university of groningen / faculty of economics and business

14005-OPERA

Periodic review and continuous ordering

Dennis R.J. Prak
Ruud H. Teunter
Jan Riezebos

SOM is the research institute of the Faculty of Economics & Business at the University of Groningen. SOM has six programmes:
- Economics, Econometrics and Finance
- Global Economics & Management
- Human Resource Management & Organizational Behaviour
- Innovation & Organization
- Marketing
- Operations Management & Operations Research

Research Institute SOM
Faculty of Economics & Business
University of Groningen

Visiting address:
Nettelbosje 2
9747 AE Groningen
The Netherlands

Postal address:
P.O. Box 800
9700 AV Groningen
The Netherlands

T +31 50 363 7068/3815

www.rug.nl/feb/research

university of groningen / faculty of economics and business

Periodic review and continuous ordering

Dennis R.J. Prak

Ruud H. Teunter
University of Groningen
r.h.teunter@rug.nl

Jan Riezebos
University of Groningen

Periodic review and continuous ordering

D.R.J. Prak, R.H. Teunter, and J. Riezebos

February 7, 2014

Abstract

There exist many inventory control studies that consider either continuous review & continuous ordering, or periodic review & 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 consider this case of periodic review and continuous ordering. 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 with no salvage cost for inventory remaining at the end of the period. In this setting we derive a base-line optimal order path, which is 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 for some time 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 and possibly counterintuitively, 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 for Normally distributed demand that adhering to the optimal ordering path instead can lead to cost reductions of 30% to 60% compared to periodic ordering, i.e. ordering exclusively at review moments.

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 et al. (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 et al. (2011) proposed an iterative method for a specific review structure and Zhang et al. (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.

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 continuous demand process. We remark that discrete demand processes can be analyzed in the same way, but the analysis and expressions are somewhat lengthier and do not provide additional insights. We therefore restrict the exposition to continuous demand.

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 et al. (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. 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, i.e. 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}. 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 at time T 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 $\mathscr{O}_t$ $(0 \leq t < T)$ at any time t during a review period. Note that since demands 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 $\mathscr{O}_t$ 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 $(0,t)$.

So, the total expected cost per cycle is

$$\int_0^T \left[h\mathrm{E}(\mathscr{O}_t - D_{t+L})^+ + p\mathrm{E}(\mathscr{O}_t - D_{t+L})^-\right] \mathrm{d}t,$$

where $(x)^+ = \max\{0,x\}$ and $(x)^- = \max\{0,-x\}$. Obviously, the best solution for the whole period is found by 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 $\mathscr{O}_t$ is the one that minimizes

$$\min_{\mathscr{O}_t} \left\{h\mathrm{E}(\mathscr{O}_t - D_{t+L})^+ + p\mathrm{E}(\mathscr{O}_t - D_{t+L})^-\right\}. \tag{1}$$

We easily get

$$\mathrm{E}(\mathscr{O}_t - D_{t+L})^+ = \int_{-\infty}^{\mathscr{O}_t} (\mathscr{O}_t - s)\mathrm{d}F_{D_t}(s) = \int_{-\infty}^{\mathscr{O}_t}\int_{s}^{\mathscr{O}_t} \mathrm{d}x\mathrm{d}F_{D_{t+L}}(s) = \int_{-\infty}^{\mathscr{O}_t}\int_{-\infty}^{x} \mathrm{d}F_{D_{t+L}}(s)\mathrm{d}x = \int_{-\infty}^{\mathscr{O}_t} F_{D_{t+L}}(x)\mathrm{d}x,$$

and similarly

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

Using

$$\frac{\mathrm{d}}{\mathrm{d}\mathscr{O}_t}\int_{-\infty}^{\mathscr{O}_t} F_{D_{t+L}}(x)\mathrm{d}x = F_{D_{t+L}}(\mathscr{O}_t),$$

and

$$\frac{\mathrm{d}}{\mathrm{d}\mathscr{O}_t}\int_{\mathscr{O}_t}^{\infty} [1 - F_{D_{t+L}}(x)]\,\mathrm{d}x = -\left[1 - F_{D_{t+L}}(\mathscr{O}_t)\right],$$

we obtain the first order condition for (1) as

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

So, the optimal order-up-to level $\tilde{\mathscr{O}}_t$ for a specific time $t \in [0,T)$, must satisfy

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

or

$$\tilde{\mathscr{O}}_t = F^{-1}_{D_{t+L}}\left(\frac{p}{p+h}\right). \tag{2}$$

Please note that $\tilde{\mathscr{O}}_t$ 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 $\tilde{\mathscr{O}}_t$ during the whole review period ($t \in [0,T)$). Therefore, applying (2) for all $t \in [0,T)$ 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 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

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

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. Figure 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 $t + L$.

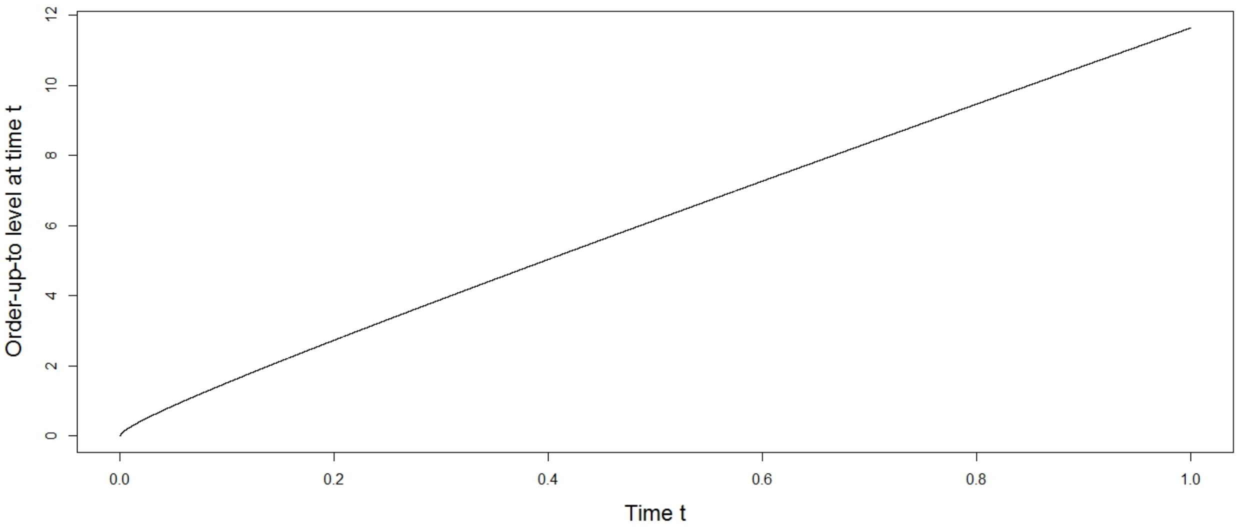

Figure 1: *Illustration of the base-line policy for Normally distributed demand with mean $\mu = 10$ and standard deviation $\sigma = 2$ per time unit ($h = 1$, $p = 10$, $L = 0$)*

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 safety 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 that the costs from $T+L$ onward, i.e. for future cycles, are only affected by the maximum order-up-to level $\bar{S}$ in a review period. 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.

So, let us first consider the best policy given a maximum order-up-to level $\bar{S} \le \tilde{\mathcal{O}}_T$ and starting stock level $\underline{S} \le \bar{S}$, 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 $\tilde{\mathcal{O}}_t$ is reached as soon as possible. If $\underline{S} \le \tilde{\mathcal{O}}_0$ this is achieved immediately at time 0 by ordering up to $\tilde{\mathcal{O}}_0$ at that time. If $\underline{S} > \tilde{\mathcal{O}}_0$, then no orders are placed until the first time that a point a is reached where $\tilde{\mathcal{O}}_a = \underline{S}$. This yields the starting inventory position $\underline{S}' = \max\{\underline{S}, \tilde{\mathcal{O}}_0\}$. Subsequently, orders at time t must be placed according to $\tilde{\mathcal{O}}_t$ until time b where $\tilde{\mathcal{O}}_b = \bar{S}$, at which point ordering should be stopped and the next inventory review must be awaited.

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

$$\hat{\mathcal{O}}_t = \begin{cases} \underline{S} & 0 \le t \in [0,a) \\ \tilde{\mathcal{O}}_t & t \in [a,b) \\ \bar{S} & t \in [b,T), \end{cases} \tag{4}$$

where

$$\begin{array}{lll} a & \text{solves} & \tilde{\mathscr{O}}_a = \underline{S}', \text{ and} \\ b & \text{solves} & \tilde{\mathscr{O}}_b = \bar{S}. \end{array}$$

Observe that if $\underline{S} \leq \tilde{\mathscr{O}}_0$, then $a = 0$. An illustration of such a policy can be found in Figure 3 of Section 4. It is obvious from the above analysis (and Figure 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

$$\begin{aligned} TC(\underline{S},\bar{S}) &= h\left[\int_0^a \mathrm{E}(\underline{S}-D_{t+L})^+\mathrm{d}t + \int_a^b \mathrm{E}(\tilde{\mathscr{O}}_t - D_{t+L})^+\mathrm{d}t + \int_b^T \mathrm{E}(\bar{S}-D_{t+L})^+\mathrm{d}t\right] \\ &\quad +p\left[\int_0^a \mathrm{E}(\underline{S}-D_{t+L})^-\mathrm{d}t + \int_a^b \mathrm{E}(\tilde{\mathscr{O}}_t - D_{t+L})^-\mathrm{d}t + \int_b^T \mathrm{E}(\bar{S}-D_{t+L})^-\mathrm{d}t\right], \end{aligned}$$

This total cost function can be rewritten as

$$\begin{aligned} TC(\underline{S},\bar{S}) &= h\left[\int_0^a\int_{-\infty}^{\underline{S}} F_{D_{t+L}}(x)\mathrm{d}x\mathrm{d}t + \int_a^b\int_{-\infty}^{\tilde{\mathscr{O}}_t} F_{D_{t+L}}(x)\mathrm{d}x\mathrm{d}t + \int_b^T\int_{-\infty}^{\bar{S}} F_{D_{t+L}}(x)\mathrm{d}x\mathrm{d}t\right] \\ &\quad +p\left[\int_0^a\int_{\underline{S}}^{\infty} [1-F_{D_{t+L}}(x)]\,\mathrm{d}x\mathrm{d}t + \int_a^b\int_{\tilde{\mathscr{O}}_t}^{\infty} [1-F_{D_{t+L}}(x)]\,\mathrm{d}x\mathrm{d}t + \int_b^T\int_{\bar{S}}^{\infty} [1-F_{D_{t+L}}(x)]\,\mathrm{d}x\mathrm{d}t\right]. \end{aligned}$$

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 $\tilde{\mathscr{O}}_0$ at the start of a period, we get the expected cost per period as

$$\begin{aligned} ETC(\bar{S}) &= \int_0^{\infty} TC\left(\bar{S}-x,\bar{S}\right) f_{D_T}(x)\mathrm{d}x \\ &= \int_0^{\bar{S}-\tilde{\mathscr{O}}_0} TC\left(\bar{S}-x,\bar{S}\right) f_{D_T}(x)\mathrm{d}x + \mathrm{P}(D_T > \bar{S}-\tilde{\mathscr{O}}_0)TC(\tilde{\mathscr{O}}_0,\bar{S}). \end{aligned}$$

By minimizing this expected cost, the optimal value $\bar{S}$ can be determined numerically for any demand process.

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 $D_r \sim N(\mu r, \sigma^2 r)$, 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.

Figure 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

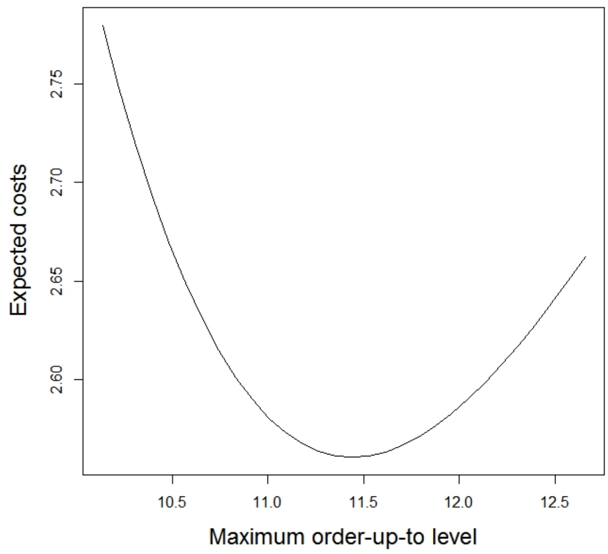

Figure 2: *Base case example: Derivation of the optimal $\bar{S}$ for Normally distributed demand with mean $\mu = 10$ and standard deviation $\sigma = 2$ per time unit ($p = 10$, $h = 1$, $L = 0$)*

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 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 Figure 3a. In this particular cycle, the observed stock level at $t = 0$ is $\underline{S} = 1.44$, which is the expected value of $\underline{S}$, 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 $\bar{S}$. Figure 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 Figure 4a and Figure 4b 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 $\hat{\tilde{O}}_T$ increases from 12.7 for $\mu = 10$ to 27.7 for $\mu = 25$ and 52.7 for $\mu = 50$. Figures 4a and 4b show that the optimal maximum order-up-to level converges towards $\tilde{O}_T$ 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 Figure 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% has led to an increase in the base-line order-up-to level of over 30%, whereas the optimal order-up-to level increases with approximately 20%. 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 neglection of future holding costs. Let us consider the optimal order path if backordering becomes less costly relative to holding costs. In Figure 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

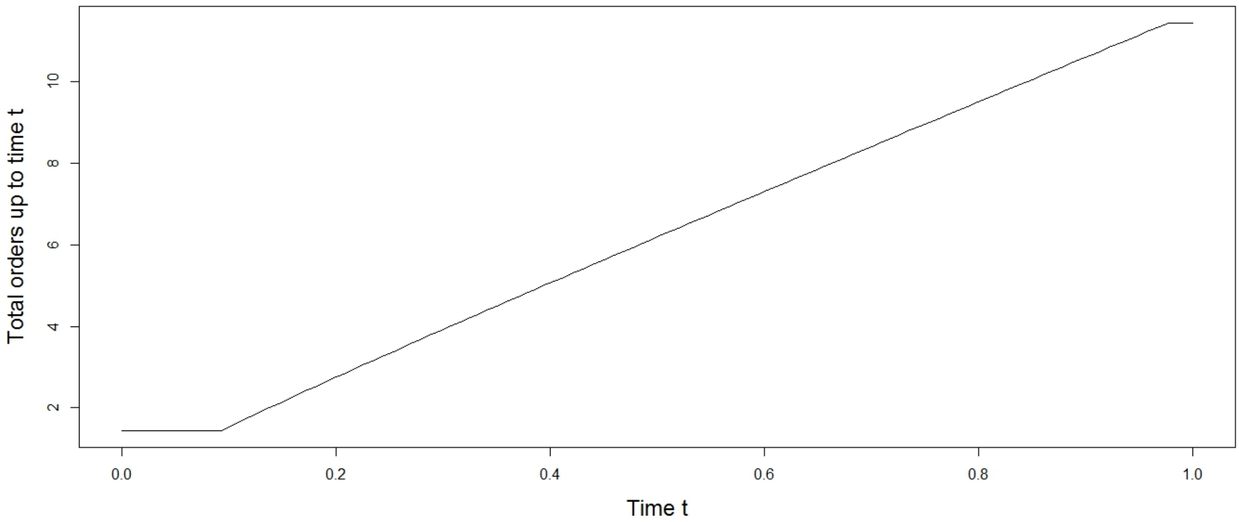

(a) Order path

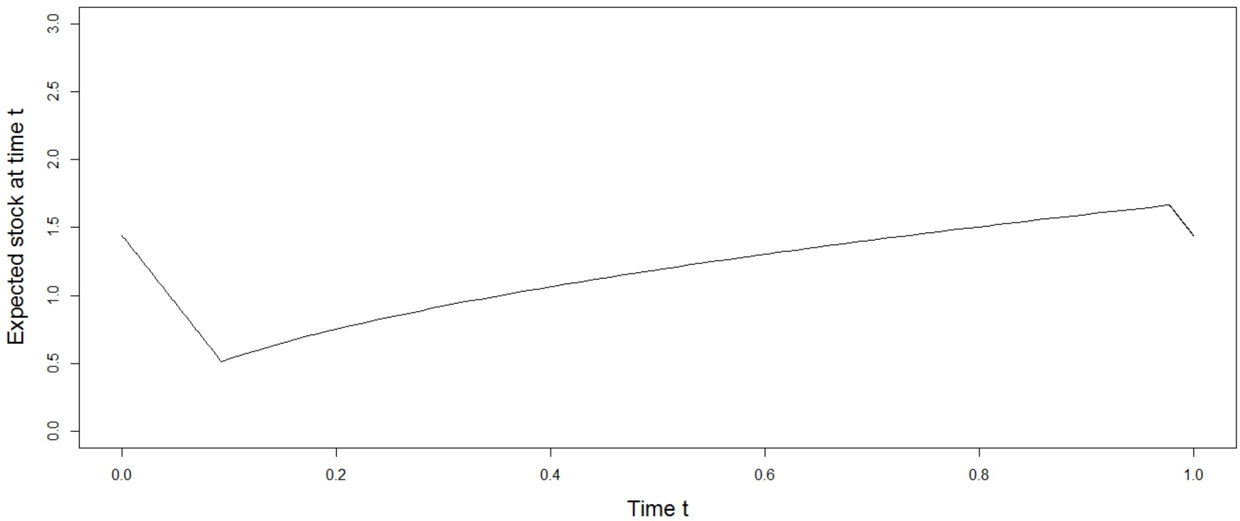

(b) Expected stock path

Figure 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* $\underline{S} = 1.44$ *and corresponding optimal* $\bar{S} = 11.44$

$\bar{S} \approx 10.45$. 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.

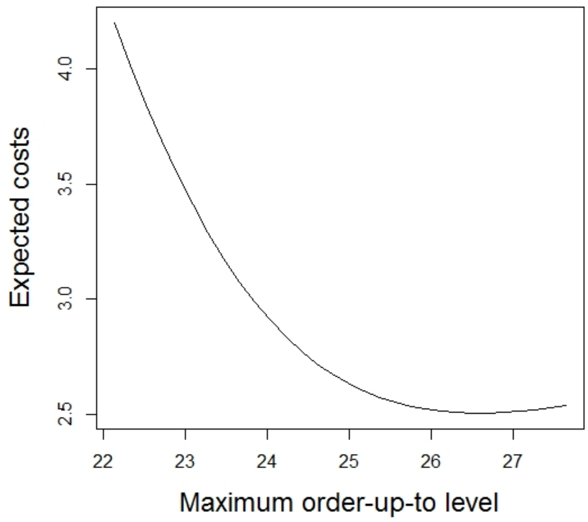

(a) Mean demand μ per time unit increased to 25

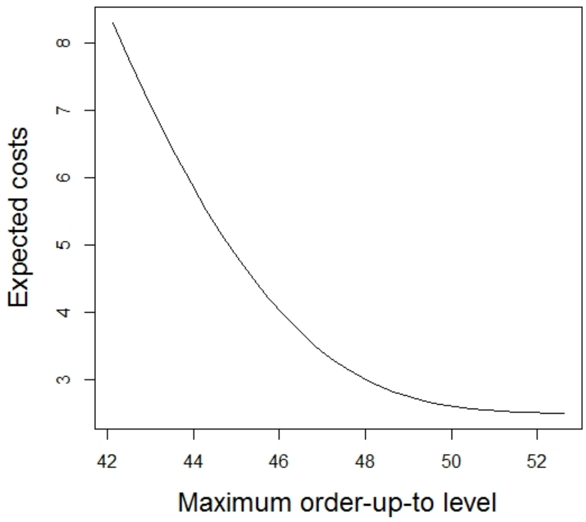

(b) Mean demand μ per time unit increased to 50

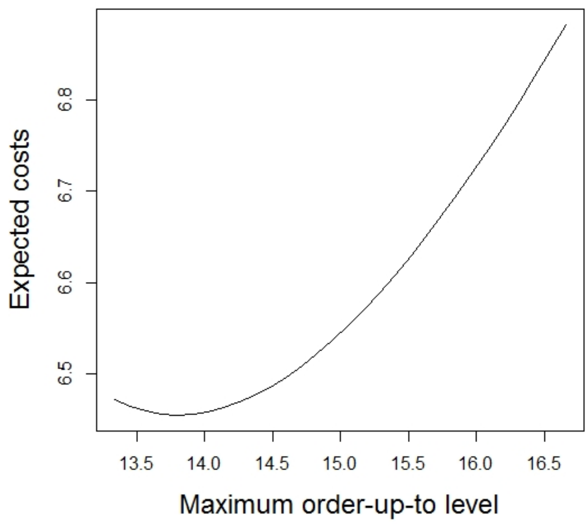

(c) Demand variance σ per time unit increased to 5

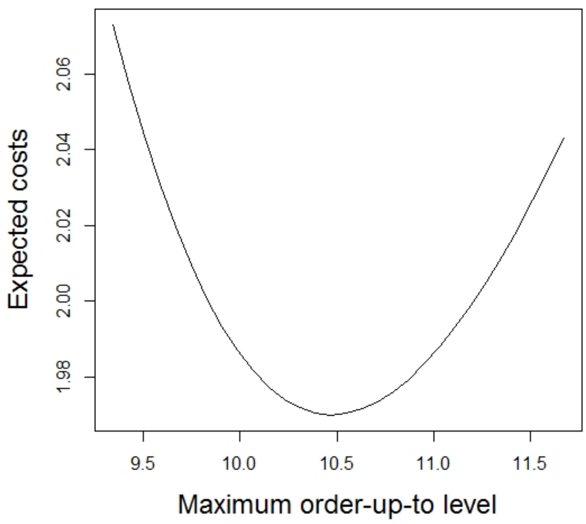

(d) Backorder cost p per unit per time unit decreased to 4

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

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 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 $\bar{S}_p$ for periodic ordering by noting that total cycle costs are given by

$$
\begin{aligned}
& h\int_0^T \mathrm{E}(\bar{S}_p - D_{t+L})^+ \mathrm{d}t + p\int_0^T \mathrm{E}(\bar{S}_p - D_{t+L})^- \mathrm{d}t \\
= \; & h\int_0^T \int_{-\infty}^{\bar{S}_p} F_{D_{t+L}}(x)\mathrm{d}x\mathrm{d}t + p\int_0^T \int_{\bar{S}_p}^{\infty} [1 - F_{D_{t+L}}(x)]\mathrm{d}x\mathrm{d}t,
\end{aligned}
$$

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

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

For the base case, this gives $\bar{S}_p = 9.6$ and an expected cycle cost of 5.9. Note that $\bar{S}_p$ 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 $\hat{\mathscr{O}}_t$ 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%. For $\sigma = 5$, the periodic maximum order-up-to level is $\bar{S}_p = 11.7$ and the expected cycle cost is 9.3. Figure 5 shows again that $\hat{\mathscr{O}}_t$ orders more in total, but total costs are only 6.45, a reduction by 31%. 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.

5 Summary & 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.

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 to 60% are achievable by using continuous ordering instead of periodic ordering under periodic review. Timing of ordering is essential in order to realise such a huge cost improvement. This paper shows that it is generally better to postpone ordering until some time 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* (second 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. and C. Tellkamp (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. and S. Garstka (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., N. DeHoratius, and Z. Ton (2001). Execution: The missing link in retail operations. *California Management Review 43*(3), 136 – 152.

Roushdy, B., N. Sobhy, A. Abdelhamid, and A. Mahmoud (2011). Inventory control for a joint replenishment problem with stochastic demand. *World Academy of Science, Engineering & Technology*, 156 – 160.

Silver, E.A., D.F. Pyke, and R. Peterson (1998). *Inventory management and production planning and scheduling* (third ed.). Hoboken: John Wiley & Sons.

Yano, C. A. and H. L. Lee (1995). Lot sizing with random yields: A review. *Operations Research 43*(2), 311 – 334.

Zhang, R., I. Kaku, and Y. Xiao (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.

List of research reports

12001-HRM&OB: Veltrop, D.B., C.L.M. Hermes, T.J.B.M. Postma and J. de Haan, A Tale of Two Factions: Exploring the Relationship between Factional Faultlines and Conflict Management in Pension Fund Boards

12002-EEF: Angelini, V. and J.O. Mierau, Social and Economic Aspects of Childhood Health: Evidence from Western-Europe

12003-Other: Valkenhoef, G.H.M. van, T. Tervonen, E.O. de Brock and H. Hillege, Clinical trials information in drug development and regulation: existing systems and standards

12004-EEF: Toolsema, L.A. and M.A. Allers, Welfare financing: Grant allocation and efficiency

12005-EEF: Boonman, T.M., J.P.A.M. Jacobs and G.H. Kuper, The Global Financial Crisis and currency crises in Latin America

12006-EEF: Kuper, G.H. and E. Sterken, Participation and Performance at the London 2012 Olympics

12007-Other: Zhao, J., G.H.M. van Valkenhoef, E.O. de Brock and H. Hillege, ADDIS: an automated way to do network meta-analysis

12008-GEM: Hoorn, A.A.J. van, Individualism and the cultural roots of management practices

12009-EEF: Dungey, M., J.P.A.M. Jacobs, J. Tian and S. van Norden, On trend-cycle decomposition and data revision

12010-EEF: Jong-A-Pin, R., J-E. Sturm and J. de Haan, Using real-time data to test for political budget cycles

12011-EEF: Samarina, A., Monetary targeting and financial system characteristics: An empirical analysis

12012-EEF: Alessie, R., V. Angelini and P. van Santen, Pension wealth and household savings in Europe: Evidence from SHARELIFE

13001-EEF: Kuper, G.H. and M. Mulder, Cross-border infrastructure constraints, regulatory measures and economic integration of the Dutch – German gas market

13002-EEF: Klein Goldewijk, G.M. and J.P.A.M. Jacobs, The relation between stature and long bone length in the Roman Empire

13003-EEF: Mulder, M. and L. Schoonbeek, Decomposing changes in competition in the Dutch electricity market through the Residual Supply Index

13004-EEF: Kuper, G.H. and M. Mulder, Cross-border constraints, institutional changes and integration of the Dutch – German gas market

13005-EEF: Wiese, R., Do political or economic factors drive healthcare financing privatisations? Empirical evidence from OECD countries

13006-EEF: Elhorst, J.P., P. Heijnen, A. Samarina and J.P.A.M. Jacobs, State transfers at different moments in time: A spatial probit approach

13007-EEF: Mierau, J.O., The activity and lethality of militant groups: Ideology, capacity, and environment

13008-EEF: Dijkstra, P.T., M.A. Haan and M. Mulder, The effect of industry structure and yardstick design on strategic behavior with yardstick competition: an experimental study

13009-GEM: Hoorn, A.A.J. van, Values of financial services professionals and the global financial crisis as a crisis of ethics

13010-EEF: Boonman, T.M., Sovereign defaults, business cycles and economic growth in Latin America, 1870-2012

13011-EEF: He, X., J.P.A.M Jacobs, G.H. Kuper and J.E. Ligthart, On the impact of the global financial crisis on the euro area

13012-GEM: Hoorn, A.A.J. van, Generational shifts in managerial values and the coming of a global business culture

13013-EEF: Samarina, A. and J.E. Sturm, Factors leading to inflation targeting – The impact of adoption

13014-EEF: Allers, M.A. and E. Merkus, Soft budget constraint but no moral hazard? The Dutch local government bailout puzzle

13015-GEM: Hoorn, A.A.J. van, Trust and management: Explaining cross-national differences in work autonomy

13016-EEF: Boonman, T.M., J.P.A.M. Jacobs and G.H. Kuper, Sovereign debt crises in Latin America: A market pressure approach

13017-GEM: Oosterhaven, J., M.C. Bouwmeester and M. Nozaki, The impact of production and infrastructure shocks: A non-linear input-output programming approach, tested on an hypothetical economy

13018-EEF: Cavapozzi, D., W. Han and R. Miniaci, Alternative weighting structures for multidimensional poverty assessment

14001-OPERA: Germs, R. and N.D. van Foreest, Optimal control of production-inventory systems with constant and compound poisson demand

14002-EEF: Bao, T. and J. Duffy, Adaptive vs. eductive learning: Theory and evidence

14003-OPERA: Syntetos, A.A. and R.H. Teunter, On the calculation of safety stocks

14004-EEF: Bouwmeester, M.C., J. Oosterhaven and J.M. Rueda-Cantuche, Measuring the EU value added embodied in EU foreign exports by consolidating 27 national supply and use tables for 2000-2007

14005-OPERA: Prak, D.R.J., R.H. Teunter and J. Riezebos, Periodic review and continuous ordering www.rug.nl/feb