing policy can perform arbitrarily bad compared to the optimal cost, and does not admit a worst-case guarantee. This example illustrates that the policy fails to make a good ordering decision, when there is a potential impulse in demand with a positive but small probability. Thus, the policy may incur potentially a very high backlogging cost.

4. Randomized Cost-Balancing Policy

One of the difficulties in the stochastic lot-sizing problem is the need to balance the nonlinear fixed ordering cost against the backlogging cost that may have large spikes because of the variability of the demands. The new policy we propose aims to strike a better balance between these costs

by *randomization*. The policy is called *randomized cost-balancing policy*. To strike this balance the policy employs randomized decision rules. That is, in each period, the decision whether to order and how much to order is based on a suitably chosen randomized decision rule; the policy chooses among various order quantities with certain respective probabilities. Before the description of the new policy, we briefly discuss a marginal cost accounting scheme that is used to employ the policy. This cost accounting scheme was introduced by Levi et al. (2007).

4.1. Marginal Cost Accounting Scheme

Following Levi et al. (2007), we next describe an alternative cost accounting scheme that is called *marginal cost accounting scheme*. Unlike (1) that decomposes the cost by periods, 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 decision, and are only affected by future demands. This may include costs in subsequent periods.

Focus first 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 decisions. Specifically, the total marginal holding cost associated with the decision to order q_s units in period s is defined to be $\sum_{j=s+L}^T h_j (q_s - (D_{[s,j]} - x_s)^+)^+$. Note that at the time the order q_s is made, 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., $b_{s+L} (D_{[s,s+L]} - (x_s + q_s))^+$, is also affected only by the future demands. This leads to a marginal cost accounting scheme. For each feasible policy PL , let H_t^{PL} be the holding cost incurred by the Q_t^{PL} units ordered in period t (for $t = 1, \dots, T$) over the interval $[t, T]$, and let Π_t^{PL} be the backlogging cost associated with period t , i.e., the cost incurred a lead time ahead in period $t + L$ ($t = 1 - L, \dots, T - L$). That is,

$$H_t^{PL} = H_t(Q_t^{PL}) = \sum_{j=t+L}^T h_j (Q_t^{PL} - (D_{[t,j]} - X_t)^+)^+,$$

$$\Pi_t^{PL} = \Pi_t(Q_t^{PL}) = b_{t+L} (D_{[t,t+L]} - (X_t + Q_t^{PL}))^+.$$

Let $\mathcal{C}(PL)$ be again the cost of the policy PL . Clearly, we have

$$\mathcal{C}(PL) = \sum_{t=1-L}^0 \Pi_t^{PL} + H_{(-\infty,0]} + \sum_{t=1}^{T-L} (K \cdot \mathbb{1}(Q_t^{PL} > 0) + H_t^{PL} + \Pi_t^{PL}), \quad (3)$$

where $H_{(-\infty,0]}$ denotes the total expected holding cost incurred by units ordered before period 1. We note that the first two expressions $\sum_{t=1-L}^0 \Pi_t^{PL}$ and $H_{(-\infty,0]}$ are not affected by any decision (i.e., they are the same for any feasible policy and each realization of the demands) and, therefore, we will omit them. Since they are nonnegative, this will not affect our approximation results. Also, observe that without loss of generality, we can assume that $Q_t^{PL} = H_t^{PL} = 0$ for any policy PL in each period $t = T - L + 1, \dots, T$, since nothing ordered in these periods can be used within the given planning horizon. We now can write the effective cost of a policy PL as

$$\mathcal{C}(PL) = \sum_{t=1}^{T-L} (K \cdot \mathbb{1}(Q_t^{PL} > 0) + H_t^{PL} + \Pi_t^{PL}). \quad (4)$$

4.2. Description of the 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 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. We denote the randomized cost-balancing policy by RB . The decision in each period, whether to order and how much to order, is based on the following quantities.

- Compute the *balancing quantity* \hat{q}_t that balances the conditional expected marginal holding cost incurred by the units ordered against the conditional expected backloging cost in period $t + L$. That is, \hat{q}_t solves

$$E[H_t^{RB}(\hat{q}_t) | f_t] = E[\Pi_t^{RB}(\hat{q}_t) | f_t], \quad (5)$$

where H_t^{RB} and Π_t^{RB} are defined as in Section 4.1, respectively. Let $\theta_t = \theta_t(f_t) \triangleq E[H_t^{RB}(\hat{q}_t) | f_t] = E[\Pi_t^{RB}(\hat{q}_t) | f_t]$ denote the *balancing cost*. The solution to (5) is unique and can be computed efficiently via bi-section search (Levi et al. (2007)).

- Compute the *holding-cost-K quantity* \tilde{q}_t that solves $E[H_t^{RB}(\tilde{q}_t) | f_t] = K$, i.e., \tilde{q}_t is the order quantity that brings the conditional expected marginal holding cost to K . Note that \tilde{q}_t can be computed readily since $E[H_t^{RB}(\cdot) | f_t]$ is monotonically increasing.
- Compute $E[\Pi_t^{RB}(\tilde{q}_t) | f_t]$, i.e., the resulting conditional expected backlogging cost in period $t + L$ if one orders the holding-cost-K quantity \tilde{q}_t units in period t .
- Compute $E[\Pi_t^{RB}(0) | f_t]$, i.e., the conditional expected backlogging cost in period $t + L$ resulting from not ordering in period t .

Based on the above quantities computed, the following randomized rule is used in each period t . Let P_t denote our ordering probability which is a priori random. With the observed information set f_t , the ordering probability $p_t = P_t | f_t$ in period t is defined differently in the two cases below.

Case (I)

If the balancing cost exceeds K , i.e., $\theta_t \geq K$, the *RB* policy orders the balancing quantity $q_t^{RB} = \hat{q}_t$ with probability $p_t = 1$. The intuition is that when $\theta_t \geq K$, the fixed ordering cost K is less dominant compared to marginal holding and backlogging costs. Moreover, if the *RB* policy does not place an order, the conditional expected backlogging cost is potentially large. Thus, it is worthwhile to order the balancing quantity $q_t^{RB} = \hat{q}_t$ with probability $p_t = 1$.

Case (II)

If the balancing cost is less than K , i.e., $\theta_t < K$, the *RB* policy orders the holding-cost-K quantity (i.e., $q_t^{RB} = \tilde{q}_t$) with probability p_t and nothing with probability $1 - p_t$. That is,

$$q_t^{RB} = \begin{cases} \tilde{q}_t, & \text{with probability } p_t \\ 0, & \text{with probability } 1 - p_t \end{cases} \quad (6)$$

The probability p_t is computed by solving the following equation

$$p_t K = p_t \cdot E[\Pi_t^{RB}(\tilde{q}_t) | f_t] + (1 - p_t) \cdot E[\Pi_t^{RB}(0) | f_t]. \quad (7)$$

The underlying reason behind the choice of this particular randomization in (7) is that the policy perfectly balances the three types of costs, namely, the marginal holding cost, the marginal backlogging cost and the fixed ordering cost associated with the period t . In particular, since we order

the holding-cost- K quantity with probability p_t and nothing with probability $1 - p_t$, the conditional expected marginal holding cost in this case is

$$E[H_t^{RB}(q_t^{RB}) | f_t] = p_t E[H_t^{RB}(\tilde{q}_t) | f_t] + (1 - p_t) E[H_t^{RB}(0) | f_t] = p_t K. \quad (8)$$

By the construction of p_t in Equation (7), the conditional expected backloging cost is

$$E[\Pi_t^{RB}(q_t^{RB}) | f_t] = p_t E[\Pi_t^{RB}(\tilde{q}_t) | f_t] + (1 - p_t) E[\Pi_t^{RB}(0) | f_t] = p_t K. \quad (9)$$

Since p_t is the ordering probability in Case (II), the expected fixed ordering cost is $p_t K$. It can be shown that Equation (7) has the following solution,

$$0 \leq p_t = \frac{E[\Pi_t^{RB}(0) | f_t]}{K - E[\Pi_t^{RB}(\tilde{q}_t) | f_t] + E[\Pi_t^{RB}(0) | f_t]} < 1. \quad (10)$$

The inequalities in Equation (10) follows from the fact that $\theta_t < K$ and $\tilde{q}_t > \hat{q}_t$, which implies that $E[\Pi_t^{RB}(\tilde{q}_t) | f_t] < E[\Pi_t^{RB}(\hat{q}_t) | f_t] = \theta_t < K$. Figure 1 illustrates how the RB policy computes the ordering probability p_t in Case (II) where $\theta_t < K$.

This concludes the description of the RB policy. In the next section, we shall show that the RB policy has an expected worst-case performance guarantee of 3.




Figure 1 A graphical depiction of how the RB policy computes the probability of ordering p_t when the balancing cost θ is below the fixed ordering cost K (Case (II)).

4.3. Worst-case analysis

To obtain a 3-approximation, one wishes to show that on expectation the cost of an optimal policy can ‘pay’ for at least one-third of the expected cost of the randomized cost-balancing policy. The periods are decomposed into subsets in which we will define explicitly. For certain well-behaved subsets, we want to show that the holding and backlogging costs incurred by an optimal policy can ‘pay’ for one-third of the cost incurred by the RB policy. The difficulty arises in analyzing the remaining subset of *problematic periods*, for which it is not a priori clear how to ‘pay’ for their cost. These problematic periods are further partitioned into intervals defined by each two consecutive orders placed by the optimal policy. It can be shown that the total expected cost incurred by the RB policy in problematic periods within each interval, does not exceed $3K$. This implies that the fixed ordering cost incurred by an optimal policy can ‘pay’ on expectation one-third of the cost incurred by the randomized cost-balancing policy in problematic periods. Next we discuss the details of this approach, and we defer all proofs to Electronic Companion for ease of presentation.

Let Z_t^{RB} be a random variable defined as

$$Z_t^{RB} := E[H_t^{RB}(Q_t^{RB}) | F_t] = E[\Pi_t^{RB}(Q_t^{RB}) | F_t]. \quad (11)$$

Note that Z_t^{RB} is a random variable that is realized with the information set in period t . Observe that by the construction of the RB policy, the random variable Z_t^{RB} is well-defined since the expected marginal holding costs and the expected marginal backlogging costs are always balanced. That is, the conditional expected marginal holding cost is always equal to the conditional expected backlogging cost. In the following lemma we show that the expected cost of the RB policy can be upper bounded using the Z_t^{RB} variables defined in (11).

LEMMA 1. *Let $\mathcal{C}(RB)$ be the total cost incurred by the RB policy. Then we have,*

$$E[\mathcal{C}(RB)] \leq 3 \cdot \sum_{t=1}^{T-L} E[Z_t^{RB}]. \quad (12)$$

To complete the worst-case analysis, we would like to show that the expected cost of an optimal policy denoted by OPT is at least $\sum_{t=1}^{T-L} E[Z_t^{RB}]$. This will be done by amortizing the cost of OPT

against the cost of the RB policy. In particular, we shall show that on expectation OPT pays for a large fraction of the cost of the RB policy. In the subsequent analysis, we will use a random partition of periods $t = \{1, 2, \dots, T - L\}$ to the following sets:

The set $\mathcal{T}_{1H} \triangleq \{t : \Theta_t \geq K \text{ and } Y_t^{OPT} > Y_t^{RB}\}$ consists of periods in which the balancing cost Θ_t exceeds K and the optimal policy had higher inventory position than that of the RB policy after ordering (recall that if $\Theta_t \geq K$ then the RB policy orders the balancing quantity with probability 1 and the value Y_t^{RB} is known deterministically (i.e., realized) with F_t).

The set $\mathcal{T}_{1II} \triangleq \{t : \Theta_t \geq K \text{ and } Y_t^{OPT} \leq Y_t^{RB}\}$ consists of periods in which the balancing cost exceeds K and the inventory position of the optimal policy does not exceed that of the RB policy after ordering (see the comment above regarding \mathcal{T}_{1H}).

The set $\mathcal{T}_{2H} \triangleq \left\{t : \Theta_t < K \text{ and } Y_t^{OPT} \geq X_t^{RB} + \tilde{Q}_t^{RB}\right\}$ consists of periods in which the balancing cost is less than K and, in such periods, the inventory position of the RB policy after ordering would be either X_t^{RB} if no order was placed, or $X_t^{RB} + \tilde{Q}_t^{RB}$ if the *holding-cost- K quantity* is ordered, depending on the randomized decision of the RB policy. However, the inventory position of OPT after ordering exceeds even $X_t^{RB} + \tilde{Q}_t^{RB}$. (Note again that the quantity \tilde{Q}_t^{RB} is known deterministically (i.e., realized) with F_t .)

Analogous to \mathcal{T}_{2H} , the set $\mathcal{T}_{2II} \triangleq \{t : \Theta_t < K \text{ and } X_t^{RB} \geq Y_t^{OPT}\}$ consists of periods in which the inventory position of OPT after ordering is below X_t^{RB} .

The set $\mathcal{T}_{2M} \triangleq \left\{t : \Theta_t < K \text{ and } X_t^{RB} < Y_t^{OPT} < X_t^{RB} + \tilde{Q}_t^{RB}\right\}$ consists of periods in which the balancing cost is less than K and the inventory position of OPT after ordering is within $(X_t^{RB}, X_t^{RB} + \tilde{Q}_t^{RB})$. Thus, whether the RB policy or OPT has more inventory depends on whether the RB policy placed an order.

Note that the sets $(\mathcal{T}_{1H} - \mathcal{T}_{2M})$ are disjoint and the union makes a complete set. Conditioning on f_t , it is already known which part of the partition period t belongs.

Next we will show that the total holding cost incurred by OPT is higher than the marginal holding cost incurred by the RB policy in periods that belong to $\mathcal{T}_{1H} \cup \mathcal{T}_{2H}$, and that the total

backlogging cost incurred by OPT is higher than the backlogging cost incurred by the RB policy associated with periods within $\mathcal{T}_{1\Pi} \cup \mathcal{T}_{2\Pi}$.

LEMMA 2. *The overall holding cost and backlogging cost incurred by OPT are denoted by H^{OPT} and Π^{OPT} , respectively. Then we have, with probability 1,*

$$H^{OPT} \geq \sum_t H_t^{RB} \cdot \mathbb{1}(t \in \mathcal{T}_{1H} \cup \mathcal{T}_{2H}), \quad \Pi^{OPT} \geq \sum_t \Pi_t^{RB} \cdot \mathbb{1}(t \in \mathcal{T}_{1\Pi} \cup \mathcal{T}_{2\Pi}). \quad (13)$$

Note that the periods in the set \mathcal{T}_{2M} introduce some uncertainties in the relation between the inventory positions after ordering of the RB policy and OPT . Thus, we are unable to carry out an analysis similar to Lemma 2. For this reason, we call \mathcal{T}_{2M} a *problematic* set of periods. Naturally, we also define the *non-problematic* set of periods to be $\mathcal{T}_N = \mathcal{T}_{1H} \cup \mathcal{T}_{1\Pi} \cup \mathcal{T}_{2\Pi} \cup \mathcal{T}_{2H}$. The analysis of the problematic periods in the set \mathcal{T}_{2M} will be done in two steps. In the first step, we will conceptually create a *bank account* A that will be used to pay some of the cost of the RB policy in these problematic periods. In particular, for each period $t \in \mathcal{T}_{2M}$, we *borrow* an amount of Z_t^{RB} from the bank account. Thus, the total amount of borrowing from the bank is given by $A = \sum_{t \in \mathcal{T}_{2M}} Z_t^{RB}$, and so $E[A] = E[\sum_t Z_t^{RB} \cdot \mathbb{1}(t \in \mathcal{T}_{2M})]$.

The following lemma shows that, with the borrowed amount A from the bank, the overall holding cost and backlogging cost incurred by OPT exceed $\sum_{t=1}^{T-L} E[Z_t^{RB}]$. The next step will be to show that $E[A]$ is at most the expected fixed ordering cost incurred by OPT . That is,

$$E[A] \leq E \left[\sum_{t=1}^{T-L} K \cdot \mathbb{1}(Q_t^{OPT} > 0) \right]. \quad (14)$$

LEMMA 3. *The expected holding cost and backlogging cost incurred by OPT plus the expected amount borrowed from the bank account A are at least $\sum_{t=1}^{T-L} E[Z_t^{RB}]$. That is, the following inequality holds*

$$E[(H^{OPT} + \Pi^{OPT}) + A] \geq \sum_{t=1}^{T-L} E[Z_t^{RB}]. \quad (15)$$

By Lemmas 1 and 3, the overall holding and backlogging costs incurred by OPT , *plus* the borrowed amount A from the bank, account on expectation for *one-third* of the overall expected costs incurred by the RB policy. To complete the worst-case analysis, we will show in Lemma 4 that

the expected amount borrowed from the bank account does not exceed the expected fixed ordering cost incurred by OPT , i.e., $E \left[\sum_{t=1}^{T-L} K \cdot \mathbf{1}(Q_t^{OPT} > 0) \right]$. We will highlight the key steps involved in proving this lemma. We decompose the problematic periods in the set \mathcal{T}_{2M} into intervals between ordering points of OPT , and we want to show that, for each such interval, the fixed ordering cost K incurred by OPT will cover the expected amount borrowed from the bank in periods that belong to set \mathcal{T}_{2M} . Conditioning on f_T^- (the entire evolution of the system excluding the randomized decisions of the RB policy), we construct a decision tree based on the randomized decisions of the RB policy. We then show that, by a tree traversal argument and Lemma 5, the expected borrowing from the problematic nodes (which belong to the set \mathcal{T}_{2M}) within an interval between ordering points of OPT does not exceed K .

LEMMA 4. *The following inequality holds*

$$E[A] \leq E \left[\sum_{t=1}^{T-L} K \cdot \mathbf{1}(Q_t^{OPT} > 0) \right]. \quad (16)$$

In other words, the expected borrowing $E[A]$ is less than the total expected fixed ordering cost incurred by OPT .

LEMMA 5. *Let $\{p_l\}_{l=1}^{\infty}$ satisfy the condition $0 \leq p_l \leq 1$ for all l . Then the following inequality holds,*

$$p_1^2 + \sum_{l=2}^{\infty} \left\{ \left(\prod_{s=1}^{l-1} (1 - p_s) \right) p_l \left(\sum_{k=1}^l p_k \right) \right\} \leq 1. \quad (17)$$

As an immediate consequence of Lemmas 3 and 4, we obtain the following lemma and theorem.

LEMMA 6. *Let $\mathcal{C}(OPT)$ be the total cost incurred by the cost-balancing policy RB . Then we have,*

$$E[\mathcal{C}(OPT)] \geq \sum_{t=1}^{T-L} E[Z_t^{RB}]. \quad (18)$$

THEOREM 1. *For each instance of the stochastic lot-sizing problem, the expected cost of the randomized cost-balancing policy RB is at most three times the expected cost of an optimal policy OPT , i.e.,*

$$E[\mathcal{C}(RB)] \leq 3 \cdot E[\mathcal{C}(OPT)]. \quad (19)$$

It should be noted that our analysis remains valid for nonstationary K_t (i.e., $\alpha K_{t+1} \leq K_t$ as is commonly assumed in the literature). Again, without loss of generality, assume that the discount factor $\alpha = 1$. One can show that, similar to (16), the expected borrowing

$$E[A] \leq E \left[\sum_{t=1}^{T-L} K_t \cdot \mathbb{1}(Q_t^{OPT} > 0) \right]. \quad (20)$$

The idea is that the expected borrowing of the RB policy from the problematic nodes within an interval between ordering points of OPT does not exceed K_t with the assumption of decreasing fixed ordering cost. The rest of arguments readily carries through.

5. Numerical Experiments

The randomized cost-balancing policies described above can be parameterized to obtain general classes of policies, respectively. The worst-case analysis discussed above can then be viewed as choosing parameter values that perform well against any possible instance. In contrast, 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 can refine 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 under the model of *advanced demand information* proposed by Gallego and Özer (2001) and Özer and Wei (2004). To the best of our knowledge, these are the few papers that report computational results (by brute-force backward induction algorithm) on the stochastic lot-sizing problem with correlated demands.

5.1. Parameterized policies.

We describe a class of parameterized policies involving parameters β , γ and η where β controls the holding-cost- βK quantity, γ controls the ratio of marginal holding costs and backlogging costs and η controls the level of expected backlogging cost resulting from not ordering.

- The balancing quantity \hat{q}_t that solves $E[H_t^{RB}(\hat{q}_t) | f_t] = \gamma \cdot E[\Pi_t^{RB}(\hat{q}_t) | f_t] := \theta_t$.
- The holding-cost- βK quantity \tilde{q}_t that solves $E[H_t^{RB}(\tilde{q}_t) | f_t] = \beta \cdot K$.

- Compute $E[\Pi_t^{RB}(\tilde{q}_t) | f_t]$, and $\eta \cdot E[\Pi_t^{RB}(0) | f_t]$.

(I) If $\theta_t \geq \beta \cdot K$, the *RB* policy orders $q_t^{RB} = \hat{q}_t$ with probability $p_t = 1$ in period t .

(II) If $\theta_t < \beta \cdot K$, the *RB* policy orders $q_t^{RB} = \tilde{q}_t$ with probability p_t and order nothing with probability $1 - p_t$ in period t , where the probability $p_t = \frac{\eta \cdot E[\Pi_t^{RB}(0) | f_t]}{\beta \cdot K - E[\Pi_t^{RB}(\tilde{q}_t) | f_t] + \eta \cdot E[\Pi_t^{RB}(0) | f_t]}$.

Since T is relatively small, we also introduce an end-of-horizon rule. Suppose we are in period t , we estimate the total expected cumulative backlogging cost (assuming no orders are placed) over the interval $[t, T]$. If this amount is less than K , we do not order in period t .

5.2. Experiment Design

Under advance information model, the demand vector in each period t is observed as $D_t = (D_{t,t}, \dots, D_{t,t+N})$ where $D_{t,s}$ represents order placed by customers during period t for future periods $s \in \{t, \dots, t+N\}$ and N is the length of the information horizon over which we have advance demand information. Note that D_t is a random vector and is realized only at the end of period t . At the beginning of period t , the demand to prevail in a future period s ($s \geq t$) can be divided into two parts: the observed demand vector $\sum_{r=s-N}^{t-1} D_{r,s}$ and the unobserved demand vector $\sum_{r=t}^s D_{r,s}$. As a result, this introduces a correlation between period demands (however the conditional joint distribution of the future demands is known in each period t). The state space of the proposed dynamic programming formulation contains the inventory position and the observed demand vector which explodes exponentially with the length of the information horizon N when $N > L + 2$. Gallego and Özer (2001) verified some structural properties of the dynamic program via numerical studies for a number of small instances. The experiments that we performed expand their numerical studies by incorporating non-zero lead times as well as longer planning horizons. Following the methodology of Aviv and Federgruen (2001), we generated a total of 90 instances to test the quality of the randomized-balancing heuristics compared to the optimal cost. The instances we used have the following combination of parameters: $T = 12, 15$, $L = 0, 1, 2$, $N = L + 2$, $K = 0, 5, 50, 100$, $h = 1, 2, 3, 6$, $p = 1, 3, 6, 9$ and $(D_{t,t}, D_{t,t+1}, D_{t,t+2})$ are modeled by Poisson random variables with mean $\lambda_0, \lambda_1, \lambda_2$.

5.3. 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, $U = 10$ and $\Delta = 0.1$ 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 an upper bound on the balancing and hold-cost-K quantities. For each $t = 1, \dots, T$, 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 $O(T^2 \log \hat{U})$, 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 $T = 12$, the average CPU time per test instance on a Pentium 1.58GHz PC is 233s. In contrast, the dynamic programming algorithm takes 1840s on average per test instance.

5.4. Numerical results

The numerical results with $(T, L) = (12, 0)$, $(T, L) = (12, 2)$ and $(T, L) = (15, 0)$ are tabulated in Table EC.1, Table EC.2 and Table EC.3, respectively (refer to Electronic Companion). The (*) in both tables indicates that the designated parameters can take arbitrary numbers without affecting the optimal values of the parameterized policy. It is observed that $(\beta^*, \eta^*) = (*, *)$ in all instances where $K = 0$, since the holding-cost- β^*K quantity is trivially 0 and therefore the algorithm only considers the balancing quantities. In some instances where K is relatively large and the holding-cost- β^*K quantity is near-optimal, it is observed that $\gamma^* = (*)$ implying that the algorithm only orders the holding-cost- β^*K quantities. For the rest of instances, the algorithm uses both the balancing quantity and the holding-cost-K quantity.

In the case where $L = 0$, on average the parameterized RB policy performs within 4.6% and always within 7% of the optimal cost for $T = 12, 15$. The numerical results show that the performance of the parameterized RB policy is insensitive to the planning horizon T . Moreover, the optimal parameters in the parameterized RB policy are intuitive: β controls the quantity of each order; γ controls the ratio in which the marginal holding cost is balanced against the marginal backlogging cost; η controls the weight put on the do-nothing backlogging cost resulted from not ordering. The optimal $\eta^* = 9$ coincides with the ratio of p to h , which implies that more weight should be put on backlogging cost so that the ordering probability can be increased. The optimal $\gamma^* = 2$ suggests that the marginal holding cost should be twice the backlogging cost. The optimal β^* is close to 1 when K is large, implying that using the holding-cost- K quantity is near optimal. The unparameterized RB policy (i.e., $(\beta, \gamma, \eta) = (1, 1, 1)$) performs on average within 27% and always within 50% error of optimal cost, which is significantly better than the theoretical worst-case performance guarantee of 3. The cost ratio is observed to be decreasing in the magnitude of fixed ordering cost K . In the case where $L = 2$, the parameterized RB policy performs on average within 10% and always within 16% error of the optimal cost. The optimal parameters are similar to those in $L = 0$. The deviation from the optimal cost is resulted from stocking more inventory units by the RB policy, as the lead time induces more uncertainty in future demands. The unparameterized RB policy performs within 50% (on average 29%) error of optimal cost. It is also noted that the average CPU time of running the RB policy is insensitive to the planning horizon T .

Acknowledgments

The authors thank the Associate Editor and the anonymous referees for a handful of constructive comments that improved the content and exposition of this paper. The research of the authors was partially supported by NSF grants DMS-0732175 and CMMI-0846554 (CAREER Award), an AFOSR award FA9550-08-1-0369, an SMA grant and the Buschbaum Research Fund of MIT.

References

- Aviv, Y., A. Federgruen. 2001. Capacitated multi-item inventory systems with random and seasonally fluctuating demands: Implications for postponement strategies. *Management Science* **47**(4) 512–531.

- Bollapragada, S., T. E. Morton. 1999. A simple heuristic for computing nonstationary (s,S) policies. *Operations Research* **47**(4) 576–584.
- Borodin, A., R. El-Yaniv. 1998. *Online Computation and Competitive Analysis*. Cambridge University Press.
- Bouchard, B., N. E. Karoui, N. Touzi. 2005. Maturity randomization for stochastic control problems. *The Annals of Applied Probability* **15**(4) 2575–2605.
- Cheng, F., S. Sethi. 1997. Optimality of (s, S) policies in inventory models with markovian demand. *Operations Research* **45**(6) 931–939.
- Federgruen, A., Y. S. Zheng. 1991. An efficient algorithm for computing an optimal (r, Q) policy in continuous review stochastic inventory systems. *Operations Research* **40**(4) 808–813.
- Federgruen, A., P. H. Zipkin. 1984. An efficient algorithm for computing optimal (s, S) policies. *Operations Research* **32**(6) 1268–1285.
- Gallego, G., Ö. Özer. 2001. Integrating replenishment decisions with advance demand information. *Management Science* **47**(10) 1344–1360.
- Gavirneni, S. 2001. An efficient heuristic for inventory control when the customer is using a (s, S) policy. *Operations Research Letters* **28**(4) 187–192.
- Guan, Y., A. J. Miller. 2008a. A polynomial time algorithm for the stochastic uncapacitated lot-sizing problem with backlogging. *IPCO 2008, volume 5035 of Lecture Notes in Computer Science* 450–462.
- Guan, Y., A. J. Miller. 2008b. Polynomial time algorithms for stochastic uncapacitated lot-sizing problems. *Operations Research* **56**(5) 1172–1183.
- Huang, K., S. Küçükyavuz. 2008. On stochastic lot-sizing problems with random lead times. *Operations Research Letters* **36**(3) 303–308.
- Levi, R., G. Janakiraman, M. Nagarajan. 2008a. A 2-approximation algorithm for stochastic inventory control models with lost-sales. *Mathematics of Operations Research* **33**(2) 351–374.
- Levi, R., M. Pál, R. O. Roundy, D. B. Shmoys. 2007. Approximation algorithms for stochastic inventory control models. *Mathematics of Operations Research* **32**(2) 284–302.
- Levi, R., R. O. Roundy, D. B. Shmoys, M. Sviridenko. 2008b. A constant approximation algorithm for the one-warehouse multi-retailer problem. *Management Science* **54**(4) 763–776.

- Levi, R., R. O. Roundy, D. B. Shmoys, V. A. Truong. 2008c. Approximation algorithms for capacitated stochastic inventory models. *Operations Research* **56**(5) 1184–1199.
- Levi, R., R. O. Roundy, V. A. Truong. 2005. Provably near-optimal balancing policies for multi-echelon stochastic inventory control models. Working Paper.
- Motwani, R., P. Raghavan. 1995. Randomized algorithms. *Cambridge University Press* .
- Özer, Ö., W. Wei. 2004. Inventory control with limited capacity and advance demand information. *Operations Research* **52**(6) 988–1000.
- Prandini, M., J. Lygeros, A. Nilim, S. Sastry. 1999. Randomized algorithms for probabilistic aircraft conflict detection. *IEEE Conference on Decision and Control* **3** 2444–2449.
- Rust, J. 1997. Using randomization to break the curse of dimensionality. *Econometrica* **65**(3) 487–516.
- Scarf, H. 1960. The optimality of (S,s) policies in the dynamic inventory problem. *Mathematical Methods in Social Sciences* .
- Shmoys, D., C. Swamy. 2006. An approximation scheme for stochastic linear programming and its application to stochastic integer programs. *Journal of the ACM* **53**(6) 978–1012.
- Shmoys, D., K. Talwar. 2008. A constant approximation algorithm for the a priori traveling salesman problem. *Integer Programming and Combinatorial Optimization* **5035** 331–343.
- Song, J., P. H. Zipkin. 1993. Inventory control in a fluctuating demand environment. *Operations Research* **41**(2) 351–370.
- Teo, C. P., D. Bertsimas. 1996. Improved randomized approximation algorithms for lot-sizing problems. *Proceedings of the 5th International IPCO Conference on Integer Programming and Combinatorial Optimization* 359–373.
- Veinott, A. 1966. On the optimality of (s,S) inventory policies: new conditions and a new proof. *SIAM J. Appl. Math* **14**(5) 1067–1083.

Electronic Companion

EC.1. Proofs of Technical Lemmas and Theorems

LEMMA 1. Let $\mathcal{C}(RB)$ be the total cost incurred by the RB policy. Then we have,

$$E[\mathcal{C}(RB)] \leq 3 \cdot \sum_{t=1}^{T-L} E[Z_t^{RB}]. \quad (\text{EC.1})$$

Proof of Lemma 1. Using the marginal cost accounting in Equation (4) and standard arguments of conditional expectations, we express

$$\begin{aligned} E[\mathcal{C}(RB)] &= \sum_{t=1}^{T-L} E[H_t^{RB}(Q_t^{RB}) + \Pi_t^{RB}(Q_t^{RB}) + K \cdot \mathbf{1}(Q_t^{RB} > 0)] \\ &= \sum_{t=1}^{T-L} E[E[H_t^{RB}(Q_t^{RB}) + \Pi_t^{RB}(Q_t^{RB}) + K \cdot \mathbf{1}(Q_t^{RB} > 0) \mid F_t]] \\ &= \sum_{t=1}^{T-L} E[2Z_t^{RB} + P_t K] \leq 3 \sum_{t=1}^{T-L} E[Z_t^{RB}]. \end{aligned} \quad (\text{EC.2})$$

The third equality follows directly from (11). To establish the first inequality in (EC.2) above, we shall show that $Z_t \geq P_t K$ almost surely. That is, for each $f_t \in F_t$, $z_t \geq p_t K$. Given any information set f_t , all the quantities x_t , θ_t , ψ_t , ϕ_t and p_t defined above are known deterministically. We split the analysis into two cases:

1. If $\theta_t \geq K$, then $q_t^{RB} = \hat{q}_t$ (the *balancing quantity*) with probability $p_t = 1$ implying $z_t = \theta_t \geq K$.

The claim follows.

2. If $\theta_t < K$, then $q_t^{RB} = \tilde{q}_t$ (the *holding-cost-K quantity*) with probability p_t and $q_t^{RB} = 0$ with $1 - p_t$. Thus, by Equations (8) and (9), we have $z_t = p_t K$, and the claim follows.

This concludes the proof of the lemma. \square

LEMMA 2. The overall holding cost and backlogging cost incurred by OPT are denoted by H^{OPT} and Π^{OPT} , respectively. Then we have, with probability 1,

$$H^{OPT} \geq \sum_t H_t^{RB} \cdot \mathbf{1}(t \in \mathcal{F}_{1H} \cup \mathcal{F}_{2H}), \quad \Pi^{OPT} \geq \sum_t \Pi_t^{RB} \cdot \mathbf{1}(t \in \mathcal{F}_{1\Pi} \cup \mathcal{F}_{2\Pi}). \quad (\text{EC.3})$$

Proof of Lemma 2. The proof is identical to Lemmas 4.2 and 4.3 in Levi et al. (2007). \square

LEMMA 3. *The expected holding cost and backlogging cost incurred by OPT plus the expected amount borrowed from the bank account A are at least $\sum_{t=1}^{T-L} E[Z_t^{RB}]$. That is, The following inequality holds*

$$E[(H^{OPT} + \Pi^{OPT}) + A] \geq \sum_{t=1}^{T-L} E[Z_t^{RB}]. \quad (\text{EC.4})$$

Proof of Lemma 3. Using linearity of expectation, it suffices to show

$$E[H^{OPT} + \Pi^{OPT}] \geq \sum_{t=1}^{T-L} E[\mathbf{1}(t \in \mathcal{I}_N) \cdot Z_t^{RB}]. \quad (\text{EC.5})$$

Using Lemma 2 and standard arguments of condition expectations, we have

$$\begin{aligned} E[H^{OPT}] &\geq E\left[\sum_t H_t^{RB} \cdot \mathbf{1}(t \in \mathcal{I}_{1H} \cup \mathcal{I}_{2H})\right] \\ &= E\left[E\left[\sum_t H_t^{RB} \cdot \mathbf{1}(t \in \mathcal{I}_{1H} \cup \mathcal{I}_{2H}) \mid F_t\right]\right] \\ &= E\left[\sum_t Z_t^{RB} \cdot \mathbf{1}(t \in \mathcal{I}_{1H} \cup \mathcal{I}_{2H})\right]. \end{aligned} \quad (\text{EC.6})$$

Similarly, we also have

$$E[\Pi^{OPT}] \geq E\left[\sum_t Z_t^{RB} \cdot \mathbf{1}(t \in \mathcal{I}_{1\Pi} \cup \mathcal{I}_{2\Pi})\right]. \quad (\text{EC.7})$$

Equation (EC.5) follows from summing up Equations (EC.6) and (EC.7). \square

LEMMA 4. *The following inequality holds*

$$E[A] \leq E\left[\sum_{t=1}^{T-L} K \cdot \mathbf{1}(Q_t^{OPT} > 0)\right]. \quad (\text{EC.8})$$

In other words, the expected borrowing $E[A]$ is less than the total expected fixed ordering cost incurred by OPT.

Proof of Lemma 4. First we define the *reduced information set* f_t^- to be the information up to period t *excluding* the randomized decisions of the RB policy over $[1, t-1]$. In particular, given the entire evolution of demand f_T^- , the sequence of orders placed by OPT is known deterministically. Let $1 \leq t_1 < t_2 < \dots < t_n \leq T-L$ be the periods in which OPT placed $n = n \mid f_T^-$ orders sequentially. Let $t_0 = 0$ and $t_{n+1} = T-L+1$. We shall show that there are no problematic periods within (t_0, t_1) and that, for each $i = 1, \dots, n$, the expected borrowing within the interval $[t_i, t_{i+1})$ does not exceed K . That is,

$$(t_0, t_1) \cap \mathcal{T}_{2M} = \emptyset, \tag{EC.9}$$

$$E \left[\sum_{t \in [t_i, t_{i+1}) \cap \mathcal{T}_{2M}} Z_t^{RB} \mid f_T^- \right] \leq K. \tag{EC.10}$$

It is important to note that f_T^- does not include the randomized decisions of the *RB* policy. Thus, the set \mathcal{T}_{2M} is still random and so is the amount borrowed from the bank. In particular, the expectation in Equation (EC.10) is taken with respect to the randomized decisions of the *RB* policy. Equations (EC.10) and (EC.9) imply that, for each f_T^- ,

$$E \left[\sum_{t \in \mathcal{T}_{2M}} Z_t^{RB} \mid f_T^- \right] \leq K \cdot n \mid f_T^- = K \cdot n, \tag{EC.11}$$

and therefore

$$E[A] \leq K \cdot E[N] = E \left[\sum_{t=1}^{T-L} K \cdot \mathbf{1}(Q_t^{OPT} > 0) \right]. \tag{EC.12}$$

Thus, it suffices to prove Equations (EC.10) and (EC.9). Figure EC.1 gives a graphical interpretation of Equation (EC.10), i.e., we want to show that the fixed ordering cost K incurred by *OPT* in period t_i will cover the expected amount borrowed from the bank in periods that belong to set \mathcal{T}_{2M} within the interval $[t_i, t_{i+1})$.




Figure EC.1 Decomposition of the problematic periods in the set \mathcal{T}_{2M} into intervals between ordering points of *OPT*

Proof of Equation (EC.9). We first show that Equation (EC.9) holds. Recall the definition $\mathcal{T}_{2M} = \{t : \Theta_t < K \text{ and } X_t^{RB} < Y_t^{OPT} \leq X_t^{RB} + \tilde{Q}_t^{RB}\}$. Since at the beginning of the planning horizon, it is assumed that every feasible policy will have the same initial inventory position, it follows that if period t is in \mathcal{T}_{2M} , *OPT* must have placed an order and overtaken the inventory position of the *RB* policy. (The two policies face the same sequence of demands.) However, (t_0, t_1) denotes the set of periods in which *OPT* has not placed any order yet. Thus, the intersection of these two sets is empty.

Proof of Equation (EC.10). Next we show that Equation (EC.10) holds. Recall that f_T^- denotes an entire evolution of the system excluding the randomized decisions of the RB policy. Given the entire evolution of demands f_T^- , construct a decision tree based on the randomized decisions of the RB policy. The root node corresponding to period 1 contains the information set $f_1 = f_1^- \in f_T^-$. The tree is built in layers, each corresponding to a period, where the number of nodes in layer t is 2^{t-1} numbered $l = 1, \dots, 2^{t-1}$. In particular, a node l in period (layer) t corresponds to some information set $f_t \in \mathcal{F}_t$ which includes the realized reduced information set $f_t^- \subseteq f_T^-$, and the realized randomized decisions up to period $t-1$ of the RB policy. Therefore it is known whether under this state period t belongs to the set \mathcal{T}_{2M} or not.

The edges in the tree represent the different (randomized) decisions that the RB policy may make with their respective probabilities. Each path from the root to a specific node corresponds to a sequence of realized randomized ordering decisions made by the RB policy. For example, consider again some node l in period (layer) t in which the RB policy will order \tilde{q}_{tl}^{RB} units with probability p_{tl} and nothing with probability $1 - p_{tl}$; then the node l in period t (denoted by tl) will have two edges to two children nodes in the next period $t+1$ each containing its distinctive ordering information. Conceptually one can think about the decision tree as a collection of independent coins, each corresponding to a node in the tree. The coin corresponding to node l at layer (period) t has probability of success (ordering) p_{tl} .

Next we partition the nodes in the tree into *problematic* nodes (pn nodes), i.e., nodes that correspond to a pair (t, f_t) for which $t \in \mathcal{T}_{2M}$, and *non-problematic* nodes (nn nodes). An example of a general decision tree is illustrated in Figure EC.2.

Focus now on a specific time interval $[t_i, t_{i+1})$. Suppose we have constructed the tree from period 1 to T ; the number of nodes and paths are clearly finite (possibly exponential). Let the set \mathcal{G} to be the set of all possible outcomes of the randomized decisions in all nodes in layers within the interval $[1, t_i - 1]$ and in all the nn nodes within the interval $[1, T]$. In particular, each $g \in \mathcal{G}$ corresponds to a specific set of outcomes in all nodes in layers (periods) within the interval $[1, t_i - 1]$ and in all the nn nodes in the tree. Using the terminology of coins proposed before, g corresponds to the




Figure EC.2 An example of a general decision tree

outcome of the respective subset of coins corresponding to all nodes within $[1, t_i - 1]$ and all nn nodes within $[1, T]$.

Conditioning on some $g \in \mathcal{G}$ induces a path from the root of the tree (in period 1) up to the earliest pn node, say j , where j corresponds to the period (layer) of that node. Here we abuse the notation ignoring the index of the node within layer j . (Namely, the exact value will be je for some e .) It is straightforward to see that $j \geq t_i$. If j falls outside the interval $[t_i, t_{i+1})$, i.e., $j \geq t_{i+1}$, it follows that there are no pn nodes within the interval $[t_i, t_{i+1})$, and there is no borrowing over the interval. Assume now that j falls within the interval $[t_i, t_{i+1})$ (j can possibly be in period (layer) t_i). We will show that the expected borrowing does not exceed K . That is,

$$E \left[\sum_{s \in [j, t_{i+1}) \cup \mathcal{T}_{2M}} Z_s^{RB} \mid f_T^-, g \right] \leq K. \tag{EC.13}$$

The proof of Equation (EC.10) will then follow.

Recall that node j corresponds to some information set $f_j \in \mathcal{F}_j$. It follows that the starting inventory position x_j^{RB} and the corresponding *holding-cost- K* quantity \tilde{q}_j^{RB} are known deterministically. Conditioning on g , the only uncertainty in the evolution of the system depends on the randomized decisions made in pn nodes within $[j, t_{i+1})$. Consider the sub-tree induced by conditioning on g . The




Figure EC.3 An example of a decision subtree: focus on the interval $[t_i, t_{i+1})$ and some $g \in \mathcal{G}$, j is the earliest period in which a problematic node (pn) occurs. According to g , there are two possible outcomes whenever a problematic node (pn) is reached, and there is only one possible outcome whenever a non-problematic node (nn) is reached. If a problematic node (pn) orders, there will not be further borrowing until the next order of OPT in period t_{i+1} .

non-problematic nodes (nn nodes) in the sub-tree have only one outgoing edge that corresponds to the decision (order/no-order) specified by g to that node. The problematic nodes (pn nodes) have two outgoing edges corresponding to the order/no-order decisions, respectively. (Recall that g does not specify the decisions in these nodes.) Moreover, each pn node $s \in [j, t_{i+1})$ is associated with the probability p_s of ordering. (We again abuse the notation introduced before and omit the index e of the node within the layer/period.) An example of a decision subtree specified by some $g \in \mathcal{G}$ is illustrated in Figure EC.3. Any sequence of randomized outcomes corresponding to the decisions in the pn nodes induces a path of evolution of the system. The resulting cumulative borrowing from the bank account A , corresponding to this path, is equal to K times the sum of probabilities associated with the pn nodes in this path. (For each pn node s in the path, the borrowing is equal to $p_s K = z_s$.)

Next we claim that the sub-tree defined above includes at most one pn node in each layer (period).

This follows from the fact that any path between two pn nodes r, s such that $j \leq r < s < t_{i+1}$ in the tree includes only no-ordering edges of pn nodes. To see why the latter is true, observe that if an order is placed by the RB policy in a pn node, the resulting inventory position of the RB policy is higher than OPT . Since both policies face the same sequence of demands, the RB policy will not have higher inventory position than OPT at least until the next order placed by OPT . This excludes the existence of pn nodes in subsequent periods until OPT places another order, i.e., beyond period $t_{i+1} - 1$.

In light of the latter observation, we re-number all the pn nodes in the sub-tree as $1, 2, \dots, M$ (where 1 corresponds to j , specified before). Moreover, it follows that the probability to arrive at node $m = 1, \dots, M$ and borrow $p_m K$ is equal to $\prod_{s=1}^{m-1} (1 - p_s)$. (This probability corresponds to no-ordering decisions in all the pn nodes prior to m .) The total expected borrowing is then

$$K \cdot \left\{ p_1^2 + \sum_{m=2}^M \left\{ \left(\prod_{s=1}^{m-1} (1 - p_s) \right) p_m \left(\sum_{k=1}^m p_k \right) \right\} \right\}. \quad (\text{EC.14})$$

Observe that the probability to borrow exactly $K \cdot \sum_{k=1}^m p_k$ is equal to $\left(\prod_{s=1}^{m-1} (1 - p_s) \right) p_m$. Moreover, we have already shown that the expression in (EC.14) is bounded above by K (see Lemma 5). This concludes the proof of the lemma. \square

LEMMA 5. Let $\{p_l\}_{l=1}^{\infty}$ satisfy the condition $0 \leq p_l \leq 1$ for all l . Then the following inequality holds,

$$p_1^2 + \sum_{l=2}^{\infty} \left\{ \left(\prod_{s=1}^{l-1} (1 - p_s) \right) p_l \left(\sum_{k=1}^l p_k \right) \right\} \leq 1. \quad (\text{EC.15})$$

Proof of Lemma 5. We construct an increasing sequence $\{a_m\}$ where

$$a_m = p_1^2 + \sum_{l=2}^m \left\{ \left(\prod_{s=1}^{l-1} (1 - p_s) \right) p_l \left(\sum_{k=1}^l p_k \right) \right\}. \quad (\text{EC.16})$$

For each m , if we replace p_m by 1, we get

$$\bar{a}_m = p_1^2 + \sum_{l=2}^{m-1} \left\{ \left(\prod_{s=1}^{l-1} (1 - p_s) \right) p_l \left(\sum_{k=1}^l p_k \right) \right\} + \left(\prod_{s=1}^{m-1} (1 - p_s) \right) \left(1 + \sum_{k=1}^{m-1} p_k \right), \quad (\text{EC.17})$$

such that $a_m \leq \bar{a}_m$. Next we will show by induction that $\bar{a}_m \leq 1$ for all m from which the proof of

the lemma follows. It is straightforward to verify $\bar{a}_1, \bar{a}_2 \leq 1$. Assume that $\bar{a}_m \leq 1$ for some $m \in \mathbb{Z}^+$, we will show that $\bar{a}_{m+1} \leq 1$.

$$\begin{aligned}
\bar{a}_{m+1} &= p_1^2 + \sum_{l=2}^m \left\{ \left(\prod_{s=1}^{l-1} (1-p_s) \right) p_l \left(\sum_{k=1}^l p_k \right) \right\} + \left(\prod_{s=1}^m (1-p_s) \right) \left(1 + \sum_{k=1}^m p_k \right) \quad (\text{EC.18}) \\
&= a_{m-1} + \left(\prod_{s=1}^{m-1} (1-p_s) \right) p_m \left(\sum_{k=1}^m p_k \right) + \left(\prod_{s=1}^m (1-p_s) \right) \left(1 + \sum_{k=1}^m p_k \right) \\
&= a_{m-1} + \left(\prod_{s=1}^{m-1} (1-p_s) \right) \left[\left(1 + \sum_{k=1}^m p_k \right) (1-p_m) + p_m \sum_{k=1}^m p_k \right] \\
&= a_{m-1} + \left(\prod_{s=1}^{m-1} (1-p_s) \right) \left(1 + \sum_{k=1}^{m-1} p_k \right) = \bar{a}_m \leq 1.
\end{aligned}$$

Hence the claim follows by induction. \square

EC.2. Performance of the proposed algorithms

The first two columns specify the test instances, namely, fixed ordering cost K , per-unit holding cost h , per-unit backlogging cost p and demand rate vector λ . The third column shows the cost incurred by the optimal policy. The fourth column shows the optimal parameters of parametrized RB policy. The fifth column shows the cost incurred by the parameterized RB policy. The sixth column shows the cost ratio of the parameterized RB policy to the optimal policy. The seventh column shows the cost of unparameterized RB policy (i.e., the original policy without parameter optimization). The eighth columns shows the cost ratio of the unparameterized RB policy to the optimal policy.

(K,h,p)	Demands ($\lambda_0, \lambda_1, \lambda_2$)	Cost of <i>OPT</i>	Optimal ($\beta^*, \gamma^*, \eta^*$)	Cost of param. <i>RB</i>	Cost Ratio	Cost of unparam. <i>RB</i>	Cost Ratio
(0,1,9)	(4,1,4)	46.85	(* ,2,*)	49.18	1.0497	58.30	1.2444
(0,1,9)	(4,1,2)	46.39	(* ,2,*)	49.30	1.0627	55.24	1.1908
(0,1,9)	(4,1,1)	46.20	(* ,2,*)	47.81	1.0348	54.26	1.1745
(0,1,9)	(3,1,2)	41.02	(* ,2,*)	41.41	1.0095	49.40	1.2043
(0,1,9)	(2,1,3)	32.88	(* ,2,*)	34.42	1.0468	41.51	1.2625
(0,1,9)	(1,1,4)	24.74	(* ,2,*)	26.40	1.0671	31.40	1.2692
(5,1,9)	(4,1,1)	102.66	(0.2,2,9)	108.28	1.0547	135.37	1.3186
(5,1,9)	(1,1,4)	86.47	(0.2,2,9)	90.70	1.0489	128.70	1.4884
(5,1,1)	(4,1,1)	71.35	(0.4,1,1)	75.42	1.0570	84.13	1.1791
(100,1,9)	(5,1,0)	427.81	(0.9,* ,9)	451.68	1.0558	605.10	1.4144
(100,1,9)	(4,1,1)	424.81	(0.9,* ,9)	449.65	1.0585	601.29	1.4154
(100,1,9)	(3,1,2)	421.76	(0.9,* ,9)	443.12	1.0506	595.10	1.4110
(100,1,9)	(2,1,3)	418.63	(0.9,* ,9)	443.64	1.0597	611.48	1.4607
(100,1,9)	(1,1,4)	415.49	(0.8,* ,9)	437.36	1.0526	618.36	1.4883
(100,1,9)	(0,1,5)	412.29	(0.8,* ,9)	435.65	1.0567	593.88	1.4404

Table EC.1 Numerical results with lead time $L = 0$ and finite horizon $T = 12$.

(K,h,p)	Demands ($\lambda_0, \lambda_1, \lambda_2$)	Cost of <i>OPT</i>	Optimal ($\beta^*, \gamma^*, \eta^*$)	Cost of param. <i>RB</i>	Cost Ratio	Cost of unparam. <i>RB</i>	Cost Ratio
(0,1,9)	(4,1,4)	93.81	(* ,2,*)	98.32	1.0481	120.14	1.2807
(0,1,9)	(4,1,2)	88.27	(* ,2,*)	94.25	1.0677	108.24	1.2262
(0,1,9)	(4,1,1)	85.48	(* ,2,*)	90.21	1.0553	93.97	1.0993
(0,1,9)	(3,1,2)	80.04	(* ,2,*)	89.73	1.1211	90.40	1.1294
(0,1,9)	(2,1,3)	73.98	(* ,1.5,*)	84.42	1.1411	90.99	1.2625
(0,1,9)	(1,1,4)	70.96	(* ,1.5,*)	81.40	1.1471	87.60	1.2345
(5,1,9)	(4,1,1)	137.66	(0.2,2,9)	153.97	1.1185	161.10	1.1703
(5,1,9)	(1,1,4)	121.47	(0.2,2,9)	140.26	1.1525	148.47	1.2223
(5,1,1)	(4,1,1)	78.18	(0.4,1,1)	90.42	1.1566	97.47	1.2467
(100,1,9)	(5,1,0)	434.30	(0.9,* ,9)	479.03	1.1030	614.17	1.4142
(100,1,9)	(4,1,1)	431.87	(0.9,* ,9)	466.33	1.0798	611.96	1.4170
(100,1,9)	(3,1,2)	429.41	(0.9,* ,9)	453.24	1.0555	551.00	1.2832
(100,1,9)	(2,1,3)	426.86	(0.9,* ,9)	451.17	1.0570	644.13	1.5090
(100,1,9)	(1,1,4)	424.25	(0.9,* ,9)	466.43	1.0994	623.56	1.4698
(100,1,9)	(0,1,5)	421.56	(0.9,* ,9)	461.65	1.0951	595.40	1.4124

Table EC.2 Numerical results with lead time $L = 2$ and finite horizon $T = 12$.

(K, h, p)	Demands $(\lambda_0, \lambda_1, \lambda_2)$	Cost of OPT	Optimal $(\beta^*, \gamma^*, \eta^*)$	Cost of param. RB	Cost Ratio	Cost of unparam. RB	Cost Ratio
(0,1,9)	(4,1,4)	57.71	(* , 2 , *)	58.23	1.0090	61.92	1.0730
(0,1,9)	(4,1,2)	57.71	(* , 2 , *)	58.36	1.0113	60.94	1.0560
(0,1,9)	(4,1,1)	57.71	(* , 2 , *)	58.30	1.0102	60.38	1.0463
(0,1,9)	(3,1,2)	50.19	(* , 2 , *)	51.49	1.0259	53.62	1.0683
(0,1,9)	(2,1,3)	41.27	(* , 2 , *)	41.96	1.0167	43.63	1.0572
(0,1,9)	(1,1,4)	30.55	(* , 2 , *)	30.88	1.0108	31.66	1.0363
(5,1,9)	(4,1,1)	128.17	(0.2, 2, 9)	133.91	1.0448	166.10	1.2959
(5,1,9)	(1,1,4)	101.70	(0.2, 2, 9)	107.34	1.0555	148.85	1.4636
(5,1,1)	(4,1,1)	86.07	(0.4, 1, 1)	90.51	1.0516	104.24	1.2111
(100,1,9)	(5,1,0)	535.14	(1.1, * , 9)	566.23	1.0581	663.61	1.2401
(100,1,9)	(4,1,1)	533.51	(1.1, * , 9)	570.65	1.0696	659.29	1.2358
(100,1,9)	(3,1,2)	529.77	(1.1, * , 9)	566.09	1.0686	682.76	1.2888
(100,1,9)	(2,1,3)	523.94	(1.1, * , 9)	555.57	1.0604	729.15	1.3917
(100,1,9)	(1,1,4)	520.03	(1.0, * , 9)	550.36	1.0583	744.45	1.4316
(100,1,9)	(0,1,5)	516.05	(1.0, * , 9)	550.65	1.0670	711.22	1.3782

Table EC.3 Numerical results with lead time $L = 0$ and finite horizon $T = 15$.