European Journal of Operational Research 241 (2015) 381–390

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Production, Manufacturing and Logistics

Parametric replenishment policies for inventory systems with lost sales and fixed order cost

Marco Bijvank[a,*], Sandjai Bhulai[b], Woonghee Tim Huh[c]

[a] Haskayne School of Business, University of Calgary, 2500 University Dr. NW, Calgary, AB T2N 1N4, Canada
[b] Department of Mathematics, Faculty of Sciences, VU University Amsterdam, De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands
[c] Sauder School of Business, University of British Columbia, 2053 Main Mall, Vancouver, BC V6T 1Z2, Canada

ARTICLE INFO

Article history:
Received 19 August 2013
Accepted 15 September 2014
Available online 28 September 2014

Keywords:
Inventory management
Lost sales
Order cost
Replenishment policy

ABSTRACT

In this paper we consider a single-item inventory system with lost sales and fixed order cost. We numerically illustrate the lack of a clear structure in optimal replenishment policies for such systems. However, policies with a simple structure are preferred in practical settings. Examples of replenishment policies with a simple parametric description are the (s, S) policy and the (s, nQ) policy. Besides these known policies in literature, we propose a new type of replenishment policy. In our modified (s, S) policy we restrict the order size of the standard (s, S) policy to a maximum. This policy results in near-optimal costs. Furthermore, we derive heuristic procedures to set the inventory control parameters for this new replenishment policy. In our first approach we formulate closed-form expressions based on power approximations, whereas in our second approach we derive an approximation for the steady-state inventory distribution. As a result, the latter approach could be used for inventory systems with different objectives or service level constraints. The numerical experiments illustrate that the heuristic procedures result on average in 2.4 percent and 1.8 percent cost increases, respectively, compared to the optimal replenishment policy. Therefore, we conclude that the heuristic procedures are very effective to set the inventory control parameters.

1. Introduction

Two types of assumptions are commonly used to model the behavior of customers toward inventory stock outs; either a customer waits until a new order is delivered, or the customer does not buy the product. They are called backordering and lost sales, respectively. In contrast to industrial settings, excess demand is often lost in retail environments. Inventory models with a backorder assumption have received by far the greatest attention in inventory literature, possibly because of the simplicity of the optimal policy for such models with periodic reviews (Scarf 1960, Chap. 13; Karlin & Scarf, 1958, Chap. 10). In a periodic review model with fixed order cost, an optimal policy prescribes to raise the inventory position (inventory on hand plus inventory on order minus backorders) to an order-up-to level S when it falls down to or below a reorder level s at a review instant. This replenishment policy is abbreviated as the (s, S) policy. However, there is much less understanding of (near) optimal replenishment policies for inventory systems where excess demand is lost.

The lost-sales inventory systems with no fixed order cost have originally been formulated by Karlin and Scarf (1958) and Morton (1969). More recently, this has been reformulated with a new state definition by Zipkin (2008b), who uses it to prove that the minimal cost function is $L^\natural$-convex. This result implies that the optimal order quantities are monotone decreasing in the inventory position and the optimal order quantities are more sensitive to recent orders. Alternative replenishment policies are proposed by, e.g., Morton (1971), Levi, Janakiraman, and Nagarajan (2008), Huh, Janakiraman, Muckstadt, and Rusmevichientong (2009), Bijvank and Johansen (2012). Zipkin (2008a) has performed a numerical comparison of several of these policies. He concludes that myopic policies perform better than order-up-to policies in general. However, Huh et al. (2009) show that order-up-to policies are asymptotically optimal when the penalty cost for a lost-sales occurrence becomes high. Such policies are also popular in practice since they are fairly simple to implement, whereas myopic policies are less insightful and they require extra computational effort at each review instant. Furthermore, Johansen and Thorstenson (2008) introduce restricted order-up-to policies, in which they impose an upper limit on the order size. This policy results in near-optimal order quantities and it is easy to implement in practice. The goal of this paper is to extend these policies to include fixed order costs resulting in a new type of replenishment policy. We also

* Corresponding author. Tel.: +1 403 220 8453.
E-mail addresses: marco.bijvank@haskayne.ucalgary.ca (M. Bijvank), s.bhulai@vu.nl (S. Bhulai), tim.huh@sauder.ubc.ca (W. Tim Huh).

http://dx.doi.org/10.1016/j.ejor.2014.09.018
0377-2217/

discuss how the parameters in the new replenishment policy should be set.

In a lost-sales inventory system, if a fixed order cost is incurred with each order and the lead time is zero, then the (s, S) policy is proven to be optimal since the dynamics of a backorder system and a lost-sales system are essentially the same (see, e.g., Veinott & Wagner, 1965; Veinott, 1966; Shreve, 1976; Bensoussan, Crouhy, & Proth, 1983; Cheng & Sethi, 1999). More recently, Xu, Bisi, and Dada (2010) derive new bounds on the best reorder level and order-up-to level, whereas Li and Yu (2012) provide insights on the optimal policies when there are extra constraints imposed on the problem by showing that the objective function is quasiconcave. In case of positive lead times, there is no simple optimal replenishment policy. Nahmias (1979) is one of the first authors to study such an inventory system with positive lead times. Hill and Johansen (2006) show by means of a numerical example that the optimal policy is neither an (s, S) policy nor an (s, Q) policy. In the latter policy a fixed number of Q units are ordered when the inventory position reaches (or drops below) reorder level s at a review instant. This policy is also investigated by Johansen and Hill (2000) under the assumption that at most one order can be outstanding at any time. Fractional lead times are considered by Kapalka, Katircioglu, and Puterman (1999) and Chiang (2007), where the lead time is assumed to be shorter than the review period. When there is no assumption on the number of outstanding orders or on the lead time, we are the first authors to our knowledge to provide numerical results on the performance of (s, S) and (s, nQ) policies in comparison to the optimal replenishment policy for inventory systems with lost sales. As our analysis will show, such parametric policies perform reasonably well compared to the optimal policy (which has no clear structure).

Another contribution to the literature is the proposal of a new and more general type of replenishment policy, which we obtain by imposing an upper bound on the order size in the (s, S) policy. Therefore, we call it the modified (s, S) policy. Since (s, S) policies are already very common in many practical settings (Caplin & Leahy, 2010), this new policy can easily be applied. We numerically show that this policy significantly outperforms the parametric policies known in literature.

In our final contribution to the literature we propose different heuristic procedures to set the inventory control parameters for the new modified (s, S) policy. We pursue two different approaches that are known in literature to determine inventory control parameters. In the first approach, we use power approximations to formulate closed-form expressions to set the parameter values for an (s, S) policy. Ehrhardt (1979) and Ehrhardt and Mosier (1984) have developed a similar approach when excess demand is backordered instead of lost. In the second approach, we derive expressions to approximate the entire steady-state distribution of the inventory levels to compute the performance measures of interest, such as the long-run expected total costs and service level. Subsequently, these approximate expressions are used to find the best inventory control parameter values. This extends the approach taken by Bijvank and Johansen (2012), where no fixed order costs are involved. The advantage of this latter approach is that it can be used for inventory systems with a different objective function or with service level constraints.

In summary, it is well known that simple replenishment policies with a parametric structure (such as (s, S) and (s, Q) policies) are not optimal to control inventory levels when excess demand is lost. However, such policies are commonly used (Caplin & Leahy, 2010). Our contributions are the following: (i) to provide an analysis which illustrates that parametric replenishment policies perform close to optimal for lost-sales inventory systems, (ii) to propose a new (parametric) replenishment policy that performs near optimal under the lost-sales assumption, and (iii) to develop heuristic procedures to set the parameter values of this new policy. In particular, we provide a numerical study when (not) to use certain policies or heuristic procedures.

The organization of this paper is as follows. In Section 2 we introduce all notation and assumptions to model the inventory system. In Section 3 we present an analytical framework to study optimal and parametric replenishment policies. The computational effort to analyze the performance of lost-sales inventory systems can be quite excessive, especially when the lead time is relatively long. Therefore, we propose heuristic procedures in Section 4 to determine the inventory control parameters for the modified (s, S) policy. In Section 5 we compute and compare the performance of all replenishment policies and the heuristic procedures. Section 6 contains our concluding remarks.

2. Notation and assumptions

In this section, we provide our notation and assumptions. The time between two reviews is called a review period. Its length is denoted by R. At each review instant the inventory status is checked to decide whether a new order should be placed. A new order arrives after a constant lead time L. Consequently, $L = lR$ where l is the number of review periods between order placement and order delivery. For ease of notation we assume l to be an integer value (i.e., the lead time is an integral multiple of the review period length). This assumption is relaxed in Appendix B.

For our exposition, we assume demand to be discrete as well as independent and identically distributed over time. The demand during a time period of length τ is a stochastic random variable D_τ with probability mass function $g_\tau(\cdot)$, mean $E[D_\tau] = \mu_\tau$ and variance $\text{Var}[D_\tau] = \sigma_\tau^2$. We also define

$$\mathcal{G}_\tau^0(i) = Pr(D_\tau < i) = \sum_{d=0}^{i-1} g_\tau(d),$$

$$\mathcal{G}_\tau^1(i) = E[(i - D_\tau)^+] = \sum_{d=1}^{i} \mathcal{G}_\tau^0(d),$$

with $(A)^+ = \max\{A, 0\}$.

We denote the set of non-negative integers by $\mathbb{N}_0$, and the set of all integers between m and n by $\mathbb{N}_{m,n} = \{i \in \mathbb{N}_0 \,|m \le i \le n\} = \{m, m+1, \ldots, n\}$. Furthermore, $\mathbb{N}_{m,n}^l$ is defined as the l-fold Cartesian product of $\mathbb{N}_{m,n}$.

In our Markov chain description of the inventory system, the state at a review instant after order delivery but before ordering is denoted by $(i, \mathbf{y})$, where i is the inventory on hand and $\mathbf{y}$ is a vector with components y_k, $k \in \mathbb{N}_{1,l-1}$. Component y_k is the quantity ordered $l - k$ review periods before, which is to be delivered after k review periods. We let $\mathbf{F}(\mathbf{y})$ denote the vector obtained from $\mathbf{y}$ by removing its first component. The state of the Markov chain at the next review instant is denoted by $(j, \mathbf{z})$, where j is the updated on-hand inventory level and $\mathbf{z} = (\mathbf{F}(\mathbf{y}), y_l)$. The component y_l represents the quantity ordered at the current review moment. The replenishment policy prescribes how y_l depends on $(i, y_1, \ldots, y_{l-1})$. We denote the actual demand during a review period by d. Hence, $j = (i - d)^+ + y_1$. Consequently, the state space is an l-dimensional vector where the first component specifies the on-hand inventory level and the remaining $l - 1$ components represent the orders outstanding.

To complete the Markov chain description, we have to specify the one-step transition probabilities between the different states of the inventory system. The transition probabilities from state $(i, \mathbf{y})$ to state $(j, \mathbf{F}(\mathbf{y}), y_l)$ are denoted by $P_{(i,\mathbf{y}),(j,\mathbf{F}(\mathbf{y}),y_l)}$. The choice of y_l depends on the replenishment policy, which we discuss further in Section 3.

To compute the expected total costs over a review period, we express the expected holding and penalty costs incurred over the review period by $c(i)$ when the on-hand inventory level equals i units at the beginning of this period. Let h denote the unit holding cost per unit time and p the unit penalty cost for each lost demand. We assume linear and proportional holding and penalty costs. Consequently,

$$c(i) = hE[(i - D_R)^+] + pE[(D_R - i)^+], \tag{1}$$

where $E[(D_R - i)^+] = E[D_R] - i + E[(i - D_R)^+] = \mu_R - i + \mathcal{G}_R^1(i)$. Thus, $c(i) = p(\mu_R - i) + (p + h)\mathcal{G}_R^1$. Other cost structures are considered in Appendix B .

The performance measure of interest for this inventory system is the long-run expected total costs, which can be computed by value iteration (Puterman, 2005). Let $V_n(i, \mathbf{y})$ denote the expected total costs incurred over n review periods when the system starts in state $(i, \mathbf{y})$ and the system incurs no costs at and after the nth review period. Consequently, $V_0(i, \mathbf{y}) = 0$ and, recursively for $n \geq 1$, we define

$$V_n(i, \mathbf{y}) = c(i) + \sum_{j, y_l} P_{(i,\mathbf{y}),(j,\mathbf{F}(\mathbf{y}),y_l)}\{K\delta(y_l) + V_{n-1}(j, \mathbf{F}(\mathbf{y}), y_l)\}, \tag{2}$$

where K denotes the fixed order cost, and the indicator function $\delta(y_l)$ is one when $y_l > 0$ and zero when $y_l = 0$. We are interested in the long-run performance. The value-iteration algorithm with accuracy number ε repeats to increase n by one and compute the value functions in Eq. (2) until $M_n - m_n < \varepsilon$, where

$$m_n = \min_{(i,\mathbf{y})}\{V_n(i, \mathbf{y}) - V_{n-1}(i, \mathbf{y})\},$$

$$M_n = \max_{(i,\mathbf{y})}\{V_n(i, \mathbf{y}) - V_{n-1}(i, \mathbf{y})\}.$$

When this value-iteration algorithm is stopped after the nth iteration, then $(m_n + M_n)/2$ cannot deviate more than 100ε percent from the long-run average costs per review period. The numerical results reported in Section 5 are computed with $\varepsilon = 1\text{E}{-}4$. In the next section we specify the model for different replenishment policies.

3. Optimal and parametric policies

As mentioned in the previous section, the actual performance of the inventory system depends on the replenishment policy. In this section, we formulate the optimal policy as well as alternative replenishment policies that have a parametric structure such that they are easy to implement in practice. Only minor modifications are necessary to model these alternative policies compared to the optimal policy. For each policy, we describe the order quantities and recursive expression to perform the value-iteration algorithm as presented at the end of Section 2. More details on the state space and the transition probabilities for the Markov decision model are described in Appendix A.

3.1. Optimal policy

To find the optimal replenishment policy and the associated expected total costs, we model the system as a Markov chain with an infinite state space. More specifically, the state space at a review instant equals $\mathcal{S} = \{(i, \mathbf{y}) \in \mathbb{N}_0^l\}$. The one-step transition probabilities between the states at subsequent reviews depend on the entire state description. Therefore, we use the following backward recursion formulation to solve this Markov decision problem

$$V_n(i, \mathbf{y}) = c(i) + \min_{y_l \geq 0}\left\{K\delta(y_l) + \sum_{d=0}^{i-1} g_R(d)V_{n-1}(i - d + y_1, \mathbf{F}(\mathbf{y}), y_l) + \left(1 - \mathcal{G}_R^0(i)\right)V_{n-1}(y_1, \mathbf{F}(\mathbf{y}), y_l)\right\}, \quad \text{for } n \geq 1 \tag{3}$$

and $V_0(i, \mathbf{y}) = 0$.

To illustrate the structure of an optimal policy, we consider the following example in which $R = 1$, $L = 2$, $h = 1$, $p = 14$ and $K = 5$. Note that $l = 2$ in this example. Consequently, at most one order is outstanding at a review instant and $\mathbf{y} = \{y_1\}$. The customer's demand is assumed to follow a Poisson process with mean $\mu_R = 5$. Table 1 presents the optimal order quantities y_l' when the on-hand inventory level equals i units at a review instant (first column) and the inventory on order is y_1 (first row). The corresponding expected total costs are 11.46. This example illustrates two interesting aspects. First, the optimal policy prescribes never to order when the inventory position equals 18 units or more (this corresponds to a reorder level of 17 units). Second, the optimal order quantities are not always decreasing in the inventory position. This is in contrast to lost-sales inventory systems with continuous reviews, where the rate of decrease is always less than or equal to one unit (Johansen & Thorstenson, 1993).

Table 1
The optimal order quantities y_l' when the on-hand inventory level equals i and the outstanding order has size y_1 at a review instant when $R = 1, L = 2, h = 1, p = 14, K = 5$ and demand has a Poisson distribution with mean 5. The highlighted order quantities coincide with the best (s, S) policy (light gray) and the best (s, nQ) policy (dark gray).

i \ y_1	0	1	2	3	4	5	6	7	8	9	10	11
0	8	8	8	8	13	13	12	12	11	10	9	8
1	8	8	8	8	13	13	12	12	11	10	9	8
2	8	8	8	8	13	13	12	12	11	10	9	8
3	8	8	8	8	13	13	12	12	11	10	9	8
4	8	8	8	13	13	12	12	11	11	10	9	0
5	8	8	8	13	13	12	12	11	10	10	9	0
6	8	8	13	13	12	12	12	11	10	9	8	0
7	8	13	13	12	12	12	11	6	6	6	0	0
8	13	13	13	12	12	11	7	6	6	0	0	0
9	13	13	12	12	7	7	7	7	6	0	0	0
10	13	12	12	7	7	7	7	6	0	0	0	0
11	12	12	7	7	7	7	6	0	0	0	0	0
12	12	7	7	7	7	6	0	0	0	0	0	0
12	7	7	7	7	6	0	0	0	0	0	0	0
14	7	7	7	6	0	0	0	0	0	0	0	0
15	7	7	6	0	0	0	0	0	0	0	0	0
16	7	6	0	0	0	0	0	0	0	0	0	0
17	6	0	0	0	0	0	0	0	0	0	0	0

3.2. The (s, S) policy

We now consider the (s, S) policy since such a policy is optimal when excess demand is backordered instead of lost. In this policy, the order size is such that the inventory position is raised to level S when it is at or below reorder level s. Hence,

$$y_l = \begin{cases} 0, & \text{if } i + \sum_{k=1}^{l-1} y_k > s, \\ S - i - \sum_{k=1}^{l-1} y_k, & \text{otherwise.} \end{cases} \tag{4}$$

The backward recursion of Eq. (2) equals

$$V_n(i, \mathbf{y}) = c(i) + K\delta(y_l) + \sum_{d=0}^{i-1} g_R(d)V_{n-1}(i - d + y_1, \mathbf{F}(\mathbf{y}), y_l) + \left(1 - \mathcal{G}_R^0(i)\right)V_{n-1}(y_1, \mathbf{F}(\mathbf{y}), y_l), \quad \text{for } n \geq 1. \tag{5}$$

The values of reorder level s and order-up-to level S have to be found such that the expected total costs are minimized. Let us denote these values by $\bar{s}$ and $\bar{S}$, respectively. In the example of Table 1, this corresponds to $\bar{s} = 17$ and $\bar{S} = 23$ (the order quantities in light gray) with expected total costs equal to 11.62 (1.35 percent cost increase compared to the optimal policy).

3.3. The (s, nQ) policy

Next, we consider the (s, nQ) policy with fixed batch sizes. In this policy, the minimal integer multiple of Q is ordered such that the inventory position exceeds s after ordering. Hence,

$$y_l = \max\left\{Q \times \left\lceil \left(s + 1 - \left(i + \sum_{k=1}^{l-1} y_k\right)\right)^+ \Big/ Q \right\rceil, 0\right\}, \tag{6}$$

where $\lceil x \rceil$ denotes the smallest integer larger than or equal to $x \in \mathbb{R}$. Note that if $Q > s$, then the order size is either zero or equal to Q units. Eq. (5) can be used for the value function, where y_l is given by Eq. (6). For the (s, nQ) replenishment policy, the values of s and Q that minimize the expected total costs are denoted by $\hat{s}$ and $\hat{Q}$, respectively. In the example of Table 1, this corresponds to $\hat{s} = 17$ and $\hat{Q} = 7$ (the order quantities in dark gray) with expected total costs equal to 11.56 (0.80 percent cost increase compared to the optimal policy).

3.4. The modified (s, S) policy

When we consider the optimal replenishment policy in Table 1, it is clear that the policy is neither an (s, S) policy nor an (s, nQ) policy. However, some states in the optimal policy comply with either one of the two policies. Therefore, we propose a modified (s, S) policy with reorder level s, order-up-to level S and upper limit q on the order quantity. For instance, the order quantities that are light gray in the lower half of Table 1 and the order quantities that are dark gray correspond to the best modified (s, S) policy where the reorder level is 17, the order-up-to level equals 23 and the maximum order size is 7. More formal, this policy prescribes to issue a replenishment order at a review instant where the size of the order equals

$$y_l = \begin{cases} 0, & \text{if } i + \sum_{k=1}^{l-1} y_k > s, \\ S - i - \sum_{k=1}^{l-1} y_k, & \text{if } S - q < i + \sum_{k=1}^{l-1} y_k \le s, \\ q, & \text{otherwise.} \end{cases} \tag{7}$$

Note that when $q \ge S$ the policy equals the standard (s, S) policy, and when $S \ge s + q > 2s$ it corresponds to the (s, nQ) policy with $Q = q$ and $n = 1$. The latter statement is true since the order size equals q units when $S - s \ge q$ and the inventory position after order placement exceeds the reorder level when $q > s$. Consequently, the modified (s, S) policy is a generalization of both (s, S) and (s, Q) policies.

Again, Eq. (5) can be used to express the value function, where y_l is represented by Eq. (7). For the modified (s, S) policy, the values of s, S, and q that minimize the expected total costs are denoted by s^*, S^* and q^*, respectively. In the example of Table 1, this corresponds to $s^* = 17$, $S^* = 23$ and $q^* = 7$ with expected total costs equal to 11.50 (0.30 percent cost increase compared to the optimal policy).

4. Heuristic procedures

It requires numerical search procedures to find optimal values of the inventory control parameters for the policies described in the previous section. This can require extensive computational effort, especially for long lead times. Johansen and Hill (2000) propose a heuristic procedure to determine near-optimal values of s and Q for the (s, nQ) policy. No procedure is available for (s, S) replenishment policies in a lost-sales context. Therefore, the goal of this section is to find good heuristic procedures to set near-optimal values for the control parameters of the (s, S) policy. At the end of this section, we also discuss how these heuristic procedures can be used for the modified (s, S) policy. The performance of all heuristic procedures is tested in Section 5.

We propose two different heuristic procedures to set the replenishment control parameters of the (s, S) policy. First, we use regression methods to formulate power approximations. Second, we derive expressions to approximate the steady-state distribution of the inventory levels that are subsequently used to approximate the performance measures of interest, such as the expected total costs and service level.

4.1. Power approximation based on regression

Power approximations have been first proposed in the inventory literature by Ehrhardt (1979) in the context of computing near-optimal (s, S) values for inventory systems with backordering and fixed ordering cost. In this approach regression models were fitted to a grid of 288 parameter settings to express the reorder level by s_p and the difference $S - s$ by Δ_p. Ehrhardt and Mosier (1984) revised the approach in which the same numerical analysis is performed to fit a power series of the form:

$$\Delta_p = c_0 \mu_R^{c_1} (K/h)^{c_2} (1 + \sigma_{L+R}^2 / \mu_R^2)^{c_3}, \tag{8}$$

$$s_p = c_4 \mu_{L+R} + \sigma_{L+R}(c_5/z + c_6 + c_7 z), \tag{9}$$

where c_i, $i = 0, 1, \ldots, 7$ are constants to be fitted to optimal policy data and z is given by $z = \sqrt{\Delta_p/(\sigma_{L+R} p/h)}$. Ehrhardt and Mosier (1984) impose $c_2 = 1 - c_1$ for rescalability reasons.

We use the same functional forms and a similar numerical setting as used by Ehrhardt and Mosier (1984) to perform the numerical analysis and set c_i. Demand is generated from two types of distributions: Poisson and negative binomial with variance-to-mean (VTM) ratios of 3 or 5. The average demand per review period equals 1, 2, 4 or 8. We consider systems where the lead time equals 1 or 3 review periods. The value of the unit holding cost is a redundant parameter which is set at unity. The unit penalty cost is 4, 9, 24 or 99, and the fixed order cost is 16, 32 or 64. All combinations of these parameter settings result in 288 training instances. We have calculated the best reorder level and order-up-to level for each problem instance by full enumeration. These values are used as data for least-square regression fits to determine the value of the constants c_i in Eqs. (8) and (9), which result in $c_0 = 1.259, c_1 = 0.478, c_2 = 0.522, c_3 = 0.081, c_4 = 1.027, c_5 = 0.185, c_6 = 0.713$ and $c_7 = -2.536$. When Δ_p/μ_R is sufficiently large (i.e., greater than 1.5), we let $s = s_p$ and $S = s_p + \Delta_p$. Otherwise, we set

$$S_0 = \underset{S \ge 0}{\operatorname{argmin}} \left\{ \sum_{i=1}^{S} c(i) g_L(S - i) + c(0) \left[1 - \mathcal{G}_L^0(S) \right] \right\}, \tag{10}$$

which corresponds to the best order-up-to level for a backorder system with $s = S - 1$, and let $s = \min\{s_p, S_0 - 1\}$ and $S = \min\{s_p + \Delta_p, S_0\}$.

4.2. Steady-state approximations

The basic motivation behind the second heuristic procedure is to modify the demand distribution with an adjustment factor for lost demands in comparison to a backorder system. This is an extension of the approach taken by Bijvank and Johansen (2012) for a lost-sales inventory system when no fixed order costs are involved. This heuristic procedure consists of several steps. First, we derive approximate expressions for the equilibrium distribution of the inventory position and on-hand inventory level. Next, we use these distributions to develop closed-form expressions to analyze the performance of the inventory system. In the final step, the value of the adjustment factor is derived.

In this heuristic procedure, we relate the on-hand inventory level to the inventory position at a review instant at time T, denoted by $IL(T)$ and $IP(T)$, respectively. A superscript '+' after IL is used to indicate the inventory level after an order has been delivered at time T, whereas a '+' after IP denotes the inventory position after ordering at time T. When this is not the case, it is denoted by superscript '−'. Furthermore, $D(T, T + L)$ denotes the demand during the time period $[T, T + L)$. Note that all the orders outstanding at time T are delivered before or at time $T + L$. Thus all the units in the inventory position $IP^+(T)$ contribute to the inventory level $IL^+(T + L)$. The demand in period $[T, T + L)$, however, depletes this inventory level. In the backorder model, we have the following relationship,

$$IL^+(T + L) = IP^+(T) - D(T, T + L), \tag{11}$$

whereas in the lost-sales model this is not true since only the satisfied demand should be subtracted from the inventory position. Next, we

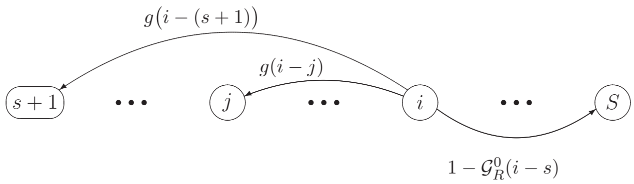

Fig. 1. The one-step transitions including their probabilities of the inventory position after ordering at a review in the backorder model.

use this observation to derive approximate expressions to describe the steady-state behavior of IP(T) and IL(T) for the lost-sales model, denoted by IP and IL, respectively.

In the (s, S) policy, when the inventory position reaches or drops below reorder level s, an order is placed at the next review instant such that the inventory position is raised to S. Hence, the inventory position is always between $s+1$ and S after ordering at a review instant, i.e., $\text{IP}^+(T) \in \mathbb{N}_{s+1,S}$. To find the steady-state distribution of the inventory position after ordering at time T, we first consider the backorder model. When excess demand is backordered,

$$\text{IP}^-(T+R) = \text{IP}^+(T) - D(T, T+R), \tag{12}$$

and

$$\text{IP}^+(T+R) = \begin{cases} S, & \text{if } \text{IP}^-(T+R) \le s, \\ \text{IP}^-(T+R), & \text{otherwise.} \end{cases} \tag{13}$$

Consequently, the inventory position after ordering in the backorder model can be modeled as a one-dimensional Markov chain with state space $\mathbb{N}_{s+1,S}$. The one-step transitions between the states in this Markov chain are graphically represented in Fig. 1. The steady-state probabilities for $\text{IP}^+(T)$ are denoted by π^+_{IP} and given by the solution to

$$\begin{aligned} \pi^+_{\text{IP}}(j) &= \sum_{i=j}^{S} g_R(i-j)\pi^+_{\text{IP}}(i), \quad j = s+1, \ldots, S-1 \\ \pi^+_{\text{IP}}(S) &= g_R(0)\pi^+_{\text{IP}}(S) + \sum_{i=s+1}^{S} \left(1 - \mathcal{G}_R^0(i-s)\right)\pi^+_{\text{IP}}(i) \end{aligned} \tag{14}$$

$$\sum_{j=s+1}^{S} \pi^+_{\text{IP}}(j) = 1.$$

These expressions for the backorder model are used to derive approximations for the steady-state distribution of the inventory position IP^+ in the lost-sales model. We adjust the distribution for the demand during the review period in Eq. (14) with a factor $\tilde{c}_{s,S}$ to account for the fact that the inventory position only decreases when demand is satisfied and actually depletes the on-hand inventory level. This means that when the demand is relatively high during a certain time period, it is less likely that all demand can be satisfied (i.e., the probability that the inventory level will decrease with the same amount as the demand should be lower than the probability according to the demand distribution). Similarly, when the demand is relatively small during a certain period, the probability that the inventory level decreases with the same amount is close to the corresponding demand probability. As a result, the average inventory position will be higher compared to the same inventory system when excess demand is backordered. Hence, for the lost-sales model

$$\tilde{\pi}^+_{\text{IP}}(j) = \sum_{i=j}^{S} \tilde{c}_{s,S} g_R(i-j)\tilde{\pi}^+_{\text{IP}}(i), \quad j = s+1, \ldots, S-1 \tag{15}$$

$$\tilde{\pi}^+_{\text{IP}}(S) = \tilde{c}_{s,S} g_R(0)\tilde{\pi}^+_{\text{IP}}(S) + \sum_{i=s+1}^{S} \left(1 - \tilde{c}_{s,S}\mathcal{G}_R^0(i-s)\right)\tilde{\pi}^+_{\text{IP}}(i) \tag{16}$$

$$\sum_{j=s+1}^{S} \tilde{\pi}^+_{\text{IP}}(j) = 1, \tag{17}$$

where we denote the approximate steady-state distribution of IP^+ for the lost-sales system by $\tilde{\pi}^+_{\text{IP}}$. The solution to this set of equations is given by

$$\tilde{\pi}^+_{\text{IP}}(j) = \frac{f(S-j)}{\sum_{i=s+1}^{S} f(S-i)}, \qquad j = s+1, \ldots, S, \tag{18}$$

where

$$f(j) = \begin{cases} 1, & \text{if } j = 0, \\ \sum_{i=1}^{j} \alpha(i) f(j-i), & \text{if } j > 0, \end{cases}$$

and

$$\alpha(j) = \frac{\tilde{c}_{s,S} g_R(j)}{1 - \tilde{c}_{s,S} g_R(0)}.$$

The next step is to relate the inventory position to the on-hand inventory level. The approximate steady-state distribution of IL^+ can be derived conditionally on IP^+ similar to Eq. (11)

$$\tilde{\pi}^+_{\text{IL}}(j) = \begin{cases} \sum_{i=j}^{S} \tilde{c}_{s,S} g_L(i-j)\tilde{\pi}^+_{\text{IP}}(i), & \text{if } j > 0, \\ \sum_{i=j}^{S} (1 - \tilde{c}_{s,S}\mathcal{G}_L^0(i))\tilde{\pi}^+_{\text{IP}}(i), & \text{if } j = 0. \end{cases} \tag{19}$$

Similarly, for IL^- by combining Eqs. (11) and (12)

$$\tilde{\pi}^-_{\text{IL}}(j) = \begin{cases} \sum_{i=j}^{S} \tilde{c}_{s,S} g_{L+R}(i-j)\tilde{\pi}^+_{\text{IP}}(i), & \text{if } j > 0, \\ \sum_{i=j}^{S} (1 - \tilde{c}_{s,S}\mathcal{G}_{L+R}^0(i))\, \tilde{\pi}^+_{\text{IP}}(i), & \text{if } j = 0. \end{cases} \tag{20}$$

We use these closed-form expressions to approximate the performance of the lost-sales inventory system with an (s, S) policy. For instance, the expected on-hand inventory level after order delivery is approximated by

$$\begin{aligned} \tilde{\text{IL}}^+_{s,S} &= \sum_{i=0}^{S} i\tilde{\pi}^+_{\text{IL}}(i) = \sum_{i=1}^{S}\sum_{j=i}^{S} i\tilde{c}_{s,S} g_L(j-i)\tilde{\pi}^+_{\text{IP}}(j) \\ &= \sum_{j=0}^{S} \tilde{c}_{s,S} E[(j-D_L)^+]\tilde{\pi}^+_{\text{IP}}(j) = \sum_{j=0}^{S} \tilde{c}_{s,S}\mathcal{G}_L^1(j)\tilde{\pi}^+_{\text{IP}}(j), \end{aligned}$$

where the first and third equalities follow from the definition of an expected value, the second equality follows from Eq. (19), and the last equality follows from the definition of $\mathcal{G}_L^1(\cdot)$.

Similarly, the approximation for the on-hand inventory level before order delivery equals

$$\tilde{\text{IL}}^-_{s,S} = \sum_{i=0}^{S} i\tilde{\pi}^-_{\text{IL}}(i) = \sum_{j=0}^{S} \tilde{c}_{s,S}\mathcal{G}_{L+R}^1(j)\tilde{\pi}^+_{\text{IP}}(j). \tag{21}$$

The difference between the two average inventory levels expresses the increase of the on-hand inventory level due to an order. Hence, the approximate average order size equals $\tilde{\text{IL}}^+_{s,S} - \tilde{\text{IL}}^-_{s,S}$. Consequently, the fraction of demand lost can be approximated by

$$\tilde{A}_{s,S} = 1 - \frac{\tilde{\text{IL}}^+_{s,S} - \tilde{\text{IL}}^-_{s,S}}{\mu_R} = 1 - \frac{\sum_{j=0}^{S} \tilde{c}_{s,S}\left[\mathcal{G}_L^1(j) - \mathcal{G}_{L+R}^1(j)\right]\tilde{\pi}^+_{\text{IP}}(j)}{\mu_R}. \tag{22}$$

Note that Eq. (22) is the same as

$$\tilde{A}_{s,S} = \sum_{j=0}^{S} \tilde{\pi}^+_{\text{IL}}(j)\frac{E[(D_R - j)^+]}{E[D_R]}. \tag{23}$$

Finally we approximate the long-run expected total costs by

$$\begin{aligned}\tilde{C}_{s,S} &= K\sum_{i=0}^{s}\tilde{\pi}_{\mathrm{IP}}^{-}(i) + \sum_{i=0}^{S}c(i)\tilde{\pi}_{\mathrm{IL}}^{+}(i)\\ &= K\left[\sum_{i=0}^{S}\left[1-\tilde{c}_{s,S}\mathcal{G}_R^0(i)\right]\tilde{\pi}_{\mathrm{IP}}^{+}(i) + \sum_{j=1}^{s}\sum_{i=j}^{S}\tilde{c}_{s,S}g_R(i-j)\tilde{\pi}_{\mathrm{IP}}^{+}(i)\right]\\ &\quad + \sum_{i=0}^{S}c(i)\tilde{\pi}_{\mathrm{IL}}^{+}(i)\\ &= K\sum_{i=s+1}^{S}\left[1-\tilde{c}_{s,S}\mathcal{G}_R^0(i-s)\right]\tilde{\pi}_{\mathrm{IP}}^{+}(i) + \sum_{i=0}^{S}c(i)\tilde{\pi}_{\mathrm{IL}}^{+}(i),\end{aligned} \tag{24}$$

where $c(i)$ is given by Eq. (1), and the third equality follows after rewriting the terms since many of them cancel out.

The final step of this heuristic procedure is to determine $\tilde{c}_{s,S}$. Note that there is not a unique value of $\tilde{c}_{s,S}$. Therefore, we propose two procedures to set the value of $\tilde{c}_{s,S}$:

(a) The first approach is the most simple, just set $\tilde{c}_{s,S} = 1$. This means that no adjustments are made for lost-sales. However, this does not correspond to a backorder model (in contrast to the procedure by Bijvank and Johansen (2012) when no fixed order costs are involved), since the inventory levels cannot be negative in Eqs. (19)–(24).

(b) The second approach to set the value of $\tilde{c}_{s,S}$ is based on the average order size. The average order size equals the expected increase in the on-hand inventory level due to an order delivery ($\tilde{\mathrm{IL}}_{s,S}^{+} - \tilde{\mathrm{IL}}_{s,S}^{-}$), but also the expected increase in the inventory position due to order placement at a review instant ($\tilde{\mathrm{IP}}_{s,S}^{+} - \tilde{\mathrm{IP}}_{s,S}^{-}$). In this approach to determine the value of $\tilde{c}_{s,S}$, we equate both amounts such that the following equality should be satisfied:

$$\sum_{i=s+1}^{S}\tilde{c}_{s,S}\left[\mathcal{G}_L^1(i) - \mathcal{G}_{L+R}^1(i)\right]\tilde{\pi}_{\mathrm{IP}}^{+}(i) = \sum_{i=s+1}^{S}\left[i - \tilde{c}_{s,S}\mathcal{G}_R^1(i)\right]\tilde{\pi}_{\mathrm{IP}}^{+}(i). \tag{25}$$

The left side expresses the difference $\tilde{\mathrm{IL}}_{s,S}^{+} - \tilde{\mathrm{IL}}_{s,S}^{-}$, whereas the right side expresses $\tilde{\mathrm{IP}}_{s,S}^{+} - \tilde{\mathrm{IP}}_{s,S}^{-}$. Note that $\tilde{\pi}_{\mathrm{IP}}^{+}(j)$ depends on $\tilde{c}_{s,S}$. Therefore, there is no explicit expression for $\tilde{c}_{s,S}$. With a bisection method we can find the value of $\tilde{c}_{s,S}$ such that Eq. (25) is satisfied.

As a result, the expected total costs are approximated by Eq. (24) for given values of s and S. With a numerical search procedure it is easy to find those values that minimize $\tilde{C}_{s,S}$.

4.3. Modified (s, S) policy

Thus far in this section we have only considered inventory systems with a standard (s, S) policy. The same reorder level and order-up-to level could be used for the modified (s, S) policy. However, the value of the maximum order quantity q should be specified as well. We relate the order-up-to level to the length of a replenishment cycle to determine this value. Let N denote the average number of review periods in between two consecutive orders when q does not play a role. This number depends on the actual order size, which consists of $S - s$ units plus the *undershoot* (i.e., the amount of inventory below reorder level s at a review instant). The average undershoot in inventory systems with backorders has been studied in the literature (see, e.g., Baganha, Pyke, & Ferrer, 1996; Tijms & Groenevelt, 1984).This amount can be approximated by $(\sigma_R^2 + \mu_R^2)/(2\mu_R)$. Hence,

$$N \approx \max\left\{1, \frac{S-s}{\mu_R} + \frac{\sigma_R^2 + \mu_R^2}{2\mu_R^2}\right\}. \tag{26}$$

It takes on average N review periods before a new order is placed and an extra L time units before it gets delivered, thus the time period between ordering and the delivery of the next order equals $(L + RN)$ time units. The inventory position after ordering should at most cover the demand during this period. Consequently, the maximum inventory position per time unit in a replenishment cycle equals S divided by $L + RN$. We multiply this amount by the average length of a replenishment cycle to get the maximum number of units required in a replenishment cycle. That is, we propose to set the maximum order quantity q equal to $SRN/(L + RN)$ rounded to the nearest integer. However, the maximum order size should at least be $S - s$, otherwise the policy corresponds to an (s, Q) policy. Consequently,

$$q = \max\left\{\mathrm{round}\left(S\frac{RN}{L+RN}\right), S-s\right\}. \tag{27}$$

We remark that this value of q is not an upper bound on the optimal order quantity. It represents our upper bound for the order size in the modified (s, S) policy (as a counterexample consider the example discussed in Section 3).

Numerical results of the heuristic procedures for the different replenishment policies are discussed in the next section.

5. Numerical results

The goal of this section is twofold. First, we compare the performance of the different replenishment policies discussed in Section 3. Second, the performance of the heuristic procedures is investigated to find near-optimal values of reorder level s, order-up-to level S and maximum order quantity q for the (modified) (s, S) policy. Pseudo-code for the implementation of the heuristic procedures can be found in Appendix D.

The performance of the policies is illustrated for a test bed similar to Zipkin (2008a) and Huh et al. (2009), but with fixed order cost. We set the review period length $R = 1$, the lead time $L \in \{1, 2, 3, 4\}$, the holding cost $h = 1$, the penalty cost $p \in \{4, 9, 14, 19, 39, 99\}$ and the fixed order cost $K \in \{5, 10, 25, 50\}$. For the demand distributions, we consider the expected demand $\mu_R \in \{2.5, 5, 10, 20\}$ and a variance-to-mean (VTM) ratio of 1 for a Poisson distribution, and VTM $\in \{2, 4\}$ for a negative binomial distribution. Due to the exponential growth of the state space in the lead time, the computer memory limited us to exclude instances where $L = 4$ and $\mu_R = 20$ (see also Zipkin, 2008a). This results in a test bed with 1080 instances (1152 minus 72).

5.1. Comparison of parametric policies

For the (s, S) policy, we compute the reorder level $\bar{s}$ and the order-up-to level $\bar{S}$ that minimize the expected total costs as described in Section 3.2. Similarly, $\hat{s}$ and $\hat{Q}$ minimize the expected total costs for the (s, nQ) policy according to Section 3.3. For the modified (s, S) policy, we compute the reorder level s^*, order-up-to level S^*, and maximum order quantity q^* that minimize the average costs (see Section 3.4). In addition, we also consider the modified (s, S) policy with reorder level $\bar{s}$, order-up-to level $\bar{S}$ and maximum order quantity $\bar{q}$ set by Eq. (27). Furthermore, we report the average costs C' and the average fill rate β' for the optimal replenishment policy, which is computed according to Section 3.1. The *fill rate* is defined as the fraction of customer demand satisfied by stock on hand (see also Appendix B.3). For completeness, we also include the performance when the optimal (s, S) policy for the backorder model is used in our lost-sales system (denoted by subscript BO).

Table 2
Comparison of replenishment policies for various values of lost-sales parameter p: fill rate for the optimal replenishment policy (β'), cost increase (CI) for each replenishment policy compared to the optimal policy, and CPU time (in seconds) to find the best inventory control parameter values for each policy.

		β' (percent)	CI_{BO} (percent)	$CI_{\bar{s},\bar{S}}$ (percent)	$CI_{\hat{s},n\hat{Q}}$ (percent)	CI_{s^*,S^*,q^*} (percent)	$CI_{\bar{s},\bar{S},\bar{q}}$ (percent)	T'	$T_{\bar{s},\bar{S}}$	$T_{\hat{s},n\hat{Q}}$	T_{s^*,S^*,q^*}
p	4	66.07	17.79	2.46	1.67	0.95	1.85	242.81	31.78	70.15	381.73
	9	91.07	6.12	1.44	1.41	0.74	1.09	72.76	36.29	60.05	388.42
	14	94.40	3.36	1.11	1.30	0.62	0.83	64.32	45.27	63.63	428.37
	19	95.97	2.68	0.91	1.22	0.54	0.68	68.70	52.47	65.07	422.60
	39	98.16	1.65	0.55	1.01	0.34	0.43	68.05	70.79	67.38	500.13
	99	99.32	0.92	0.30	0.81	0.19	0.24	65.74	94.25	59.68	613.72
Mean		90.82	5.00	1.08	1.22	0.55	0.82	92.06	55.91	64.13	458.24
Standard deviation		18.25	16.67	1.22	1.09	0.75	0.95	452.79	157.89	268.73	1531.06

Table 3
Comparison of replenishment policies for various values of indicator $\sqrt{2K/\mu_R hL^2}$, where N represents the number of instances corresponding to the value of the indicator in the given range.

$\sqrt{2K/\mu_R hL^2}$	N	CI_{BO} (percent)	$CI_{\bar{s},\bar{S}}$ (percent)	$CI_{\hat{s},n\hat{Q}}$ (percent)	CI_{s^*,S^*,q^*} (percent)	$CI_{\bar{s},\bar{S},\bar{q}}$ (percent)	T'	$T_{\bar{s},\bar{S}}$	$T_{\hat{s},n\hat{Q}}$	T_{s^*,S^*,q^*}
[0.00–0.25]	36	9.90	3.21	3.18	0.55	1.04	106.90	504.52	423.93	1708.53
[0.25–0.50]	198	6.97	2.27	2.08	1.05	1.49	71.55	89.33	130.12	437.54
[0.50–0.75]	180	5.71	1.49	1.48	1.02	1.28	185.57	85.82	65.00	1237.76
[0.75–1.00]	144	6.35	1.02	1.36	0.70	0.95	291.98	22.91	37.29	688.43
[1.00–1.25]	108	3.70	0.69	0.85	0.50	0.76	25.14	26.54	33.31	75.76
[1.25–1.50]	90	3.17	0.34	0.54	0.08	0.37	0.81	3.65	1.96	2.80
[1.50–2.00]	108	5.49	0.25	0.76	0.03	0.22	1.91	3.28	45.85	3.19
[2.00–3.00]	108	2.44	0.17	0.47	0.02	0.11	0.15	2.59	4.28	1.28
[3.00–6.50]	108	0.78	0.10	0.29	0.01	0.09	0.01	1.34	0.32	1.39

The results for all 1080 problem instances in our test bed are summarized in Table 2. Besides the relative cost increase (*CI*) compared to the optimal policy, it represents the expected CPU times T (in seconds) to find the best control parameters for each policy. The rows represent the average cost increase compared to the optimal policy for a specific value of the penalty cost (that is, the average over 180 instances), where the last two rows report the average and standard deviation over all test instances. The performance of each policy does not only depend on the penalty cost, but also on the time between order placements in relation to the lead time. We use the economic order quantity divided by the average demand as indicator for the length of a replenishment cycle. Next, we divide this by L to relate the cycle length to the lead time. Therefore, we consider $\sqrt{2\mu_R K/h}/\mu_L = \sqrt{2K/\mu_R hL^2}$. Note that the inverse is an indicator for the expected number of orders outstanding. Table 3 provides the average cost increase for each policy compared to the optimal policy when this variable is within a certain range. The second column indicates the number of instances that fall within each range, whereas the other columns are the same as in Table 2.

The results in these two tables illustrate that the parametric replenishment policies known from literature perform reasonably well compared to the optimal policy. However, it is important to include the lost-sales behavior when the values of these policy parameters are determined. For most instances, the (s, S) policy performs better than the (s, nQ) policy and the modified (s, S) policy improves performance even more. It is interesting to see that the (s, S) policy performs better for large values of p, whereas the (s, nQ) policy does not show the same magnitude of improvement when p increases. The modified (s, S) policy performs well for all values of p. Another observation is that the performance of these policies deviates more from the optimal cost when the value of $\sqrt{2K/\mu_R hL^2}$ decreases (that is, when K decreases or L increases). Especially when multiple orders can be outstanding (roughly when $\sqrt{2K/\mu_R hL^2} < 1$), the modified (s, S) policy significantly outperforms the other traditional parametric policies. Furthermore, we observe that the heuristic rule to set the maximum order size in the modified (s, S) policy expressed in Eq. (27) performs well when we compare $CI_{\bar{s},\bar{S}}$ to $CI_{\bar{s},\bar{S},\bar{q}}$.

Table 4
The frequency that the cost of each replenishment policy deviates from the optimal cost within a given range.

Range (percent)	(s_{BO}, S_{BO}) (percent)	$(\bar{s}, \bar{S})$ (percent)	$(\hat{s}, n\hat{Q})$ (percent)	(s^*, S^*, q^*) (percent)	$(\bar{s}, \bar{S}, \bar{q})$ (percent)
[0–1]	32.89	62.51	53.60	80.44	69.51
[1–2]	18.41	19.85	27.33	12.94	20.52
[2–3]	11.60	9.11	12.18	5.27	5.75
[3–4]	8.34	4.51	3.26	0.86	2.59
[4–5]	6.04	2.40	2.49	0.29	1.25
[5–6]	4.41	0.96	0.86	0.19	0.10
[6–7]	3.26	0.38	0.29		0.19
[7–8]	3.16	0.29			0.10
[8–9]	2.21				
[9–10]	1.34				
[10–12]	2.88				
[12–14]	1.53				
[14–16]	1.05				
[16–20]	0.86				
[20–25$^+$]	2.01				

Besides a focus on average performance, it is also interesting to see whether the average cost increase is consistent across the different test instances. This is partially reported by the standard deviations in Table 2, but in Table 4 we provide the percentage of instances for which the expected total costs of a replenishment policy deviate within a certain range from the expected costs of the optimal policy. It becomes clear that the use of a backorder model in the lost-sales system is a bad strategy. On average the cost increase is 5.0 percent over all instances, but in almost 23 percent of the instances the cost increase is more than 5 percent and even more than 10 percent in 8 percent of the instances. The costs of well-known parametric policies such as (s, S) and (s, nQ) policies are within 3 percent of the optimal costs for about 92 percent of the instances. As mentioned before,

Table 5
Comparison of replenishment policies where the parameter values are based on heuristic procedures: the average cost increase (CI) compared to the optimal policy.

	CI_{HJ2006} (percent)	CI_{JH2000} (percent)	$CI_{\tilde{s}_a,\tilde{S}_a}$ (percent)	$CI_{\tilde{s}_a,\tilde{S}_a,\tilde{q}_a}$ (percent)	$CI_{\tilde{s}_b,\tilde{S}_b}$ (percent)	$CI_{\tilde{s}_b,\tilde{S}_b,\tilde{q}_b}$ (percent)	$CI_{\tilde{s}_c,\tilde{S}_c}$ (percent)	$CI_{\tilde{s}_c,\tilde{S}_c,\tilde{q}_c}$ (percent)
p								
4	5.94	4.65	5.91	5.20	6.51	4.45	6.26	4.36
9	5.29	4.10	3.36	2.86	3.15	2.47	2.90	2.33
14	5.44	4.77	2.57	2.22	2.28	1.90	2.05	1.76
19	5.50	5.75	2.12	1.83	1.81	1.54	1.63	1.40
39	5.58	9.42	1.56	1.37	1.11	0.98	0.99	0.87
99	5.63	24.10	1.55	1.42	0.61	0.54	0.57	0.52
$\sqrt{2K/\mu_R hL^2}$								
[0.00–0.25]	47.70	51.46	8.13	5.82	9.16	7.52	7.55	5.94
[0.25–0.50]	16.96	14.20	5.32	4.42	5.31	4.38	4.37	3.44
[0.50–0.75]	3.46	19.53	3.49	3.17	2.91	2.47	2.44	2.00
[0.75–1.00]	0.56	3.89	2.58	2.39	2.05	1.54	2.04	1.50
[1.00–1.25]	0.09	1.51	1.23	1.27	1.25	0.85	1.37	0.98
[1.25–1.50]	0.01	1.37	1.12	1.00	0.84	0.51	1.18	1.10
[1.50–2.00]	0.00	1.39	1.67	1.61	0.68	0.22	1.14	0.87
[2.00–3.00]	0.00	0.92	0.77	0.71	0.47	0.14	0.89	0.70
[3.00–6.50]	0.00	0.72	0.48	0.44	0.34	0.08	0.54	0.43
mean	5.55	8.93	2.74	2.39	2.44	1.90	2.27	1.79
Standard deviation	11.78	53.81	3.16	2.82	3.12	2.59	2.73	2.18

the best modified (s, S) policies outperform these policies since its costs deviate on average 0.55 percent from the optimal cost, with an observed maximum deviation of 5.85 percent. In particular, the cost increase is within 1 percent of the optimal cost for more than 80 percent of all instances and within 3 percent for almost 99 percent of the instances. Therefore, we conclude that the modified (s, S) policy performs excellent.

Finally, when we compare the individual parameter values for the different policies (see Table C.1 in Appendix C.1), we observe that in most instances $S^* \geq \bar{S}$ due to restriction q^* on the order size in the (s, S, q) policy, and $q^* \geq \hat{Q}$ since q^* is an upper bound on the order size in the (s, S, q) policy whereas $\hat{Q}$ is the actual order size in the (s, nQ) policy.

To conclude our comparison of the different policies, we compare the computational effort to find the best parameter values for each of the policies. Tables 2 and 3 report the average CPU time in seconds to find the optimal policy, as well as the best (modified) (s, S) policy and (s, nQ) policy. All numerical results have been performed on a 2.0 gigahertz AMD Opteron 246 processor running Linux and C++. It becomes clear that the computational effort to find the optimal policy and the best values for the control parameters of the specific replenishment policies can take quite some time, especially for the modified (s, S) policy (up to 4.5 hours for one of the instances). This latter observation is to be expected, since it requires a search procedure in three dimensions. Consequently, it may be helpful to efficiently find good inventory control values with a heuristic procedure.

5.2. Comparison of heuristic procedures

The inventory control parameters based on the heuristic procedures are denoted by $\tilde{s}_i$ and $\tilde{S}_i$, where $i = a$ represents the power approximation (Section 4.1) and $i = b, c$ represents the two approaches to determine the value of $\tilde{c}_{s,S}$ in the steady-state approximation to minimize the total costs $\tilde{C}_{s,S}$ (Section 4.2). We also investigate the three corresponding modified (s, S) policies in which the reorder levels and order-up-to levels are based on the heuristic procedures for the (s, S) policy. Their maximum order quantity q (denoted by $\tilde{q}_a$, $\tilde{q}_b$ and $\tilde{q}_c$, respectively) is set according to Eq. (27). Besides our new heuristic procedures, we also compute the performance of the policy proposed by Hill and Johansen (2006) and the (s, nQ) policy proposed by Johansen and Hill (2000). These values are denoted by $HJ2006$ and $JH2000$, respectively.

Table 5 summarizes the results over all test instances similar to Tables 2 and 3. Note that we did not include the CPU times for the heuristic procedures, since they are performed within a few milliseconds. Similar to the previous section, the percentage of instances for which the expected total costs deviate within a certain range from the optimal cost is included in Table 6 and Table C.2 in Appendix C.1 contains more detailed results for 48 instances.

From these results, it becomes clear that the policy as proposed by Hill and Johansen (2006) performs close to optimal when it is expected that at most one order is outstanding (that is, when $\sqrt{2K/\mu_R hL^2} \geq 1$), which is the underlying assumption in their procedure. Unfortunately, this policy has no clear structure or parametric description. Since this policy does not provide much insight into the order quantities, it may not be preferred in practical settings. In the remainder of this section, we focus on the heuristic procedures to set the parameter values of the parametric policies.

The procedure proposed by Johansen and Hill (2000) for the (s, nQ) policy is underperforming compared to our procedures for the (modified) (s, S) policies. When we compare the results of the different heuristic procedures that we propose, we observe that the procedure based on the steady-state approximations performs better than the power approximation procedure for most instances. This is particularly true when we consider the modified (s, S) policy where the maximum order size is given by Eq. (27). In general, the first approach for the steady-state approximation procedure (i.e., $\tilde{c}_{s,S} = 1$) performs better when $\sqrt{2K/\mu_R hL^2}$ is larger than 1, whereas the second approach performs better when this amount is smaller than 1.[1] Based on the summary provided in Table 6, we conclude that the heuristic procedures for the (modified) (s, S) policies deviate within 3 percent from the optimal cost for about 73–83 percent of the instances. This illustrates again how well these heuristic procedures perform. When we zoom in on the values of the policy control parameters (see Appendix C.1), we observe (for almost all instances)

[1] We also used a power approximation to find good values of $\tilde{c}_{s,S}$. The power approximation that performed best was of the form

$$\tilde{c}_{s,S} = c_0 + c_1 K^{c_2}\left(\frac{p}{p+h}\right)^{c_3}\left(1 + \frac{\sigma^2_{R+L}}{\mu^2_R}\right)^{c_4},$$

where $c_0 = 0.667, c_1 = 0.169, c_2 = 0.142, c_3 = 0$ and $c_4 = 0$. Since the results are similar to those of the approach where $\tilde{c}_{s,S} = 1$ (see Table 5), we do not include these results.

Table 6
The frequency that the cost of each replenishment policy deviates from the optimal cost within a given range.

Range (percent)	*HJ*2006 (percent)	*JH*2000 (percent)	$(\tilde{s}_a, \tilde{S}_a)$ (percent)	$(\tilde{s}_a, \tilde{S}_a, \tilde{q}_a)$ (percent)	$(\tilde{s}_b, \tilde{S}_b)$ (percent)	$(\tilde{s}_b, \tilde{S}_b, \tilde{q}_b)$ (percent)	$(\tilde{s}_c, \tilde{S}_c)$ (percent)	$(\tilde{s}_c, \tilde{S}_c, \tilde{q}_c)$ (percent)
[0–1]	62.70	29.24	37.49	40.36	43.05	52.25	39.69	48.99
[1–2]	4.31	22.15	19.75	21.38	20.04	17.45	25.22	22.15
[2–3]	3.45	12.75	11.31	11.12	10.45	8.92	11.98	11.89
[3–4]	3.16	8.44	8.34	7.57	7.00	6.71	6.62	5.47
[4–5]	1.53	6.90	6.62	5.47	4.70	3.64	4.51	3.64
[5–6]	1.05	4.31	3.55	3.93	3.64	3.64	3.36	2.49
[6–7]	1.63	4.12	2.68	2.40	2.68	1.73	1.82	1.63
[7–8]	0.77	1.53	2.68	2.30	1.82	1.63	1.53	1.15
[8–9]	1.15	1.63	2.11	1.82	1.15	1.25	1.53	0.48
[9–10]	1.82	1.15	1.34	0.86	1.25	0.77	0.86	1.05
[10–12]	3.16	1.63	1.73	1.34	2.01	0.96	1.34	0.48
[12–14]	1.53	1.53	0.86	0.86	1.15	0.58	0.67	0.29
[14–16]	2.21	0.77	0.96	0.29	0.48	0.19	0.38	0.19
[16–20]	2.30	0.77	0.48	0.10	0.38	0.19	0.38	0.00
[20–25+]	9.20	3.07	0.10	0.19	0.19	0.10	0.10	0.10

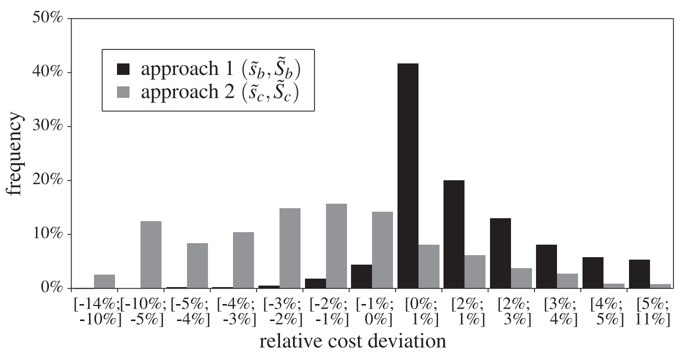

Fig. 2. The relative cost deviation of $\tilde{C}_{s,S}$ expressed by Eq. (24) compared to the actual costs for the (s, S) policy. In the first approach $\tilde{c}_{s,S} = 1$ (resulting in policy $(\tilde{s}_b, \tilde{S}_b)$), and in the second approach $\tilde{c}_{s,S}$ is based on the average order size (resulting in policy $(\tilde{s}_c, \tilde{S}_c)$).

the following relationships for the reorder level $\max\{\bar{s}, s^*\} \leq \tilde{s}_b$, and for the order-up-to level $\max\{S^*, \tilde{S}_c\} \leq \tilde{S}_b$.

Besides the performance of the heuristic procedures to find good control parameter values, we also compare Eq. (24) to approximate the expected total costs against the actual expected total costs. Fig. 2 illustrates the frequencies that the approximated costs for (s, S) policies given by Eq. (24) deviate from the actual costs within a certain range. It is clear that the first approach approximates the expected costs within 1 percent deviations from the actual costs for more than 45 percent of the instances with a bias toward overestimating the actual costs, whereas the second approach has a bias toward underestimating the actual costs.

Based on the results in Table 5 we already concluded that the steady-state approximation procedure (columns 6–9) performs better than the heuristic procedure based on power approximations (columns 4 and 5). However, the advantage of a power approximation is that the values for the inventory control parameters are set with a closed-form expression such that it becomes possible to implement this procedure with a simple spreadsheet.

6. Conclusions

The retail industry is characterized by lost customer demand in case of a stock out. Optimal replenishment policies for lost-sales inventory systems lack any form of structure. However, parametric policies provide such a structure and insight into the order quantities, which make them preferred in literature and practice (Caplin & Leahy, 2010). Our analysis shows that parametric policies such as (s, S) and (s, nQ) policies perform in general close to optimal. However, there are situations in which these policies do not perform as great (especially when it is expected to have multiple orders outstanding). Therefore, we propose a new replenishment policy; a modified (s, S) policy in which the order size is restricted to a maximum. The expected costs for this policy are within 3 percent from the optimal costs for almost 99 percent of our instances, with an average cost increase of only 0.55 percent. Since (s, S) policies are already common in many practical settings, it does not require a lot of effort to implement an upper bound on the order size.

However, the computation times to find the best values of the inventory control parameters for these structured replenishment policies can increase rapidly. Therefore, we also propose heuristic procedures to set the parameter values for (modified) (s, S) policies. The first procedure derives closed-form expressions for the parameter values with power approximations, whereas the second procedure approximates the steady-state behavior and long-run average performance measures (e.g., expected total costs, service level). Both procedures perform well and each has its own (dis)advantages. The first procedure can easily be implemented with a simple spreadsheet but is only applicable to specific inventory models. The second procedure requires a numerical search procedure (that is performed within milliseconds though), but can be generalized to many models with different objective functions and with additional (service level) constraints (see Appendix B).

Supplementary material

Supplementary material associated with this article can be found, in the online version, at 10.1016/j.ejor.2014.09.018.

References

Baganha, M., Pyke, D., & Ferrer, G. (1996). The undershoot of the reorder point: Tests of an approximation. *International Journal of Production Economics*, *45*(1–3), 311–320.
Bensoussan, A., Crouhy, M., & Proth, J. (1983). Mathematical theory of production planning. Amsterdam: North-Holland.
Bijvank, M., & Johansen, S. (2012). Periodic review lost-sales inventory models with compound Poisson demand and constant leadtimes of any length. *European Journal of Operational Research*, *220*(1), 106–114.
Caplin, A., & Leahy, J. (2010). Economic theory and the world of practice: A celebration of the (s, s) model. *The Journal of Economic Perspectives*, *24*(1), 183–202.
Cheng, F., & Sethi, S. (1999). Optimality of state-dependent (s, S) policies in inventory models with Markov-modulated demand and lost sales. *Production and Operations Management*, *8*(2), 183–192.
Chiang, C. (2007). Optimal ordering policies for periodic-review systems with a refined intra-cycle time scale. *European Journal of Operational Research*, *177*(2), 872–881.
Ehrhardt, R. (1979). The power approximation for computing (s, S) inventory policies. *Management Science*, *25*(8), 777–786.
Ehrhardt, R., & Mosier, C. (1984). A revision of the power approximation for computing (s, S) policies. *Management Science*, *30*(5), 618–622.

Hill, R., & Johansen, S. (2006). Optimal and near-optimal policies for lost sales inventory models with at most one replenishment order outstanding. *European Journal of Operational Research*, *169*(1), 111–132.
Huh, W., Janakiraman, G., Muckstadt, J., & Rusmevichientong, P. (2009). Asymptotic optimality of order-up-to policies in lost sales inventory systems. *Management Science*, *55*(3), 404–420.
Johansen, S., & Hill, R. (2000). The (r, Q) control of a periodic-review inventory system with continuous demand and lost sales. *International Journal of Production Economics*, *68*(3), 279–286.
Johansen, S., & Thorstenson, A. (1993). Optimal and approximate (Q, r) inventory policies with lost sales and Gamma-distributed lead time. *International Journal of Production Economics*, *30–31*, 179–194.
Johansen, S., & Thorstenson, A. (2008). Pure and restricted base-stock policies for the lost-sales inventory system with periodic review and constant lead times: Technical Report. Aarhus University.
Kapalka, B., Katircioglu, K., & Puterman, M. (1999). Retail inventory control with lost sales service constraints, and fractional lead times. *Production and Operations Management*, *8*(4), 393–408.
Karlin, S., & Scarf, H. (1958). Inventory models of the Arrow-Harris-Marschak type with time lag. In K. Arrow, S. Karlin, & H. Scarf (Eds.), *Studies in the mathematical theory of inventory and production*. Stanford University Press, Stanford, CA.
Levi, R., Janakiraman, G., & Nagarajan, M. (2008). A 2-approximation algorithm for stochastic inventory control models with lost sales. *Mathematics of Operations Research*, *33*(2), 351–374.
Li, Q., & Yu, P. (2012). On the quasiconcavity of lost-sales inventory models with fixed costs. *Operations Research*, *60*(2), 286–291.
Morton, T. (1969). Bounds on the solution of the lagged optimal inventory equation with no demand backlogging and proportional costs. *SIAM Review*, *11*(4), 572–596.
Morton, T. (1971). The near myopic nature of the lagged proportional cost inventory problem with lost sales. *Operations Research*, *19*(7), 1708–1716.
Nahmias, S. (1979). Simple approximations for a variety of dynamic leadtime lost-sales inventory models. *Operations Research*, *27*(5), 904–924.
Puterman, M. (2005). *Markov decision processes: Discrete stochastic dynamic programming*. New Jersey: John Wiley & Sons.
Scarf, H. (1960). The optimality of *(s, S)* policies in the dynamic inventory problem. In K. Arrow, S. Karlin, & P. Suppes (Eds.), *Mathematical methods in the social sciences*. Stanford University Press, Stanford, CA.
Shreve, S. (1976). Abbreviated proof [in the lost sales case]. In D. Bertsekas (Ed.), *Dynamic programming and stochastic control* (pp. 105–106). New York: Academic Press.
Tijms, H., & Groenevelt, H. (1984). Simple approximations for the reorder point in periodic and continuous review (s, S) inventory systems with service level constraints. *European Journal of Operational Research*, *17*(2), 175–190.
Veinott, A., Jr. (1966). On the optimality of (s, S) inventory policies: New conditions and a new proof. *SIAM Journal of Applied Mathematics*, *14*(5), 1067–1083.
Veinott, A., Jr., & Wagner, H. (1965). Computing optimal (s, S) inventory policies. *Management Science*, *11*(5), 525–552.
Xu, Y., Bisi, A., & Dada, M. (2010). New structural properties of (s, S) policies for inventory models with lost sales. *Operations Research Letters*, *38*(5), 441–449.
Zipkin, P. (2008a). Old and new methods for lost-sales inventory systems. *Operations Research*, *56*(5), 1256–1263.
Zipkin, P. (2008b). On the structure of lost-sales inventory models. *Operations Research*, *56*(4), 937–944.