University of Groningen

Periodic review and continuous ordering

Prak, Dennis; Teunter, Ruud; Riezebos, Jan

Published in:
European Journal of Operational Research

DOI:
[10.1016/j.ejor.2014.10.026](https://doi.org/10.1016/j.ejor.2014.10.026)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2015

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Prak, D., Teunter, R., & Riezebos, J. (2015). Periodic review and continuous ordering. *European Journal of Operational Research*, 242(3), 820-827. <https://doi.org/10.1016/j.ejor.2014.10.026>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Production, Manufacturing and Logistics

Periodic review and continuous ordering

Dennis Prak, Ruud Teunter*, Jan Riezebos

Department of Operations, University of Groningen, PO Box 800, 9700 AV Groningen, The Netherlands

ARTICLE INFO

Article history:

Received 11 February 2014

Accepted 14 October 2014

Available online 22 October 2014

Keywords:

Inventory

Periodic review

Continuous ordering

Newsvendor solution

ABSTRACT

Many inventory control studies consider either continuous review and continuous ordering, or periodic review and periodic ordering. Mixtures of the two are hardly ever studied. However, the model with periodic review and continuous ordering is highly relevant in practice, as information on the actual inventory level is not always up to date while making ordering decisions. This paper will therefore treat this model. Assuming zero fixed ordering costs, and allowing for a non-negative lead time and a general demand process, we first consider a one-period decision problem without salvage cost for inventory remaining at the end of the period. In this setting we derive a base-line optimal order path, described by a simple newsvendor solution with safety stocks increasing towards the end of a review period. We then show that for the general, multi-period problem, the optimal policy in a period is to first arrive at this path by not ordering until the excess buffer stock from the previous review period is depleted, then follow the path by continuous ordering, and stop ordering towards the end to limit excess stocks for the next review period. An important managerial insight is that, typically, no order should be placed at a review moment, although this may seem intuitive and is also the standard assumption in periodic review models. We illustrate that adhering to the optimal ordering path instead can lead to cost reductions of 30–60 percent compared to pure periodic ordering.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

In the inventory control literature, the focus is often on two extreme cases: either periodic stock review and periodic ordering at that same review point, or continuous stock review and continuous order possibilities. See e.g. Axsäter (2006) and Silver, Pyke, and Peterson (1998) for discussions of such models. Mixtures of both extremes are hardly ever studied. For continuous review and periodic ordering this is not surprising, since in a single-item setting the optimal policy will be equal to the pure periodic review solution with review periods equal to the time inbetween ordering points. Therefore, the sole contributions to the literature in this setting consider multi-item models. Some work has also been done on the situation of continuous review and periodic ordering, in the specific case where multiple products are jointly replenished from the same supplier to achieve cost savings. Some of the first, concrete steps here were made by Goyal (1974), who introduced an optimal algorithm for this problem. Since then a number of others have also studied this so-called “Joint Replenishment Problem”. Recently, Roushdy, Sobhy, Abdelhamid, and Mahmoud (2011) proposed an iterative method for a specific review

structure and Zhang, Kaku, and Xiao (2012) studied this problem under correlated demands.

Interestingly and perhaps surprisingly, the other mixture, periodic review and continuous ordering, has never been studied to the best of our knowledge, at least not with “truly” continuous ordering. There have been a number of contributions where orders are allowed at a number of predefined times during a period. Two decades ago, Flynn and Garstka (1990) already formulated a model and according policies where orders are allowed to be placed at the start of sub-periods of equal length during a review period. Chiang (2001) proposes order splitting in a periodic review framework. That is, at the start of a period an order is placed, and this order arrives in batches with fixed interarrival times in the current period. This method provides a holding cost advantage, which is shown by minimizing costs under a service level constraint.

However, as mentioned before, none of the previous periodic review studies considers continuous ordering, i.e. potential ordering at any point in time, as we will do in this study. This will allow us to obtain new structural results and insights into periodic review inventory systems. Moreover, whereas models with a finite number of ordering opportunities typically have to be solved using time-consuming numerical techniques such as dynamic programming, our continuous formulation leads to simple newsvendor equations that determine the optimal ordering strategy during a review period. Interestingly, this strategy is also of a quite different nature than those proposed and studied before: it typically does not order at review moments.

* Corresponding author. Tel.: +31 (0) 50 363 8617.

E-mail addresses: d.r.j.prak@rug.nl (D. Prak), r.h.teunter@rug.nl (R. H. Teunter), j.riezebos@rug.nl (J. Riezebos).

As we are the first to explore this problem, we will assume a negligible fixed ordering cost. This allows us to study the maximum benefit of continuous over periodic ordering, and also to obtain insightful analytical results. We do so under quite general conditions of a non-negative lead time and a general demand process. Although the main focus will be on continuous processes, we will also discuss the equivalent analysis for discrete demand distributions.

In line with previous periodic review studies (including those discussed above), we assume that no (partial) inventory updates are done between reviews. Obviously, this is relevant for situations where substantial effort is required to receive such updates. Despite the current technological improvements that facilitate and automate stock counting, the assumption that inventories can be completely checked on a continuous base is often unrealistic. Raman, DeHoratius, and Ton (2001) found evidence of inventory counting inaccuracy and product misplacement, Yano and Lee (1995) studied product quality issues, Nahmias (1982) analyzed spoilage due to product perishability, and Fleisch and Tellkamp (2005) performed a simulation study in which it was found that theft has severe consequences for the optimality of inventory policies that ignore resulting inaccuracies. Nevertheless, it is still worthwhile for future research to analyze whether partial information can be used to further lower costs, compared to not using that information at all, as is assumed in our initial exploration and more generally in the periodic review literature. We will return to this issue in the concluding section.

So, in our model, orders can be placed continuously and the quantity of interest is the order-up-to level for the inventory position at each time instant. We will derive the optimal policy in two phases. In the first phase, we assume that there is only one period and there is no salvage cost for inventory remaining at the end of the period. Any remaining inventory can be discarded free of charge. Given this simplifying assumption we formulate the total cost function and minimize it with respect to the order-up-to level at each time instant. The resulting policy will serve as the base-line for phase 2, where we consider the more realistic case with multiple periods in which remaining stock from any period remains present in the next period. We show that the optimal policy during a review period is to (i) not order until excess buffer stock remaining from the previous period is depleted, (ii) then apply continuous ordering following the base-line path for some time, but (iii) stop towards the end of the period in order to limit the excess buffer for the upcoming period.

The remainder of this paper is structured as follows. In Section 2 we derive the one-period base-line policy, and thereafter in Section 3 we adjust this policy to the general multi-period setting. In Section 4 we provide numerical examples and compare the policy to the periodic review, periodic ordering system, and in Section 5 we summarize our findings, discuss insights, and give concluding remarks.

2. The one-period problem: base-line model

Consider a single review period of length $T > 0$, for which at time 0 a stock level of 0 is observed. Stock information is updated only once per review period, at the start. However, non-negative orders can be placed at any time $t \in [0, T)$ and arrive after lead time $L \geq 0$. Demand D_r over a period of length r follows a distribution characterized by the continuous cdf F_{D_r} and corresponding pdf f_{D_r} . Holding costs per unit per time unit are $h > 0$ and shortage costs per unit per time unit are $p \geq h$. Fixed ordering costs are 0, and we can freely dispose of remaining inventory. Please note that since information on demand (including theft, misplacement, etc.) is not made available between reviews, demand during a review period is not subtracted from the inventory position. That is, the inventory position at any time during a review period is defined as the starting inventory position plus all orders placed since the start of the current review period.

Any inventory strategy is characterized by the order-up-to level $\mathcal{O}_t$ ($0 \leq t < T$) at any time t during a review period. Note that since de-

mands during a review period are not subtracted from the inventory position, only strategies with non-decreasing order-up-to levels need to be considered. The aim is to find the values for $\mathcal{O}_t$ that minimize the expected cost per period. An expression for that cost is obtained based on the following observation that holds for any $t \in [0, T)$: the inventory level at time $t + L$ is equal to the inventory position at time t minus the demand in interval $(0, t + L)$. Please note that we need to subtract demands in the interval $(0, t + L)$ and not only in the interval $(t, t + L)$, different from the standard analysis of continuous review inventory systems (see e.g. Axsäter, 2006, p. 90), since our definition of the inventory position at time t does not subtract the unknown demand in period $(0, t)$.

This problem is a one-period problem in the sense that effects on inventory positions after time T (or inventory levels after time $T + L$) are not taken into account. We seek for an optimal order-up-to level for any point in time in the interval $[0, T)$, so that total expected costs due to inventory levels in the interval $[L, T + L)$ are minimized. Any inventory that remains after time $T + L$ does not incur extra costs. Despite the restriction of the decision horizon to $[0, T)$, demand D_r is still defined for $r > T$, as is required in this model.

So, the total expected cost per cycle is

$$TC = \int_0^T [hE(\mathcal{O}_t - D_{t+L})^+ + pE(\mathcal{O}_t - D_{t+L})^-] dt,$$

where $(x)^+ = \max\{0, x\}$ and $(x)^- = \max\{0, -x\}$. Obviously, no best solution for the whole period can be better than applying the optimal solution at any point during the period. Next, we therefore derive the optimal solution for a specific point in time during the period, after which we show that the point-for-point optimal solution indeed determines a feasible solution for the whole period as well.

For a specific value of t , the best value of $\mathcal{O}_t$ is the one that minimizes

$$\min_{\mathcal{O}_t} \{hE(\mathcal{O}_t - D_{t+L})^+ + pE(\mathcal{O}_t - D_{t+L})^-\}. \quad (1)$$

Using integration by parts, we easily get

$$\begin{aligned} E(\mathcal{O}_t - D_{t+L})^+ &= \int_{-\infty}^{\mathcal{O}_t} (\mathcal{O}_t - x) dF_{D_{t+L}}(x) \\ &= \int_{-\infty}^{\mathcal{O}_t} \mathcal{O}_t dF_{D_{t+L}}(x) - \int_{-\infty}^{\mathcal{O}_t} x dF_{D_{t+L}}(x) \\ &= \mathcal{O}_t F_{D_{t+L}}(\mathcal{O}_t) - \mathcal{O}_t F_{D_{t+L}}(\mathcal{O}_t) + \int_{-\infty}^{\mathcal{O}_t} F_{D_{t+L}}(x) dx \\ &= \int_{-\infty}^{\mathcal{O}_t} F_{D_{t+L}}(x) dx, \end{aligned}$$

and similarly

$$E(\mathcal{O}_t - D_{t+L})^- = \int_{\mathcal{O}_t}^{\infty} [1 - F_{D_{t+L}}(x)] dx.$$

It follows that the first order condition for (1) is

$$hF_{D_{t+L}}(\mathcal{O}_t) - p[1 - F_{D_{t+L}}(\mathcal{O}_t)] = 0.$$

As

$$\frac{d^2 TC}{d\mathcal{O}_t^2} = (p + h)f_{D_{t+L}}(\mathcal{O}_t) > 0,$$

TC is convex, and hence the found solution is indeed a minimum. So, the optimal order-up-to level $\mathcal{O}_t$ for a specific time $t \in [0, T)$, must satisfy

Fig. 1. Illustration of the base-line policy for Normally distributed demand with mean $\mu = 10$ and standard deviation $\sigma = 2$ per time unit ($h = 1, p = 10, L = 0$).

$$F_{D_{t+L}}(\tilde{\mathcal{O}}_t) = \frac{p}{h+p},$$

or

$$\tilde{\mathcal{O}}_t = F_{D_{t+L}}^{-1}\left(\frac{p}{p+h}\right). \quad (2)$$

Please note that $\tilde{\mathcal{O}}_t$ is non-decreasing in t for any non-negative demand process D_{t+L} . This implies that it is indeed feasible to achieve order-up-to level $\tilde{\mathcal{O}}_t$ at any time t during the whole review period ($t \in [0, T]$). Therefore, applying (2) for all $t \in [0, T]$, and thereby inducing a continuous stream of orders, minimizes the expected cost over the whole review period.

A common assumption in both theory and practice is that demand over some time interval follows a Normal distribution. If we indeed assume, as an example, a stationary Normal demand process with mean μ and standard deviation σ per time unit, i.e. $D_r \sim N(\mu r, \sigma^2 r)$, then (2) gives

$$\tilde{\mathcal{O}}_t = \mu(t+L) + \Phi^{-1}\left(\frac{p}{p+h}\right)\sigma\sqrt{t+L}, \quad (3)$$

where Φ is the well-tabulated standard Normal distribution function.

We remark that the order-up-to levels in (3) may not always be non-decreasing over time for the (unrealistic) case that the holding cost rate is much smaller than the backorder cost rate, and the coefficient of variation σ/μ is large. However, rather than providing an interesting special case, this is an indication that the Normal distribution is unsuitable for estimating small quantiles of highly variable demand processes (due to the significant probability of demand being negative). Nevertheless, assuming Normal demand has shown to be suitable for many real-life situations, and we will also use it in our numerical investigation in Section 4.

Note from (3) that the existence of a positive lead time does not affect the nature of the problem. A larger lead time only implies that orders need to be placed earlier and that, correspondingly, a larger safety stock is needed at any time during a review period. Fig. 1 illustrates the base-line path for the case with $D_r \sim N(10r, 4r)$ ($\mu = 10, \sigma = 2$), backorder costs $p = 10$, holding costs $h = 1$, lead time $L = 0$, and the review period normalized to unity. The figure shows how the order-up-to level increases at a (slightly) diminishing rate over time, as a combined effect in (3) of a linear increase in mean demand and a non-linear increase in variance over a period of length $t + L$.

We can equivalently study the case where demand follows a discrete distribution. Now f_{D_r} is a probability mass function rather than a density function. We find

$$TC = h \sum_{x \leq \mathcal{O}_t} (\mathcal{O}_t - x) f_{D_{t+L}}(x) + p \sum_{x > \mathcal{O}_t} (x - \mathcal{O}_t) f_{D_{t+L}}(x).$$

The first-difference function is

$$\Delta TC = h \sum_{x \leq \mathcal{O}_t} f_{D_{t+L}}(x) - p \sum_{x > \mathcal{O}_t} f_{D_{t+L}}(x).$$

It follows that ΔTC is strictly increasing and that the unique minimum of TC is achieved at the smallest value $\mathcal{O}_t$ for which

$$h \sum_{x \leq \mathcal{O}_t} f_{D_{t+L}}(x) \geq p \sum_{x > \mathcal{O}_t} f_{D_{t+L}}(x),$$

which defines the optimal policy $\tilde{\mathcal{O}}_t$ and shows convexity of TC .

3. The multi-period problem: general optimal policy

Consider the same set-up as in the previous section, but now under the more realistic assumption that there are multiple review periods and all inventory remaining at the end of a review period is carried over to the next review period. Any strategy has a maximum order-up-to level $\mathcal{O}_T$. Since the base-line policy derived in the previous section (which is optimal in the one-period problem) does not take into account the extra expected holding costs that are incurred due to the remaining stock that will be carried over to the next period, the optimal value of $\mathcal{O}_T$ may be smaller than $\tilde{\mathcal{O}}_T$. This can clearly only be achieved by deviating from the base-line path by not ordering-up-to more than a certain value $\bar{S}$. Indeed, note that the order-up-to levels in a review period $[0, T]$ determine the costs in “cycle” $[L, T + L]$, but these order-up-to levels only affect costs from $T + L$ onward through the maximum order-up-to level $\bar{S}$. This maximum order-up-to level namely determines, together with total demand in the period, the starting stock level for the next period. Given this starting stock level, and the order-up-to levels in the next period, total costs for the upcoming period are completely determined. Therefore, we will first derive the optimal policy for a given level of $\bar{S}$ in some period (and for some value of the starting stock level $\underline{S}$ in that period), and then proceed to determine the unconstrained optimal policy, where the optimal value for $\bar{S}$ is determined so that also future costs are incorporated. Since given $\bar{S}$ the optimal ordering path is determined, finding the optimal value for $\bar{S}$ is the final step in finding the optimal policy.

So, let us first consider the best policy given a maximum order-up-to level $\bar{S} \leq \tilde{\mathcal{O}}_T$ and starting stock level $\underline{S} \leq \bar{S}$, for some period. From the convexity of the cost function in (1) for any $t \in [0, T]$, it follows that costs in the interval $[L, T + L]$ are minimized by staying as close to the base-line policy as possible at any time during the

period $[0, T)$. It is easy to see that this is achieved as follows. First of all, order at the start of a period such that the path $\tilde{O}_t$ is reached as soon as possible. If $\underline{S} \leq \tilde{O}_0$ this is achieved immediately at time 0 by ordering up to $\tilde{O}_0$ at that time. If $\underline{S} > \tilde{O}_0$, then no orders are placed until the first time that a point a is reached where $\tilde{O}_a = \underline{S}$. This yields the starting inventory position $\underline{S}' = \max\{\underline{S}, \tilde{O}_0\}$. Subsequently, orders at time t must be placed according to $\tilde{O}_t$ until time b where $\tilde{O}_b = \bar{S}$ and ordering should be stopped to await the next inventory review.

So, given $\underline{S}$ and $\bar{S}$ such that $\underline{S} \leq \bar{S} \leq \tilde{O}_T$, the optimal policy in a period is to set

$$\tilde{O}_t = \begin{cases} \underline{S} & 0 \leq t \in [0, a) \\ \tilde{O}_t & t \in [a, b) \\ \bar{S} & t \in [b, T), \end{cases} \quad (4)$$

where

$$a = \min[a | \tilde{O}_a \geq \underline{S}'], \quad \text{and} \\ b = \min[b | \tilde{O}_b \geq \bar{S}].$$

An illustration of such a policy can be found in Fig. 3 of Section 4. Observe that if $\underline{S} \leq \tilde{O}_0$, then $a = 0$. The equations $\tilde{O}_a \geq \underline{S}'$ and $\tilde{O}_b \geq \bar{S}$ can generally be solved with equality. An exception arises e.g. when $\underline{S}' = 0$ and the demand distribution does not support demands of size 0. Also, if the demand distribution is discrete, then there exists a range of solutions to these equations. The chosen formulation then ensures that the smallest (integer) solution is chosen, and thus that the ordering path coincides with the base-line policy as early as possible. It is obvious from the above analysis (and Fig. 3) that the optimal level, $\bar{S}$, at which to stop ordering is independent of the starting stock level, $\underline{S}$. The latter only affects the time it takes to reach the base-line path. This implies that the optimal policy is stationary in that it applies the same maximum order-up-to level for each period, independent of the starting stock level.

What remains is to find the optimal value for $\bar{S}$. From the above discussion, it follows that the cost of the optimal policy for a period, given an initial inventory position $\underline{S}$, is given by

$$TC(\underline{S}, \bar{S}) = h \left[\int_0^a E(\underline{S} - D_{t+L})^+ dt + \int_a^b E(\tilde{O}_t - D_{t+L})^+ dt + \int_b^T E(\bar{S} - D_{t+L})^+ dt \right] \\ + p \left[\int_0^a E(\underline{S} - D_{t+L})^- dt + \int_a^b E(\tilde{O}_t - D_{t+L})^- dt + \int_b^T E(\bar{S} - D_{t+L})^- dt \right],$$

This total cost function can be rewritten as

$$TC(\underline{S}, \bar{S}) = HC(\underline{S}, \bar{S}) + BC(\underline{S}, \bar{S}),$$

where

$$HC(\underline{S}, \bar{S}) = h \left[\int_0^a \int_{-\infty}^{\underline{S}} F_{D_{t+L}}(x) dx dt + \int_a^b \int_{-\infty}^{\tilde{O}_t} F_{D_{t+L}}(x) dx dt \right. \\ \left. + \int_b^T \int_{-\infty}^{\bar{S}} F_{D_{t+L}}(x) dx dt \right],$$

and

$$BC(\underline{S}, \bar{S}) = p \left[\int_0^a \int_{\underline{S}}^{\infty} [1 - F_{D_{t+L}}(x)] dx dt \right. \\ \left. + \int_a^b \int_{\tilde{O}_t}^{\infty} [1 - F_{D_{t+L}}(x)] dx dt + \int_b^T \int_{\bar{S}}^{\infty} [1 - F_{D_{t+L}}(x)] dx dt \right].$$

Using that: (i) the stock at the end of a period is equal to $\bar{S}$ minus the demand in that period, and (ii) one orders up to $\tilde{O}_0$ at the start of a period, we get the expected cost per period as

$$ETC(\bar{S}) = \int_0^{\infty} TC(\bar{S} - x, \bar{S}) f_{D_T}(x) dx \\ = \int_0^{\bar{S} - \tilde{O}_0} TC(\bar{S} - x, \bar{S}) f_{D_T}(x) dx + P(D_T > \bar{S} - \tilde{O}_0) TC(\tilde{O}_0, \bar{S}).$$

By minimizing this expected cost, the optimal value $\bar{S}$ can be determined numerically for any demand process. Summarizing, the sequential optimization procedure is to first determine for any $\bar{S}$ the optimal path towards this maximum order-up-to level via (4), and thereafter determine given this optimal path that exists for any $\bar{S}$, the corresponding costs and the optimal value for $\bar{S}$. Note that the adaption of this analysis to discrete demand distributions is easily achieved by replacing the inner integrals in HC and BC by summations, similar to the discrete version of the base-line cost function in Section 2.

4. Sensitivity analysis and comparison with pure periodic review

4.1. Numerical examples and sensitivity analysis

The policy derived in the previous section can be applied to any continuous demand process. In this section, we consider some examples with Normally distributed demand $D_T \sim N(\mu r, \sigma^2 r)$, zero lead time, and a review period of unit length. For every example, the expected total cost function is approached numerically by replacing the integrals with variable-dependent bounds by their finite sum equivalents of sufficient length.

Fig. 2 shows the expected cycle costs as a function of $\bar{S}$ for the earlier considered case with $\mu = 10$, $\sigma = 2$, $h = 1$, $p = 10$, and $L = 0$. As can be seen, the costs are convex in $\bar{S}$ and a minimum is achieved at $\bar{S} \approx 11.44$. Recall that the optimal base-line policy orders up to 12.67. So, the optimal policy is to stop ordering in the last phase of the review period, after order-up-to level 11.44 is reached. As this level is still above the mean demand of 10 per period, there typically still is (excess) stock left at the start of a period, implying an initial

Fig. 2. Base case example: Derivation of the optimal $\bar{S}$ for Normally distributed demand with mean $\mu = 10$ and standard deviation $\sigma = 2$ per time unit ($p = 10$, $h = 1$, $L = 0$).

Fig. 3. Base case example: order path and expected stock path for Normally distributed demand with mean $\mu = 10$ and standard deviation $\sigma = 2$ per time unit ($p = 10$, $h = 1$, $L = 0$), observed $\bar{S} = 1.44$ and corresponding optimal $\bar{S} = 11.44$.

phase of the review period without ordering. This leaves a phase in between with continuous ordering up to an increasing level. This is illustrated in Fig. 3a. In this particular cycle, the observed stock level at $t = 0$ is $\bar{S} = 1.44$, which is the expected value of $\bar{S}$, given $\bar{S} = 11.44$ and $\mu = 10$. This figure shows the typical three-part structure of the optimal ordering during a review period, with a horizontal start, a slightly decreasing order speed along the base-line in the middle, and an ordering stop at a level $\bar{S}$. Fig. 3b shows the corresponding expected inventory level during the cycle. The expected stock first decreases linearly with slope $-\mu = -10$, after which ordering is started and a safety stock is built up. Finally, orders are halted and the expected stock level decreases linearly again.

Now that we have seen a full-fledged example of our inventory policy, an interesting question is how it reacts to parameter changes. In Fig. 4a and b we increase μ (in two steps), ceteris paribus. From (3) we know that an increase in μ leaves the optimal safety stock levels for the base-line policy unchanged. Therefore, the maximum order-up-to level $\hat{O}_T$ increases from 12.7 for $\mu = 10$ to 27.7 for $\mu = 25$ and 52.7 for $\mu = 50$. Fig. 4a and b shows that the optimal maximum order-up-to level converges towards $\hat{O}_T$ as μ increases. This is because a higher mean demand rate implies that excess safety stocks, if observed at the next review, can be depleted at a faster rate and are therefore less costly. Next we analyze the response of the policy to

an increase in demand uncertainty. Specifically, we increase σ to 5. See Fig. 4c. The base-line policy now orders up to 16.68, whereas the optimal order stop level is $\bar{S} \approx 13.8$. Hence, an increase in standard deviation of 150 percent has led to an increase in the base-line order-up-to level of over 30 percent, whereas the optimal order-up-to level increases with approximately 20 percent. That is, if demand uncertainty increases, then the build-up of safety stock increases in optimum, but the relative deviation from the base-line policy increases as well.

Instead of altering distribution parameters, we can also study changes in cost parameters. The drawback of the base-line policy is mainly due to neglect of future holding costs. Let us consider the optimal order path if backordering becomes less costly relative to holding costs. In Fig. 4d, the backorder cost p is decreased from 10 to 4. The base-line order level is now decreased to 11.68, whereas the optimal $\bar{S}$ is now $\bar{S} \approx 10.45$. Hence, compared to the initial case, less safety stock is built up, since a shortage is less expensive and holding costs have an effect both in the current and in the next period.

4.2. Cost savings of continuous over periodic ordering

As discussed in Section 1, most inventory systems in the literature assume either periodic review with periodic ordering, or continuous

Fig. 4. Sensitivity analysis. Base case: Normally distributed demand with mean $\mu = 10$ and standard deviation $\sigma = 2$ per time unit ($p = 10, h = 1, L = 0$).

review with continuous ordering. Under zero ordering costs and with a lead time equal to 0, the latter will always provide zero cycle costs, as stock can be kept at zero and demands can be satisfied immediately. So, the cost saving of the optimal continuous ordering policy under periodic review compared to the optimal periodic ordering policy under periodic review also indicates to what extent continuous ordering can compensate for the cost disadvantage of periodic review.

We can derive the optimal order-up-to level $\tilde{S}_p$ for periodic ordering by noting that total cycle costs are given by

$$h \int_0^T E(\tilde{S}_p - D_{t+L})^+ dt + p \int_0^T E(\tilde{S}_p - D_{t+L})^- dt$$

$$= h \int_0^T \int_{-\infty}^{\tilde{S}_p} F_{D_{t+L}}(x) dx dt + p \int_0^T \int_{\tilde{S}_p}^{\infty} [1 - F_{D_{t+L}}(x)] dx dt,$$

and so the optimal order-up-to level must satisfy

$$h \int_0^T F_{D_{t+L}}(\tilde{S}_p) dt = p \int_0^T [1 - F_{D_{t+L}}(\tilde{S}_p)] dt.$$

For the base case, this gives $\tilde{S}_p = 9.6$ and an expected cycle cost of 5.9. Note that $\tilde{S}_p$ is below the expected demand during a review period,

despite the 10 to 1 ratio of backorder cost rate vs. holding cost rate. The reason is that (safety) stocks arrive at the start of a period whereas backorders occur at the end of a period, making it costly to prevent (possible) backorders. Observe that our continuous policy $\hat{O}_t$ orders more in total, but the spreading of the orders reduces holding costs, such that total costs are only 2.56, which is a reduction by 57 percent. For $\sigma = 5$, the periodic maximum order-up-to level is $\tilde{S}_p = 11.7$ and the expected cycle cost is 9.3. Fig. 5 shows again that $\hat{O}_t$ orders more in total, but total costs are only 6.45, a reduction by 31 percent. Hence, a substantial improvement can be made by considering continuous order possibilities, but the improvement decreases in the variance of the demand distribution. With increasing demand variance, it becomes more difficult to “predict” demand during a review period and respond with the best continuous ordering plan.

4.3. Applications with Gamma distributed demand

Despite the fact that the previous numerical examples all made use of Normally distributed demand, the application of the derived policy is not restricted to this distribution. In this section we will therefore show numerical examples of applications of the policy to other probability distributions. Firstly, we consider a Gamma distribution. To preserve the base case demand mean $10t$ and standard deviation $2\sqrt{t}$, we choose shape parameter $k = 25t$ and scale parameter $\theta = 0.2$.

Fig. 5. Derivation of the optimal $\bar{S}$ for Gamma distributed demand with shape $k = 1/4$ and scale $\theta = 4$ per time unit ($p = 10, h = 1, L = 0$).

Further parameter settings are identical to the base case, so $p = 10, h = 1, L = 0$, and $T = 1$. It follows that the optimal value for $\bar{S}$ is 11.52, which is only slightly larger than the optimal value under Normally distributed demand with identical mean and standard deviation. Also the convex shape of the expected total cost function with respect to $\bar{S}$ is very similar to the case of a Normal distribution as shown in Fig. 2.

Although in the previous example we used a Gamma distribution for modeling demand, the parameter settings that we chose to get the mean and standard deviation per time unit of the previous examples were such that the shape of the resulting Gamma distribution is very similar to a Normal distribution. We now consider an example with $k = 1/4$ and $\theta = 4$, resulting in a non-unimodal, strictly decreasing density function. The demand distribution has mean 1 per time unit and standard deviation 2 per time unit. As Fig. 5 shows, also in this situation the expected total cost function has a convex shape. However, the optimum for $\bar{S}$ (0.50) is now very small compared to $\bar{Q}_T$ (3.24). This is because the probability of low demands is

very high in this model, so that a large part of the ordered stock is likely to be carried over to the next period. The base-line model does not take into account the extra holding costs that are associated with this excess stock, whereas the optimal decision for $\bar{S}$ does correct for this. In Fig. 6 we depict the associated total order path, with a starting stock of 0. It is apparent that the base-line model is followed only shortly, and subsequently ordering is stopped until the next inventory review.

5. Summary and conclusion

In this paper we have presented an optimal inventory policy under periodic review with continuous ordering, for any continuous demand distribution, for any non-negative, deterministic lead time, and for zero fixed ordering costs. The implied order paths lead to expected inventory paths consisting of two downward sloping linear parts where no orders are placed, separated by a middle part in which a safety stock is built up at a diminishing rate. An important insight is that typically no order is placed at the start of a period. Instead, excess safety stocks from the previous period are likely to remain. Periodic review ordering policies in the literature do order at a review and, in fact, typically only at a review. This makes them very ineffective, which was confirmed by substantial cost savings of continuous ordering from our numerical examples.

Another important observation from the above described ordering policy is that the build-up of safety stock is not continued until the end of a period. Although the uncertainty surrounding demand since the last review does continue to increase throughout the period, excess safety stocks increase costs in the first part of the next review period. For this reason, no ordering takes place during the last part of a review period.

Given zero ordering costs, the presented policy is applicable and optimal under the quite general conditions mentioned before. The assumption of continuous demand can be relaxed without severe adaptations, such that also discrete alternatives such as the often used Poisson distribution can be studied. Non-deterministic lead times can be approximately dealt with in the same manner as has been suggested for other inventory control systems (see e.g. Axsäter, 2011), by increasing the demand variance during the lead time and (part of the) review time. Exact analysis of inventory control models with stochastic lead times is known to be very complex, especially if order crossing is allowed.

Fig. 6. Order path for Gamma distributed demand with shape $k = 1/4$ and scale $\theta = 4$ per time unit ($p = 10, h = 1, L = 0$), observed $\underline{S} = 0$ and corresponding optimal $\bar{S} = 0.50$.

As follows from our comparisons, in a situation without fixed ordering costs, it can be very lucrative to apply continuous ordering during part of a review period. However, when ordering costs are strictly positive, then such a policy cannot be optimal anymore. An adjusted policy that limits the number of orders placed is needed. This could for instance be achieved by considering more general order level, order-up-to level policies, with both levels changing during a review period. In doing such further research, results on order splitting (e.g. by Chiang, 2001) should be taken into account. However, as the term suggests, those models still assume that orders are placed at reviews, and ordering opportunities are also typically predetermined, leaving many opportunities for future research.

Our sensitivity analysis showed that the expected costs incurred by the policy increase when the degree of randomness in demand (parametered by σ) increases. This suggests that further large cost reductions can be accomplished by integrating (partial) information on demand during the review period into the model in order to reduce the variance of the remaining, random part. One obvious model is to assume that some but not all demands are recorded, e.g. customer orders are recorded, but theft, misplacement, etc. are not. Such models provide lower bounds for demand. Similarly, upper bounds can be taken into account. For example, issues like theft cannot have a larger effect than the current stock at hand. Any method that reduces the uncertainty in demand will lead to inventory cost reductions. Recall that the ideal situation of complete continuous review with continuous ordering (and under zero lead time) leads to zero costs. Cost improvements of 30–60 percent are achievable by using continuous ordering instead of periodic ordering under periodic review. Timing of ordering is essential in order to realize such a huge cost

improvement. This paper shows that it is generally better to postpone ordering until sometime after the review moment and not order at the review moment itself. Thereby, it makes a fundamental first step in bridging the gap between inefficient periodic review and periodic ordering policies, and often unrealistic continuous total review and continuous ordering policies.

References

- Axsäter, S. (2006). *Inventory control*. (2nd ed.). New York: Springer.
- Axsäter, S. (2011). Inventory control when the lead-time changes. *Production and Operations Management*, 20(1), 72–80.
- Chiang, C. (2001). Order splitting under periodic review inventory systems. *International Journal of Production Economics*, 70(1), 67–76.
- Fleisch, E., & Tellkamp, C. (2005). Inventory inaccuracy and supply chain performance: A simulation study of a retail supply chain. *International Journal of Production Economics*, 95(3), 373–385.
- Flynn, J., & Garstka, S. (1990). A dynamic inventory model with periodic auditing. *Operations Research*, 38(6), 1089–1103.
- Goyal, S. K. (1974). Determination of optimum packaging frequency of items jointly replenished. *Management Science*, 21(4), pp. 436–443.
- Nahmias, S. (1982). Perishable inventory theory: A review. *Operations Research*, 30(4), 680–708.
- Raman, A., DeHoratius, N., & Ton, Z. (2001). Execution: The missing link in retail operations. *California Management Review*, 43(3), 136–152.
- Roushdy, B., Sobhy, N., Abdelhamid, A., & Mahmoud, A. (2011). Inventory control for a joint replenishment problem with stochastic demand. *World Academy of Science, Engineering & Technology*, 5(4), 156–160.
- Silver, E., Pyke, D., & Peterson, R. (1998). *Inventory management and production planning and scheduling*. (3rd ed.). Hoboken: John Wiley & Sons.
- Yano, C. A., & Lee, H. L. (1995). Lot sizing with random yields: A review. *Operations Research*, 43(2), 311–334.
- Zhang, R., Kaku, I., & Xiao, Y. (2012). Model and heuristic algorithm of the joint replenishment problem with complete backordering and correlated demand. *International Journal of Production Economics*, 139(1), 33–41. doi:10.1016/j.ijpe.2011.01.019.