Hedenstierna, C.P.T. and Disney, S.M., (2014), “The inventory ripple effect in just-in-time systems with auto-correlated demand”,
Pre-prints of the 18th International Working Seminar of Production Economics, Innsbruck, Austria, February 24th–28th, Vol. 3, pp207–217.

The Inventory Ripple Effect in Periodic Review Systems with Auto-correlated Demand

Carl Philip T. Hedenstierna [1] and Stephen M. Disney [2]

Logistics Systems Dynamics Group, Cardiff University, Aberconway Building,
Colum Drive, Cardiff, CF10 3EU, United Kingdom.
[1]cpthed@gmail.com, [2]DisneySM@cardiff.ac.uk

Abstract

We investigate a replenishment system with periodic inventory inspections, where at least one batch is completed between inspections, and where orders are placed once every P inspection periods. In this type of system, each ordering occasion will generate P orders for delivery in each of P periods in the future. We can keep the inventory level in each inspection period centered on the inventory norm (the safety stock level), but generating multiple orders at one point in time, with different delivery dates (effectively different lead-times), will cause the inventory variance to change over time. We call this phenomenon the inventory ripple effect. This paper identifies an Order-Up-To policy with minimum mean square error forecasts, under linear holding and backlog costs when demand is a normally distributed first-order autoregressive process. For this we identify a lower bound for the inventory ripple effect. We find that the introduction of positive autocorrelation in demand amplifies the inventory ripple effect in comparison to demand with independent and identically distributed (i.i.d.) error terms, while negatively correlated demand provides an effect smaller than that of i.i.d. demand. A time-varying safety stock setting proves optimal, being significantly more efficient than constant safety stock levels.

Keywords: Order-Up-To replenishment policy, Inventory variability, Staggered deliveries, Reorder period, Reorder cycle.

1. Introduction

Just-In-Time principles are instrumental in the pursuit of supply chain efficiency. One of these is the idea of *small batch sizes*, which keeps cycle stocks low; another is the concept of *level production*, which not only reduces the need for frequent changes in capacity, but also minimizes the (peak) capacity requirement (Shingo, 1989). One way to achieve level production is via cyclical planning, i.e. to determine and freeze orders periodically, such that the orders for some time into the future (e.g. a week, fortnight, or month) are determined at a single point in time, and not changed until the next cycle begins. Level production can then be achieved, either by collecting all overtime work to one period, while the remaining periods have a constant production volume, or by distributing all of the overtime evenly over the cycle.

Cyclical planning is commonplace. A survey of 292 Swedish companies found that 67% of the companies had daily planning, while 21% planned once per week, and 13% planned once every fortnight or less frequently (Jonsson and Mattsson, 2013). An international perspective on planning cycles can be found in Table 1.

The question is how cyclical planning affects production system performance, particularly if small batch sizes are used. If at least one batch is produced per period, and inventory inspections take place at the end of each period; then the effect of cycle stock can be ignored. Under these conditions, we can determine the orders for each period in a cycle, such that the expected inventory level assumes any desired value. Typically, this equals the safety stock level. However, the inventory variance is another matter. Because multiple orders are placed at once, and they effectively have different lead times due to the timing of delivery relative to

Company	Industry	Planning frequency	When
Anonymous	Computer consumables	Weekly	2010-2014
Tesco	Grocery	Every 8 hours or daily depending on store size	2000-2005+
Renishaw	Industrial measuring equipment	Monthly	2014
P&G	Household goods	Weekly	2000-2014
BAT	Fast moving consumer goods	Monthly	2012
Harmon Kardon	Audio equipment	Weekly	2001
Princes	Fruit Juice	Weekly	2003
TRW	Automotive engine components	Weekly	1999

Table 1. Current industrial planning cycles (*Source: Authors*)

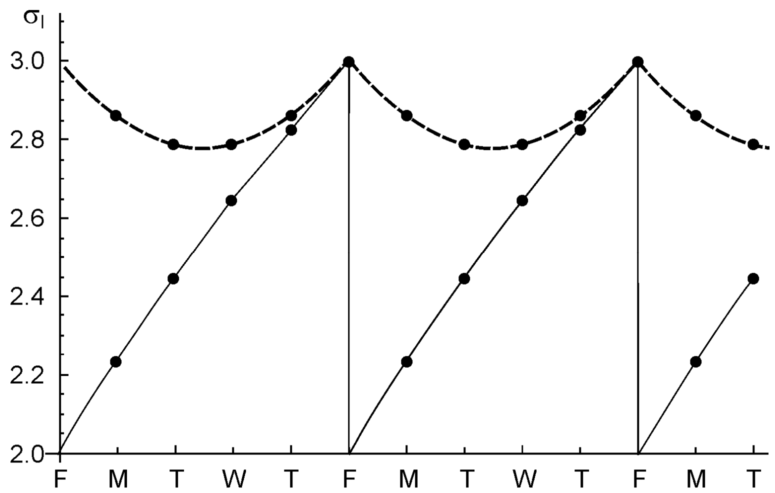

Figure 1. Inventory standard deviation for two implementations of the Order-Up-To policy when the ripple effect is present (*Hedenstierna and Disney, 2013*)

the time the plan was issued, the inventory variance will change over time. We term this the *inventory ripple effect*, as its graphical manifestation resembles ripples or saw teeth of various designs (see Figure 1).

The phenomenon is studied in detail for independent and identically distributed demand in Hedenstierna and Disney (2013), where several factors are found to determine it. These include the reorder cycle length, the overtime strategy used and eventual use of production smoothing policies. An important result is that variable safety stocks are optimal when the effect is present (this can also be inferred from Flynn, 2008).

This paper presents an investigation of the minimum inventory ripple effect that will be experienced when demand is a first-order autoregressive process. We then extend this to cases where the autocorrelation function of demand is arbitrary. This is followed by a section on analytical and numerical insights and an investigation of the optimal (time varying) versus a conventional (constant) safety stock strategies. Finally, we summarize our findings and implications.

2. Literature review

We can trace remarks about the length of the reorder cycle back to 1924, when General Motors Corporation, following a decision by Alfred P. Sloan, shifted from reviewing their production plans once every three months, to once every ten days. Sloan (1963) documents that General Motors managed to increase their inventory turnover, from two, to seven-and-a-half times per annum, due to the improvements made to the production and distribution system in the 1920’s.

The Period Batch Control (PBC) system, used in the production of the Spitfire aircraft (Burbidge, 1989) also considers frequent reordering to be important. Burbidge claims that reordering cycles should be as short as capacity permits, and that reductions in set-up time should be pursued to allow further shortening of the reorder cycle. The same line of reasoning is presented by Shingo (1989) on the implementation of Just-In-Time production. At the time, the Toyota Motor Corporation used reordering cycles of ten days, but Shingo entertained the idea of reducing it to daily or weekly cycles. Ohno (1988) also writes about long reordering cycles and their harmful effect on inventory performance.

Various models that contain reorder cycles have been presented. Several approaches to understanding the reorder cycle exist, with a fundamental one assuming that inventory inspections are synchronized with ordering, and that a single batch is produced over the order cycle. The result is the economic order quantity (EOQ) model, converted to an economic ordering cycle through division by the average demand rate (Waters, 2003). When this problem is relaxed to also consider inventory inspections at discrete points within an inventory cycle (multiple inspections per lot ordered) exact costs can be found in Chiang (2006, 2007), when the reorder period is given as an input variable. Silver and Robb (2008) seek to identify an optimum reorder period under these circumstances, but find that this type of system is sensitive even to small changes in the parameters, and that it is therefore difficult to make any general claims about how these systems can be improved.

In a model with several batches per ordering decision, Modigliani and Hohn (1955) presents a model for production planning when all cyclical demand is known at the start of the reorder cycle. They find that the reorder cycle should not be longer than one seasonal cycle, and possibly shorter if inventory costs are high. Tang and Grubbström (2002) present a different model taking into account inventory holding costs and schedule change costs; they also identify a cost-optimal reorder cycle length for their model.

Systems with multiple orders per reorder cycle, backlog and holding costs, plus fixed costs per order cycle, are covered in Flynn and Garstka (1997), who demonstrate that an optimum solution exists. Flynn (2001) describes how this applies to a multi-product scenario.

Hedenstierna and Disney (2013) present a model considering the production of at least one batch between inventory inspections, linear holding and backlog costs, reactive capacity costs, and a fixed cost per reorder cycle. They identify that the inventory variance changes over the reorder cycle (the inventory ripple effect), and find that this is dependent on both the order rate, and the way in which overtime work is planned. Longer reorder cycles increase the inventory related costs, but decrease capacity costs due to variability pooling over time. Optimum solutions are identified, and it is demonstrated that the reorder period should be as short as possible when cycle costs are absent.

Reference	**Advice given**
Burbidge (1988)	As frequently as capacity permits
Flynn (2008)	Approximate optimum
Hedenstierna & Disney (2013)	Optimize
Modigliani and Hohn (1955)	No longer than one seasonal cycle
Shingo (1989)	Ten days or weekly
Silver and Robb (2008)	No recommendation
Tang and Grubbström (2002)	Optimize
King (2012, p.216-222)	As short as possible based on "available time left over after changeovers", or linked to the EOQ model

Table 2. Current advice given on length of planning cycles

3. Model development

Consider a supply chain whose inventory level i is inspected at discrete points in time, $t \in \mathbb{N}$, and where negative inventory indicates a backlog. In every period, the inventory satisfies demand d_t and receives orders o_{t-L-1} that were released (but not necessarily planned) $L+1$ periods ago, where $L \in \mathbb{N}_0$ is the lead time exceeding the inventory inspection interval. For a lead time of zero, orders would arrive prior to the next inventory inspection. The inventory balance equation is $i_t = i_{t-1} + o_{t-1-L} - d_t$, and an accompanying Work-In-Progress (w_t) term $w_t = w_{t-1} + o_{t-1} - o_{t-L-1} = \sum_{i=1}^{L+1} o_{t-i}$. In each period a holding cost *H*, or a backlog cost *B* is applied to the inventory level according to $\$_i[t] = H(i_t)^+ + B(i_t)^-$ where $(x)^+ = \max\{0,-x\}$ and $(x)^- = \max\{0,-x\}$.

Demand is expressed as a first-order autoregressive (AR(1)) process, where the demand *d* in any given period *t* consists of a random noise element $\varepsilon_t \sim \mathcal{N}(0,\sigma_\varepsilon^2)$, plus a multiple of the stochastic components in the previous period's demand

$$d_t = \phi(d_{t-1} - \mu) + \varepsilon_t + \mu \tag{1}$$

where μ is the average periodic demand and ϕ is the autoregressive parameter, which is defined for all ϕ, but provides a stable and invertible signal only when $|\phi| < 1$. Note $|\phi| < 1$ is required for a finite demand order variance, but finite inventory variances are produced for all ϕ, Graves (1999). A practical justification for the AR(1) demand process can be found in Lee et al. (2000), who employed it following an investigation where the demand signals for over 130 products were identified as AR(1). Under these assumptions, the expected inventory cost $\$_i$ in an arbitrary period *k* is known to be

$$E\left[\$_i[k]\right] = H \cdot SS_k - \frac{H+B}{\sigma_{i,k}} \int_{-\infty}^{0} \varphi\left[\frac{x - SS_k}{\sigma_{i,k}}\right] x dx = H \cdot SS_k + \sigma_{i,k}(H+B) G\left[\frac{SS_k}{\sigma_{i,k}}\right] \tag{2}$$

where $G[x] = \varphi[x] + x(\Phi[x] - 1)$ is the standard normal loss function, SS_k is the safety stock level in period *k*, and $\sigma_{i,k}$ is the standard deviation of the inventory at time *k* (Axsäter, 2000).

$\varphi[\cdot]$ is the probability density function of the standard normal distribution, while $\Phi[\cdot]$ is the cumulative standard normal distribution function. The minimum inventory cost is obtained when $SS_k^* = \sigma_{i,k} z_i$, where $z_i = \Phi^{-1}\left[\frac{B}{B+H}\right]$, in which $\Phi^{-1}[\cdot]$ is the inverse of $\Phi[\cdot]$, applied to the critical ratio, $B/(B+H)$. When SS_k^* is used, the inventory costs are linear in $\sigma_{i,k}$ and can be expressed as

$$\$_{i,k}^*\left[\sigma_{i,k}\right] = \sigma_{i,k}\left(B+H\right)\varphi\left[z_i\right]. \tag{3}$$

We note that if the standard deviation of the inventory changes, the safety stock must be adapted so that cost optimality can be maintained.

4. The reordering mechanism

While inventory is observed every period, we assume that the production quantities are determined only once every *P* periods, when $P \mid t$, and are released for production over the next *P* periods. These orders will then be received in the periods $t+L+k$ where $1 \le k \in \mathbb{N} \le P$ denotes an arbitrary day in the inventory cycle. The relation between the inventory and ordering mechanism are described in Figure 2.

We seek an ordering policy that minimizes the inventory cost given the reordering cycle constraint on orders. Consider that a unit demand impulse that occurs in the *m*’th period, counted backwards from the point of ordering, where $1 \le m \le P$ can be corrected via order receipts no earlier than $L+m$ periods ahead of the impulse. The inventory shortfall, resulting from a unit impulse and its autocorrelation, is at the time of the first possible receipt

$$\sum_{n=0}^{L+m} \phi^n = \frac{1-\phi^{1+L+m}}{1-\phi}. \tag{4}$$

In every period, the expected inventory should equal the safety stock level. For the specific case when $k=1$, the amount to order equals all independent demand shocks (indexed by *m*) from the previous cycle multiplied by (4). Because AR(1) processes with i.i.d. shocks are Markovian, the order quantity simplifies to the first term in (6). For the remaining $P-1$ periods in the cycle, we can also set the expected inventory level equal to the safety sock by ordering the quantity provided by the minimum-mean-squared-error (MMSE) forecast, $\phi^{L+m+k-1}$.

Tsypkin (1964) provides that the variance of a process equals the sum of its squared impulse response from zero to infinity (i.e. its autocorrelation function), but since only the first $P+L$ periods are non-zero, only these will influence the inventory costs when an MMSE Order-Up-To policy is employed.

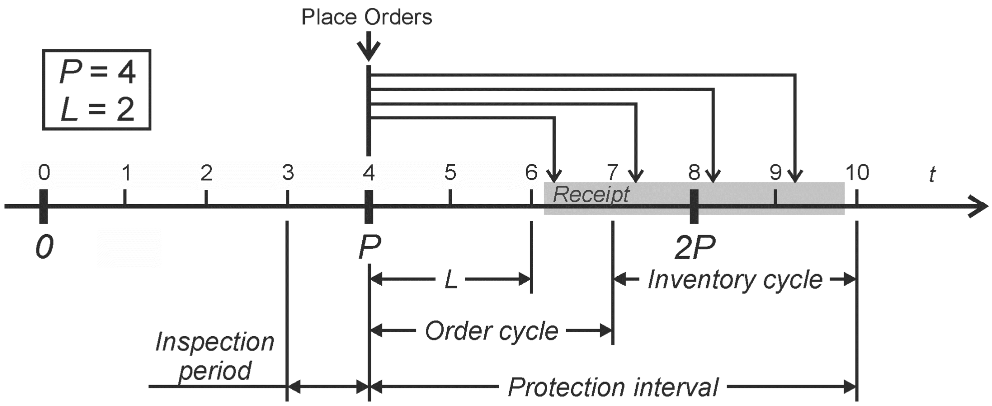

Figure 2. The sequence of events (*Hedenstierna & Disney, 2013*)

We seek to identify the inventory variance when *k* is given. The greatest possible deficit from a single impulse at a point *k* is $\sum_{x=0}^{L+k-1}\phi^x$, where the impulse has occurred at the beginning of the current protection interval. There will also be one uncorrected i.i.d. impulse for every period between this point in time, until the present. Knowing this, Tsypkin's relation (Tsypkin 1964) can be used to obtain the inventory variance ratio,

$$\frac{\sigma_{i,k}^2}{\sigma_\varepsilon^2}=\sum_{n=0}^{L+k-1}\left(\sum_{x=0}^{n}\phi^x\right)^2=\frac{\phi\left(\phi^{k+L}-1\right)\left(\phi^{1+k+L}-\phi-2\right)}{(\phi-1)^3(\phi+1)}+\frac{k+L}{(\phi-1)^2}. \tag{5}$$

From (5), we can see that the inventory variance changes with the particular period of the inventory cycle. That is, inventory levels are heteroskedactic. As (5) is used for calculating SS^* and the inventory cost, both are varying over the periods in the inventory cycle.

The variance of the orders at time *k* in the order cycle can be found in a similar manner. Consider first the initial reaction to a unit impulse, possible only when $k=1$. At this point in time, $\sum_{x=0}^{L+m}\phi^x$ units will be ordered in response to an original shock, and ϕ^{L+xP+m}, where $x\in\mathbb{N}$ will be ordered due to autocorrelation with past shocks.

At this point in time, orders will be placed such that the projected inventory deviation at the time of receipt is zero.

$$o_t=\begin{cases}\mu+\dfrac{\phi\left(\phi^{L+1}-1\right)}{\phi-1}(d_t-\mu)-i_t-w_t+S_{t+L+1} & \text{when } k=1\\ \mu+(d_{t-k+1}-\mu)\phi^{L+k}+S_{t+L+1}-S_{t+L} & \text{otherwise}\end{cases} \tag{6}$$

It is possible to eliminate *k* via the following substitution: $k=t-P\lfloor t/P\rfloor+1$, assuming that that $k=1$ occurs when $t=0$. Just as before, the variance will differ over the inspection periods in the cycle. Again, the impulse response is considered for all *P* random variables that affect each independent period, *k*. It is beneficial to consider that the order rate consists of two components: the first being orders to correct for unforeseen deviations, and the second being corrections for foreseeable deviations, i.e. the autocorrelations of demand observed in previous cycles. This second component is used for all periods, but the first component is only invoked when $k=P$.

Starting with the second component, we know that the first deviation where we can “catch up” with demand, is $k+L+1$ periods ahead (all demand occurring before this must be corrected as an inventory shortfall). Given that we have P interlacing signals, the squared impulse response is given by

$$\sum_{n=L+1+k}^{\infty} \phi^{2n} = \frac{\phi^{2(1+k+L)}}{\left(1-\phi^2\right)}. \tag{7}$$

Apart from this, there is the first component (only used for orders where $k=1$), which corrects for unforeseen deviations that have occurred in the most recent order cycle. This deviation will accumulate until it can be corrected $m+L$ periods later, giving the impulse response $\sum_{m=1}^{P}\left(\sum_{n=0}^{m+L}\phi^n\right)^2$, which when combined with (7) gives the following general expressions for the order rate variance,

$$\frac{\sigma_{o,1}^2}{\sigma_\varepsilon^2} = \left(\sum_{m=1}^{P}\left(\sum_{n=0}^{m+L}\phi^n\right)^2 + \sum_{n=k+L+1}^{\infty}\phi^{2n}\right) = \frac{P}{(\phi-1)^2} - \frac{2\phi^{2+L}\left(\phi^P-1\right)}{(\phi-1)^3} - \frac{\phi^{2+2L}\left(\phi^2+\phi^{2P}-2\phi^{1+2P}\right)}{(\phi-1)^3(\phi+1)}, \tag{8}$$

$$\frac{\sigma_{o,k}^2}{\sigma_\varepsilon^2} = \left(\sum_{n=k+L}^{\infty}\phi^{2n}\right) = \frac{\phi^{2(k+L)}}{1-\phi^2} \qquad \text{when } k \neq 1. \tag{9}$$

We note that the order rate is also heteroskedactic. In periods when $k=1$, it exhibits a rich behavior depending upon ϕ, but when $k<P$, the order rate variance is decreasing in $\{k,\ L\}\forall\phi$.

When measuring the order rate variance over every t (ignoring the cyclical heteroskedacticity), the effects of using a time-varying safety stock becomes apparent. Let $\Delta S_t = S_{t+L} - S_{t+L-1}$ denote the quantity that is required to be ordered in period t to change the safety stock to its desired level. Obviously, the variance of the ΔS_t is zero when measured once every P periods, but when measured over every t, the average variance of the orders required to change the safety stock quantity is

$$\sigma_{\Delta S}^2 = P^{-1}\sum_{t=1}^{P}\Delta S_t^2, \tag{10}$$

where we exploit the cyclical nature of the safety stock $\overline{\Delta S} = P^{-1}\sum_{t=0}^{P}\Delta S_t = 0$. Then, the total order variance, measured over every t, is $\sigma_{\Delta S}^2 + P^{-1}\sum_{k=1}^{P}\sigma_{O,k}^2$.

5. The general case for auto-correlated demand

The inventory variance can also be determined for the cyclical case if we can simply define the autocorrelation function of demand $R(t)$. The variance of the demand process is known to be

$$\frac{\sigma_d^2}{\sigma_\varepsilon^2} = \sum_{n=0}^{\infty} R^2(n). \tag{11}$$

Let us now consider how inventory is affected by demand. Let the inventory deficit after q periods of cumulative demand that results from a single unit impulse be denoted as $S(q) = -\sum_{n=0}^{q-1} R(n)$. Under an optimal policy, with the cyclical planning constraint, the maximum inventory shortfall resulting from a demand impulse that occurred k periods before the ordering occasion is $S(k+L)$. Since we have $k + L$ of these

$$\frac{\sigma_{i,k}^2}{\sigma_\varepsilon^2} = \sum_{m=1}^{k+L} S^2[m]. \tag{12}$$

Because $d \in \mathbb{R}$, we know that the inventory variance is increasing in both k and L. We can also conclude that the average standard deviation of inventory increases with P.

6. Properties of the AR(1) policy with MMSE forecasting

From our investigation above, we can identify several important properties of the AR(1)-optimal policy under cyclical reordering. While this policy may not be a natural one to implement, it is important because it puts a definitive lower bound on the inventory ripple effect. For the general case, even if demand is not AR(1), we know that $\forall\{k,\ L\}$, the inventory variance is increasing. This also implies that the average inventory standard deviation, to which the inventory cost function is linearly related, increases with P. Therefore, the optimum reorder cycle length equals one ($P^* = 1$) when no ordering or cycle costs are present.

Before investigating AR(1) demand, consider the i.i.d. case, i.e. $\phi = 0$,

$$\lim_{\phi \to 0} \frac{\sigma_i^2}{\sigma_\varepsilon^2} = \lim_{\phi \to 0} \sum_{n=0}^{L+k-1} \left(\sum_{x=0}^{n} \phi^x \right)^2 = k + L. \tag{13}$$

(13) is identical to the corresponding expression in Hedenstierna & Disney (2013), where only i.i.d demand was investigated. Let us also consider the cases which limit demand as a stable and invertible AR(1) signal, i.e. $-1 < \phi < 1$. The upper limit,

$$\lim_{\phi \to 1} \frac{\sigma_i^2}{\sigma_\varepsilon^2} = \lim_{\phi \to 1} \sum_{n=0}^{L+k-1} \left(\sum_{x=0}^{n} \phi^x \right)^2 = \frac{1}{6}(k+L)(1+k+L)(1+2k+2L), \tag{14}$$

shows that the inventory variance is a cubic function of both k and L. We can consider this as an upper bound for the inventory variance, when $0 < \phi \le 1$. The other limiting case, when $\phi = -1$ is

$$\lim_{\phi \to -1} \frac{\sigma_i^2}{\sigma_\varepsilon^2} = \lim_{\phi \to -1} \sum_{n=0}^{L+k-1} \left(\sum_{x=0}^{n} \phi^x \right)^2 = \frac{1-(-1)^{k+L}}{4} + \frac{k+L}{2}. \tag{15}$$

We notice that this variance ratio is always smaller than the i.i.d. case, except when $k = 1$. It also has the peculiar property of assuming the same value twice in a row, only increasing with every other k. If L is even, this transition occurs whenever k switches to an odd number, and conversely, when L is odd, the switch occurs when k becomes even. It is noteworthy, that the upper limit $\phi = 1$ gives an inventory variance that grows cubically with both k and L, while the i.i.d. case gives linear growth with a gradient of unity, and the lower limit, $\phi = -1$, gives a linear growth of ½, when k and L are measured in increments of two. The inventory variance of these limiting cases, and settings where demand is stationary and non-stationary, are presented in Figure 3. From numerical observations, we find that the least ripple effect is encountered in the range $-1 \le \phi < 0$.

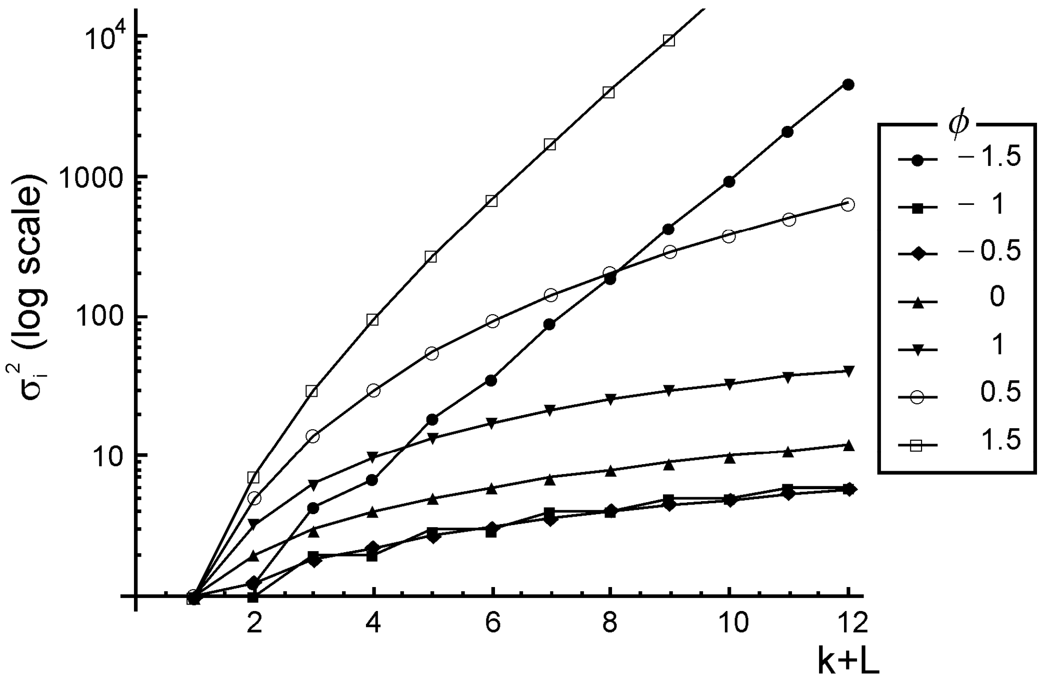

Figure 3. The inventory ripple effect under non-stationary demand

For a comprehensive look at the cases when $|\phi| < 1$, we may consult Figure 4, where the inventory variance again is given as a function of k+L and ϕ. Since the inventory ripple effect indicates multiple variances, it should be read from $L+1$ to $L+P$ on the y-axis. Noticeably, the inventory variance increases the most for large positive values of ϕ, and to a lesser degree for negative ϕ.

The influence of the ripple effect on the safety stock and its associated costs still remains to be demonstrated. Let us first assume a zero-lead time, $L = 0$, as this gives a clear illustration of the inventory ripple effect. It is also possible to consider non-zero lead times, but in that case the relative cost of the ripple effect becomes smaller. Further assume the inventory cost factors $\{H = 1,\ B = 9\}$, which imply an optimal availability of 90%. We shall consider the difference between a constant safety stock setting designed to cover the entire protection interval, i.e. $L+P$, and an optimal safety stock setting, which due to the heteroskedacticity of the inventory deviations from the norm, must be time-varying (see (3)). The results for some different reorder cycle lengths (P), is shown in Table 3.

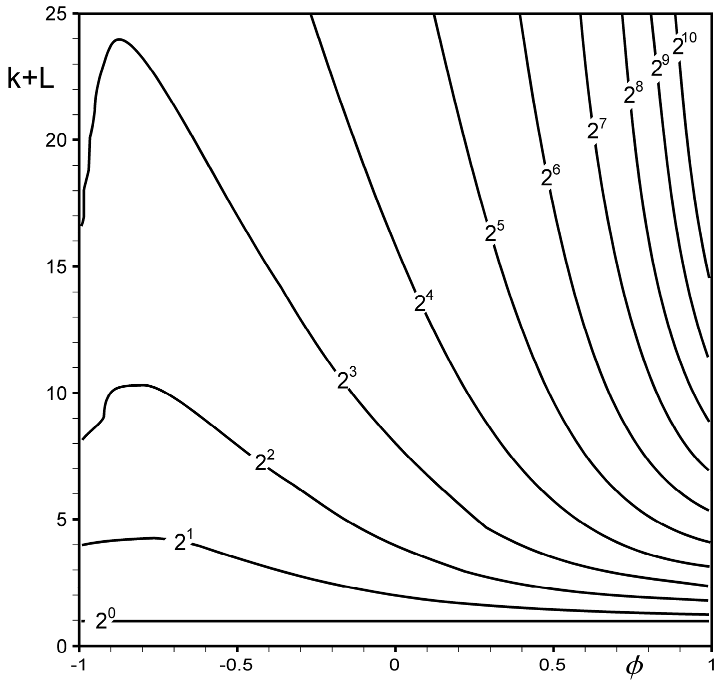

Figure 4. The inventory variance development for stable demand

		ϕ				
	P	-0.9	-0.5	0	0.5	0.9
SS_P^*	1	1.75	1.75	1.75	1.75	1.75
	5	2.39	2.47	3.31	5.32	8.98
	10	2.98	3.25	4.60	8.26	19.67
	20	3.87	4.43	6.45	12.24	40.00
SS_k^*	1	1.75	1.75	1.75	1.75	1.75
	5	2.22	2.31	2.94	4.26	6.26
	10	2.73	2.90	3.94	6.52	12.83
	20	3.45	3.82	5.41	9.75	26.18

Table 3. Costs resulting from conventional and optimal configurations.

It is clear that the optimal setting is superior to the sub-optimal one in all cases except when $P=1$ when they are equal (also note that the cost does not change with ϕ for this setting, due to the lead time being zero). Analogous to the inventory variance, we find that inventory costs are lower for strongly negatively correlated demand, and small values of ϕ, while costs increase with ϕ and *P*. For monthly planning $(P=20)$, strongly positively correlated demand has a 33% higher inventory costs with the conventional setting over the optimal one, but when demand is uncorrelated this difference is 18%, and for strongly negatively correlated demand, only 16%.

7. Concluding remarks

In many cases, supply chain planning takes place once per week or less frequently (34% of 292 Swedish companies). If planning occurs less frequently than inventory inspections, a ripple effect occurs, causing inventory deviations to change over time. Infrequent reordering is not optimal from an inventory cost perspective but when it is used, constant safety stock levels lead to suboptimal costs. Instead a time-varying safety stock lends the lowest overall costs. The inventory ripple effect is especially pronounced when demand is positively auto-correlated, which is often the case for real demand. The ripple effect we have indicated is the bare minimum an inventory system will experience under linear reordering policies. Therefore, it is likely that the ripple effect will be worse than we have indicated, for those companies that plan once per week or less frequently. Our numerical example indicates a large cost differential between supply chains that would consider the effect, and those that ignore it. It should therefore be of interest, to managers and academics alike, to find evidence of the inventory ripple effect in the real world, and to design pragmatic policies to minimize its consequences.

8. References

Axsäter, S., 2006. *Inventory Control*. Springer, New York.

Burbidge, J.L., 1988. *Production Flow Analysis for Planning Group Technology*. University Press, Oxford.

Chiang, C., 2006. Optimal ordering policies for periodic-review systems with replenishment cycles. *European Journal of Operational Research*, 170, 44–56.

Chiang, C., 2007. Optimal Ordering Policies for periodic-review systems with a refined intra-cycle time scale. *European Journal of Operational Research*, 177, 872-881.

Flynn, J. and Garstka, S., 1997. The optimal review period in a dynamic inventory model. *Operations Research*, 45 (5), 736–750.

Flynn, J., 2001. Selecting review periods for a coordinated multi-item inventory model with staggered deliveries. *Naval Research Logistics*, 48 (5), 430–449.

Graves, S.C., 1999. A single-item inventory model for a non-stationary demand process. *Manufacturing and Service Operations Management*, 1 (1), 50-61.

Hedenstierna, C.P.T. and Disney, S.M., 2013. Inventory ripple effect in periodic inventory control systems. Working Paper, Cardiff Business School.

King, P.L., (2009). *Lean for the Process Industries*. CRC Press, New York.

Lee, H.L., So, K.C. and Tang, C.S., 2000. The value of information sharing in a two-level supply chain. *Management Science*, 46 (5), 626-643.

Jonsson, P. and Mattsson, S.-A., 2013. Lagerstyrning i svensk industri: 2013 års användning, användningssätt och trender, Chalmers tekniska högskola, Gothenburg. *In Swedish.*

Modigliani, F. and Hohn, F. E., 1955. Production Planning Over Time and the Nature of the Expectation and Planning Horizon. *Econometrica*, 23 (1), 46-66.

Ohno, T., 1988. *Toyota Production System: Beyond Large-scale Manufacturing*. Productivity Press, New York.

Shingo, S., 1989. *The Toyota Production System from an Industrial Engineering Viewpoint*. Productivity Press, New York.

Silver E. S. and Robb, D. J., 2008. Some insights regarding the optimal reorder period in periodic review inventory systems. *International Journal of Production Economics*, 112, 354-366.

Sloan, A.P., 1963. *My Years with General Motors*. Doubleday, New York.

Tang, O. and Grubbström, R.W., 2002. Planning and re-planning the master production schedule under demand uncertainty. *International Journal of Production Economics*, 78, 323-334.

Tsypkin, Y.Z.,1964. *Sampling Systems Theory and its Application, Vol. 1*. Pergamon Press, Oxford.

Waters, D., 2003. *Inventory Control and Management*. John Wiley & Sons, New Jersey.