Hedenstierna, C.P., and Disney, S.M., (2016), "Inventory performance under staggered deliveries and auto-correlated demand", European Journal of Operational Research, in press.

Inventory performance under staggered deliveries and autocorrelated demand

Carl Philip T. Hedenstierna*, Stephen M. Disney

Logistics System Dynamics Group, Cardiff Business School, Cardiff University, Aberconway Building, Colum Drive, CF10 3EU, UNITED KINGDOM

Abstract

Production plans often span a whole week or month, even when independent production lots are completed every day and service performance is tallied daily. Such policies are said to use *staggered deliveries*, meaning that the production rate for multiple days are determined at a single point in time. Assuming autocorrelated demand, and linear inventory holding and backlog costs, we identify the optimal replenishment policy for order cycles of length P. With the addition of a once-per-cycle *audit cost*, we optimize the order cycle length P^* via an inverse-function approach. In addition, we characterize periodic inventory costs, availability, and fill rate. As a consequence of staggering deliveries, the inventory level becomes cyclically heteroskedastic. This manifests itself as ripples in the expected cost and service levels. Nevertheless, the cost-optimal replenishment policy achieves a constant availability by using time-varying safety stocks; this is not the case with suboptimal constant safety stock policies, where the availability fluctuates over the cycle.

Keywords: inventory, autoregressive demand, order-up-to-policy, staggered deliveries, planning cycles

1. Introduction

Our paper is motivated by a recent experience with two manufacturers who planned on a weekly basis. One of the manufacturers had a local customer,

*Corresponding author, tel: +44 (0)2920 876 310, fax: +44 (0)2920 874 301

Email addresses: HedenstiernaCP@cardiff.ac.uk (Carl Philip T. Hedenstierna), DisneySM@cardiff.ac.uk (Stephen M. Disney)

just a few hours away by truck. The other manufacturer had a customer on another continent, six weeks away by container ship. Production and distribution planning in each case was conducted once per week in SAP, which automated the order-up-to (OUT) replenishment policy calculations. For the manufacturer with the remote customer, products were loaded into 40ft containers that were dispatched at the end of the week. This matches well the literature on periodic review inventory systems, such as the base stock policy and the OUT policy.

For the manufacturer with the local customer, whilst production and distribution plans where still generated at the start of the week, daily production lots were loaded onto trucks that were dispatched to the local customer at the end of each day. Here we have a planning cycle of one week, and an inventory inspection period of one day. This scenario, with mismatched order cycles and inspection periods, is quite popular in the just-in-time / lean environment. Nevertheless, this phenomenon has received little attention and is not well understood. The case when goods can be received every day, but when orders may be placed only once every P days is called *staggered deliveries*.

1.1. Literature review

The pioneering work on staggered deliveries was done by Flynn and Garstka (1990), who developed an inventory-optimal control policy under staggered deliveries and independent and identically distributed (i.i.d.) demand. Flynn and Garstka (1990) call the problem of finding the optimal replenishment policy the *ordering problem*, while they call the problem of identifying the optimal planning cycle length the *auditing problem*. They find for the infinite horizon case, with a fixed audit cost (a cost incurred when production plans are generated), that the optimal policy is a base-stock policy (a variant of the OUT policy). This model was expanded to include a fixed cost for each cycle with a non-zero order quantity; then an (s,S)-type policy is optimal (Flynn, 2000). Prak et al. (2015) investigated the same problem with an arbitrary order cycle length (which they termed *periodic review*) and continuously staggered deliveries (termed *continuous ordering*). They found that the optimal ordering policy for the continuous-time problem is OUT.

Chiang (2001) investigated an OUT (termed $[R,S]$) inventory system with service constraints, inventory holding costs, an audit cost, and a fixed cost per receipt. The study demonstrated that lot splitting, i.e. staggering, may lead to reduced costs, particularly when the audit cost is high. A particular type of staggered delivery system was studied by Chiang (2009), who considered two variations of a staggered policy. In one of these, which we shall refer to as Chiang's simplified policy, the order quantity was arbitrary in the first period of the order cycle, and constant in the remaining periods. The other policy included a non-negative ordering constraint; it also permitted a variable order quantity (bounded from above by some constant) for several periods following the first, potentially extending to every period in the cycle.

In staggered systems, and in some periodic review systems, it is common to ignore cycle stock and to apply inventory costs to the inventory level at the end of each replenishment cycle, just before new receipts arrive. This can be compared with the non-staggered models in Silver and Robb (2008) and Chiang (2007), which permit inventory inspections (with associated costs) between two subsequent receipts.

The lean literature often suggests planning as frequently as set-up costs allow (Bicheno and Holweg, 2009, p. 155; Burbidge, 1983), but this is not always reflected in practice. A recent survey of 292 Swedish companies revealed that when using the OUT policy, 21% planned on a daily basis or more frequently, 37% planned on a weekly basis, and 42% planned fortnightly or less frequently (Jonsson and Mattsson, 2013). The same survey reported that, over the period 1993-2013, planning cycles have been shrinking. However Flynn and Garstka (1990) showed that an optimum order cycle length, $P^* > 1$, may exist if there is an audit cost — a fixed cost per planning cycle. Approximate solutions for finding P^* were given in Flynn and Garstka (1997), Flynn (2000), and Flynn (2008). For i.i.d. demands the problem of identifying P^* for coordinating multiple items was investigated in Flynn (2001). Lian et al. (2006) used simulation to find the optimal length of the planning cycle.

With autocorrelated demand, the demand in one period may be affected by random demand fluctuations that have occurred in previous periods. Autocorrelated demand has been observed in industry (Erkip et al., 1990), and it is known that autocorrelation may affect inventory cost and service levels significantly (Lee et al., 2000). Zhang (2007) considers the consequences of heteroskedastic demand on the OUT policy, whereas we notice, as did Lian et al. (2006) and Flynn (2008), that heteroskedasticity of inventory levels can be induced by the staggered delivery mechanism.

1.2. Contribution

This paper extends the inventory-optimal model of Flynn and Garstka (1990) to autocorrelated demand, and demonstrates how availability and fill rates fluctuate between the days in a delivery cycle. This is linked to the heteroskedastic inventory variance, which depends on demand autocorrelation.

Lian et al. (2006) and Flynn (2008) find that time-varying safety stocks are optimal for i.i.d. demand. We build on these results and identify the optimal time-varying safety stocks for autocorrelated demand.

We identify the inventory performance in terms of cost, availability, and fill rate when optimal safety stocks are used, and compare this with the performance of two practical (but suboptimal) constant safety stock settings. The problem of the optimal order cycle length P^* has already been considered by Flynn (2000), Flynn (2001) and Lian et al. (2006) for i.i.d. demand. We generalise this to autocorrelated demand and characterise the exact solution to P^* by an inverse function approach.

1.3. Paper structure

The structure of this paper is as follows. In section 2 we describe the model in a natural seven-day setting to introduce the concept of staggering. Then, in section 3, we derive the optimal ordering rule for arbitrary planning cycle lengths and autocorrelated demand. Here we also define how performance is measured via inventory costs, availability and the fill rate. We also provide an exact solution to the problem of finding the length of the optimal planning cycle for autocorrelated demand. Section 4 considers the staggered OUT policy under first-order autoregressive (AR(1)) demand and provides some numerical insights. First we operationalize the staggered delivery policy to demonstrate how one actually generates production plans. Then we investigate the inventory cost and service level consequences of the staggered policy. Finally we provide an example of how to find the optimal planning cycle length. Section 5 concludes. Proofs are presented in appendices.

2. Model description in a natural setting

To ease understanding, let us initially define the model in a weekly setting of seven days, where we plan once per week, but produce every day. Later, we generalize this to arbitrary planning periods of length P, but for now we consider the planning cycle to be seven days long.

Every morning the inventory level is tallied. If it is Monday, a production plan is made immediately after the inventory inspection. This production plan contains seven orders, to cover an entire week of production. This reflects the planning cycle length $P = 7$. In more general terms, staggering deliveries means that we must determine the production rate for P days once every P periods. Between two such occasions, no new production plans are calculated, as we are committed to the established plan; figure 1 illustrates this. Let t

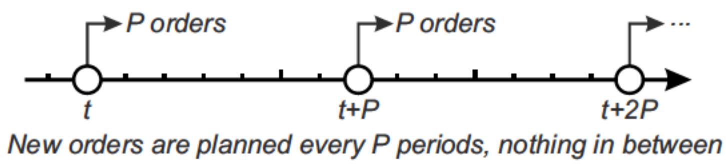

Figure 1: A total of P orders are issued every P periods.

number the individual days (periods), and let Mondays occur when t/P is an integer. Suppose it is the start of Monday morning, and that we must plan the orders for the next cycle $\{O_{t,1}, O_{t,2}, ..., O_{t,7}\}$, numbered in the same sequence as they will be produced. Every order O corresponds to a future inventory receipt R, according to

$$O_{t,k} = \begin{cases} R_{t+\tau} = R_{t+k+L} & \text{when } t/P \text{ is an integer,} \\ \varnothing & \text{otherwise;} \end{cases} \tag{1}$$

where $k \in \{1, 2, ..., P\}$ is the order release offset due to staggering, L is a non-negative integer lead time (i.e. the delay until the first order can be received), and $\tau = k + L$ is the effective lead time. In this case, $L = 4$, meaning that these orders will register as received in the periods $\{t + 5, t + 6, ..., t + 11\}$, as figure 2 illustrates.

The policy operates on the master production scheduling (MPS) level, and is therefore concerned with due-date setting. Order releases are typically handled by material requirements planning, using the MPS (generated by the policy) as input (Pinedo, 2009, p. 10). Immediately when a cycle's order quantities have been determined, the required receipt rates R_{t+k+L} are disaggregated into a detailed schedule of individual jobs. These jobs must be released so that the production completions between $t + \tau - 1$ and $t + \tau$ equal $R_{t+\tau}$. We need not know the exact timing of job releases, as long as the planned quantity arrives in the right period. The lead time L is the delay until the first order is received, reflecting the time to effect a new production plan, including the time to schedule jobs, to allocate resources, and to produce. Note that if $L = 0$, the first order would be released to production immediately, and be completed in less than one day. It would therefore contribute to the inventory level measured at time $t + 1$.

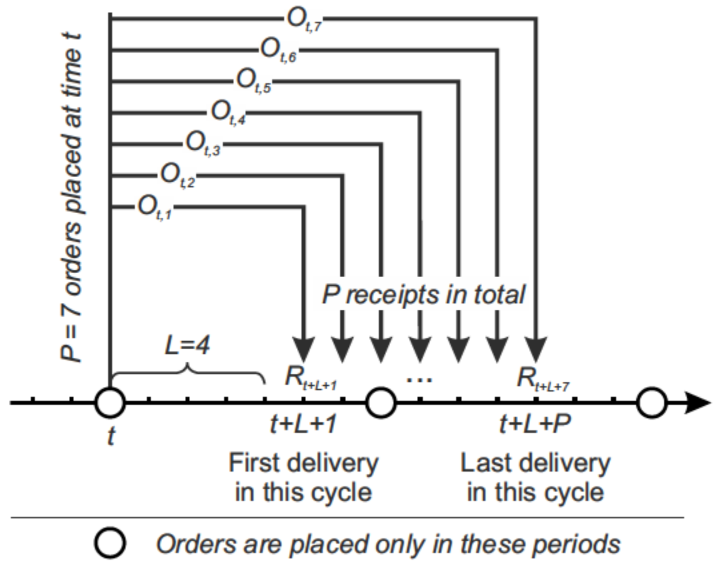

Figure 2: Staggering means that even though we plan intermittently, we receive orders every period.

The receipts resulting from the staggered deliveries are placed in inventory. We assume that there is no shrinkage, and that the inventory level increases with receipts (R_t) and decreases with demand (D_t),

$$I_t = I_{t-1} + R_t - D_t. \tag{2}$$

If there are no goods on-hand, the excess demand is backlogged and subtracted from the inventory level; then the inventory level falls below zero. Backlogged demand is satisfied immediately when new goods are received. We assume that the system is linear, therefore negative orders are permitted, reflecting costless returns (in distribution scenarios), or that goods are sold off at a price equalling the variable cost of production (in manufacturing scenarios). Negative demand is also permitted, reflecting returns from customers to our inventory. Both the storage and the production facilities have unlimited capacity.

3. The optimal ordering rule

The periodic inventory costs consist of a holding cost H per unit of inventory, and a backorder cost B per unit of unsatisfied demand, which we model as negative inventory,

$$j(I_t) = H\,(I_t)^+ + B\,(-I_t)^+ = H\,I_t + (B+H)\,(-I_t)^+ \,. \tag{3}$$

where $(X)^+ = \max(X,0)$. Let the lead-time demand be denoted as $F_{t,\tau} = \sum_{n=1}^{\tau} D_{t+n}$, and the work-in-progress as $W_{t,\tau} = \sum_{n=1}^{\tau-1} R_{t+n}$; then the inventory level can be expressed as $I_{t+\tau} = I_t + W_{t,\tau} + O_{t,k} - F_{t,\tau}$, when t/P is an integer. Equation (3) is thus convex in $O_{t,k}$, and the convexity is preserved when taking the expectation,

$$J(I_{t+\tau}) = E\left[j(I_{t+\tau})|D_t, D_{t-1}, \cdots\right]. \tag{4}$$

As $O_{t,k}$ can be set freely, it can eliminate the influence of the system state $I_t + W_{t,\tau}$ on $I_{t+\tau}$ before costs are incurred. Therefore, $J(I_{t+\tau})$ depends only on the decision variable $O_{t,k}$, and on the lead-time demand $F_{t,\tau}$. Consequently, the optimal order policy is myopic, meaning that the optimal solution for an n-period problem can be found by solving n independent single-period problems (Heyman and Sobel, 1984, p. 63-71). In practice, we need only to consider the immediate consequences of each decision, as it has no bearing on the cost incurred in other (future) periods. The myopic nature of the optimal policy can be verified by setting up a dynamic program. See Bertsekas (2005, p. 162-164) for a problem with the same cost structure, compatible with autocorrelated demand, but without staggering, and Flynn and Garstka (1990) for a dynamic program with staggering, but without demand autocorrelation. We exclude a dynamic programming formulation, as it adds complexity without providing commensurate insights.

The expected inventory cost (4) is convex, and a minimum exists because both B and H are positive. Therefore, there exists an optimal expected inventory level that minimizes $J(I_{t+\tau})$. This optimum is referred to as the *safety stock*, $I^*_{t+\tau}$. It is identified in the following lemma.

Lemma 1 (The optimal safety stock level)**.** The expected inventory cost, $J(I_{t+\tau})$, is minimized when

$$I^*_{t+\tau} = \Phi^{-1}_{I_{t+\tau}}\left(\frac{B}{B+H}\right), \tag{5}$$

where $\Phi^{-1}_{I_{t+\tau}}$ is the inverse of the inventory level's cumulative distribution function at time $t+\tau$.

Proof. The expected inventory cost J is structurally identical to the single-period newsvendor problem, with B and H representing the underage and overage costs. See Hopp and Spearman (2008, p. 67-69) or Cachon and Terwiesch (2009, p. 232-236) for a simple proof. ∎

With $I^*_{t+\tau}$ known, all that remains is to specify the policy that sets $E\left[I_{t+\tau}|D_t,\, D_{t-1}\,\cdots\right] = I^*_{t+\tau}$. This is done in the following theorem.

Theorem 2. *When there are P staggered orders per cycle, the expected inventory cost is minimized by the policy*

(a) for the first order in a cycle, when $k = 1$,

$$O_{t,1} = \hat{F}_{t,L+1} + I^*_{t+L+1} - I_t - W_{t,L+1}, \tag{6}$$

where $\hat{F}_{t,\tau} = E\left[F_{t,\tau}|D_t,\, D_{t-1},\cdots\right]$ is the forecast;

(b) for all subsequent orders in the cycle, when $k > 1$,

$$O_{t,k} = I^*_{t+\tau} - I^*_{t+\tau-1} + \hat{D}_{t,\tau}, \tag{7}$$

where $\hat{D}_{t,\tau} = E\left[D_{t+\tau}|\{D_t, D_{t-1},\ ...\}\right]$ is the singe-period forecast, made at time t, for $D_{t+\tau}$.

Proof. Given in Appendix A. ∎

The policy can be interpreted as an OUT policy with increasing lead times over the cycle and a simplified ordering rule for all periods but the first. Note that our simplified rule, (7), is different from the simplified rule in Chiang (2009), who assumes that $O_{t,2} = O_{t,3} = \cdots$. Instead, we order the single-period forecast of demand plus any desired change in safety stock, $I^*_{t+\tau} - I^*_{t+\tau-1}$. When $P = 1$ the policy simplifies to the regular OUT policy.

3.1. Demand specification

To gain further insights about the optimal policy and its dynamic performance, we shall assume that demand is autocorrelated,

$$D_t = \mu + \sum_{n=0}^{\infty} \varepsilon_{t-n}\theta_n, \tag{8}$$

where θ_n is the autocorrelation function, $\mu = E[D_t]$, and ε_t is an independent and identically distributed (i.i.d.) random variable drawn from the normal distribution. We call ε_t the *error term*. It has a mean of zero and a variance of σ_ε^2. The mechanism of such demand processes is well documented in Box and Jenkins (1976), where they are described as moving-average processes. One important property of this type of demand signal is that its variance can be obtained as $\text{Var}(D) = \sigma_\varepsilon^2 \sum_{m=0}^{\infty} \theta_m^2$.

To find I_t^*, we exploit the linear system assumption, and the assumption of normally distributed, i.i.d. random error terms. Consequently the inventory level follows a normal distribution, hence the mean and variance are sufficient to specify the inventory distribution. The mean inventory can be set to I_t^* with R_t, but the variance is a function of τ and the demand process. Before identifying the inventory variance required to calculate I_t^*, let us define the service levels.

3.2. Service levels

Not only does the optimal safety stock minimize the total cost of the system, it also sets the system's availability (Silver et al., 1998) to the critical ratio $B/(B+H)$. The availability (α), or type 1 service level, refers to the probability of not encountering a stock-out in any given period,

$$\alpha = \Pr(I_t \geq 0). \tag{9}$$

The fill rate (β), or type 2 service level, is sometimes considered a more appropriate measure in customer-facing settings, as it measures the fraction of demand fulfilled immediately from stock (Johnson et al., 1995). The exact formula for the fill rate when demand is autocorrelated and possibly negative (Disney et al., 2015) is of the following form:

$$\beta = \frac{E\left\{[\min(D_t, I_t + D_t)]^+\right\}}{E\left[(D_t)^+\right]}. \tag{10}$$

This exact fill rate takes the expectation of immediately satisfied demand, and divides it by the expected positive demand. This works well when demand is weakly stationary, but does not work for nonstationary demand. If the variables in this expression are normally distributed, we can obtain the fill rate via the following lemma.

Lemma 3. The exact fill rate for normally distributed demand, where periods with negative demand do not contribute to the fill rate, is

$$\beta = \frac{\int_{x=0}^{\infty} \varphi^{-}(x) x \, dx}{\sigma(D_t)\, G\left[-\mu/\sigma(D_t)\right]}. \tag{11}$$

Here $\sigma(D_t) = \sqrt{\text{Var}(D_t)}$ is the standard deviation of D_t, $\varphi^{-}(x)$ is the probability distribution function (pdf) of the minimum of the normally distributed bivariate random variables D_t and $(D_t + I_t)$. $G(x) = \varphi(x) - x[1 - \Phi(x)]$ is the standard normal loss function, where $\varphi(x)$ is the standard normal pdf, and $\Phi(x)$ is the standard normal cumulative density function (Axsäter, 2006, p. 91).

Proof. Given in Disney et al. (2015). ■

Remark. The pdf of the minimum of bivariate random normal variables, $\varphi^{-}(x)$, is given in Cain (1994), as $\varphi^{-}(x) = \varphi_1^{-}(x) + \varphi_2^{-}(x)$, where

$$\varphi_1^{-}(x) = \frac{\varphi\left(\frac{x - E[I_t + D_t]}{\sigma(I_t + D_t)}\right)}{\sigma(I_t + D_t)} \; \Phi\left[\frac{\rho\left(\frac{x - E[I_t + D_t]}{\sigma(I_t + D_t)}\right) - \frac{x - \mu}{\sigma(D_t)}}{\sqrt{1 - \rho^2}}\right], \tag{12}$$

$$\varphi_2^{-}(x) = \frac{\varphi\left(\frac{x - \mu}{\sigma(D_t)}\right)}{\sigma(D_t)} \; \Phi\left[\frac{\rho\left(\frac{x - \mu}{\sigma(D_t)}\right) - \frac{x - E[I_t + D_t]}{\sigma(I_t + D_t)}}{\sqrt{1 - \rho^2}}\right], \tag{13}$$

where the correlation coefficient is

$$\rho = \frac{\text{Cov}\left(I_t + D_t, D_t\right)}{\sqrt{\text{Var}(I_t + D_t)\text{Var}(D_t)}}. \tag{14}$$

It is often necessary to evaluate (11) numerically. This is usually done with software like Mathematica or Matlab, but it can also be achieved with Microsoft Excel using the macro provided in Disney et al. (2015).

To calculate the exact fill rate we must know the variances $\text{Var}(I_t)$ and $\text{Var}(I_t + D_t)$, and also the correlation coefficient ρ. For autocorrelated demand, we identify these according to

Theorem 4. *If planning took place at time $t - \tau$, the characteristics of the inventory level are as follows:*

(a) The inventory variance is

$$\mathrm{Var}(I_t) = \sigma_\varepsilon^2 \sum_{n=0}^{\tau-1} \left(\sum_{m=0}^{n} \theta_m \right)^2 . \tag{15}$$

(b) The variance of $I_t + D_t$ is

$$\mathrm{Var}(I_t + D_t) = \sigma_\varepsilon^2 \left\{ \left[\sum_{m=1}^{\tau-1} \left(\sum_{n=0}^{m-1} \theta_n \right)^2 \right] + \sum_{x=\tau}^{\infty} \theta_x^2 \right\} . \tag{16}$$

(c) The covariance between demand, D_t, and $I_t + D_t$, is

$$\mathrm{Cov}(D_t, I_t + D_t) = \sigma_\varepsilon^2 \left\{ \left[\sum_{n=1}^{\tau-1} \left(-\sum_{m=0}^{n-1} \theta_m \right) \theta_n \right] + \sum_{x=\tau}^{\infty} \theta_x^2 \right\} . \tag{17}$$

Proof. Presented in Appendix B. ■

The inventory variance (15) increases in τ, regardless of θ_t, and is finite for all demands, stationary or nonstationary. The variance of the state variable $I_t + D_t$ is also increasing in τ, but is only finite for stationary demand. The covariance (17) between demand and initial inventory exists only for stationary demand.

The main insight from (15) is that the inventory variance increases over the cycle. As inventory costs are minimized when $\Pr(I_t \geq 0) = B/(B+H)$, we find a *time-varying* safety stock to be optimal. This safety stock is increasing in τ. It is also clear from (15) that autocorrelation can amplify or attenuate inventory heteroskedasticity.

3.3. Total cost and the optimal reorder cycle length, P^*

Under normally distributed demand and linear transformations, the inventory level is also normally distributed. The expected inventory cost is

$$\begin{aligned} J(I_{t+\tau}) &= E[I_{t+\tau}]H - \frac{B+H}{\sigma_{i,k}} \int_{-\infty}^{0} \varphi\left(\frac{x - E[I_{t+\tau}]}{\sigma_{i,k}} \right) x \, dx \\ &= HE[I_{t+\tau}] + (B+H)\sigma_{i,k}\, G\left(\frac{E[I_{t+\tau}]}{\sigma_{i,k}} \right), \end{aligned} \tag{18}$$

where $E[I_{t+\tau}]$ denotes the safety stock, and $\sigma_{i,k} = \sqrt{\mathrm{Var}(I_{t+k+L})}$. As the error terms are i.i.d., $J(I_t) = J(I_{t+P})$. Therefore, the average cost is obtained by averaging over P successive periods. When the optimal safety stocks $I^*_{t+\tau}$ are used, the average cost from (18) simplifies to

$$J^*_P = \frac{1}{P}\sum_{k=1}^{P} J(t+k) = \bar{\sigma}_{i,P}(B+H)\varphi\left[\Phi^{-1}\left(\frac{B}{B+H}\right)\right], \tag{19}$$

where $\bar{\sigma}_{i,P} = P^{-1}\sum_{k=1}^{P}\sigma_{i,k}$ is the average standard deviation of the inventory level. This variable is essential for characterizing P^*.

Consider a fixed audit cost per cycle, V, leading to an average audit plus inventory cost per period of $C_P = J^*_P + V/P$. Let $\lambda = V/\psi$ where

$$\psi = V + (B+H)\varphi\left[\Phi^{-1}\left(\frac{B}{B+H}\right)\right]. \tag{20}$$

The total cost can then be expressed as a linear function of $\lambda \in [0,1]$,

$$C_P(\lambda) = \psi\left[\bar{\sigma}_{i,P} + \lambda\left(P^{-1} - \bar{\sigma}_{i,P}\right)\right]. \tag{21}$$

With this formulation, it is possible to find a cost combination λ_P for which P minimizes the total cost.

Theorem 5. *When $\sigma_{i,P}$ is increasing in P, the order cycle length P minimizes $C_P(\lambda)$ for $\lambda \in [\lambda_{P-1}, \lambda_P]$ where $\lambda_0 = 0$ and*

$$\lambda_P = 1 - \frac{1}{1 + P\left(\sigma_{i,P+1} - \bar{\sigma}_{i,P}\right)}. \tag{22}$$

Proof. Let λ_P denote the intersection $C_P(\lambda_P) = C_{P+1}(\lambda_P)$. Solving for λ_P provides (22). Suppose λ_P is increasing in P. Then, as $P = 1$ minimizes the cost for $\lambda \in [0, \lambda_1]$, the reorder period P minimizes $C_P(\lambda)$ for $\lambda \in [\lambda_{P-1}, \lambda_P]$. To see that λ_P is increasing in P, recall that $\sigma_{i,P} \leq \sigma_{i,P+1}$. This provides

$$P\sigma_{i,P+1} - \sum_{k=1}^{P}\sigma_{i,k} \leq (P+1)\,\sigma_{i,P+2} - \sum_{n=1}^{P+1}\sigma_{i,n}, \tag{23}$$

which leads to $P\left(\sigma_{i,P+1} - \bar{\sigma}_{i,P}\right) \leq (P+1)\left(\sigma_{i,P+2} - \bar{\sigma}_{i,P+1}\right)$. This is equivalent to $\lambda_P \leq \lambda_{P+1}$, completing the proof. ■

This procedure lets us specify a P and provides the range of cost configurations $[\lambda_{P-1}, \lambda_P]$ for which this P is optimal. Through this indirect approach, several properties of P^* are revealed: it is increasing in the audit cost V, and it is decreasing in the factors that drive inventory cost, namely B, H, L, and σ_ε. This follows from the influence of (15) and (19) on λ. Furthermore, given a cost balance λ and an arbitrary P, it is immediately clear if $P < P^*$, $P = P^*$, or $P > P^*$, as λ_P is increasing in P. These observations hold for generally autocorrelated demand as (15) is increasing in τ.

To find P^*, it is sufficient to identify two values P_1 and P_2, such that $\lambda_{P_1} \leq \lambda \leq \lambda_{P_2}$; a binary search between these values then provides the optimum. As an alternative, we may plot the first few λ_P, and then seek P^* graphically. This simpler approach does not guarantee that P^* will be in the range plotted, but it is nonetheless reasonable when the audit cost is moderate in relation to the inventory cost.

4. The staggered OUT policy for first-order autoregressive demand

To better understand the model, it is helpful to consider a simple case. Here we choose the AR(1) demand process. It is stationary and invertible for $|\phi| < 1$ (Box and Jenkins, 1976). The corollary below provides necessary expressions for calculating inventory costs, availability, the fill rate, and λ_P.

Corollary 6 (AR(1) demand)**.** Using $\theta_m = \phi^m$ as the autocorrelation function of demand in theorem 4, we obtain the following variance expressions:

(a) The inventory variance,

$$\mathrm{Var}\,(I_t) = \sigma_\varepsilon^2 \left[\frac{\tau}{(\phi-1)^2} + \frac{\phi(\phi^\tau - 1)(\phi^{\tau+1} - \phi - 2)}{(\phi+1)\,(\phi-1)^3} \right]. \tag{24}$$

(b) The variance of $I_t + D_t$,

$$\mathrm{Var}\,(I_t + D_t) = \sigma_\varepsilon^2 \left[\frac{\tau}{(\phi-1)^2} - \frac{2\phi^\tau}{(\phi-1)^3} + \frac{1 + \phi\,(2 - (\phi-2)\phi^{2\tau})}{(\phi-1)^3\,(\phi+1)} \right]. \tag{25}$$

(c) The covariance between $I_t + D_t$ and D_t,

$$\mathrm{Cov}(D_t, I_t + D_t) = \sigma_\varepsilon^2 \left[\frac{(\phi+1)\phi^\tau - \phi - \phi^{1+2\tau}}{(\phi-1)^2\,(\phi+1)} \right]. \tag{26}$$

Knowing the inventory variance for AR(1) demand (24), we have sufficient information to compute the optimal order quantities for each period in the planning cycle.

4.1. Determining the production quantities

To understand how this policy works in practice, consider the following numerical example using this type of demand process.

Example 1. Consider the staggered system in figure 2, where $L = 4$ and $P = 7$. The current period is $t = 0$, and we are ordering for the receipts in periods 5 through 11. Demand is known to be first-order autocorrelated (i.e. $\theta_n = \phi^n$) with $\phi = 0.70$, with a mean $\mu = 10$, and error terms that are normally distributed $\varepsilon \sim \mathcal{N}(0, 1)$. We have also observed that the inventory level is $I_0 = 5.20$, and that the work-in-progress inventory is $W_{0,5} = 41.30$. Adding these we obtain the current inventory position $I_0 + W_{0,5} = 46.50$. The optimal order quantities are calculated as follows:

1. Make the *lead-time demand forecast* for the first period. As demand is autocorrelated, our last demand observation, $D_0 = 8.71$, is sufficient to make a forecast of lead time demand $\mu + (D_0 - \mu)\sum_{n=1}^{L+1}\phi^n = 10 + (8.71 - 10) \times 1.94117 = 47.5$.

2. Make the *single period forecasts* for the remaining periods, $k = 2$ to $k = 7$, or equivalently $\tau = 6$ to $\tau = 11$. We obtain this as $F_{t,\tau} = \mu + (d_t - \mu)\phi^\tau$, which for the second order of the cycle gives $F_{t,6} = 10 + (8.71 - 10) \times (0.7)^6 = 9.85$. The remaining periods are obtained in the same way, after incrementing τ.

3. Calculate the *time-varying safety stocks*. These are of the form $I^*_{t+\tau} = \sigma_{i,k}\Phi^{-1}\left[B/\left(B + H\right)\right]$, where $\sigma_{i,k}$ is the square root of the inventory variance found in (24). Thus, the first safety stock for $k = 1$ is $I^*_5 = \sqrt{22.7923} \times \Phi^{-1}\left(0.9\right) = 6.12$. For the following periods, increment k and perform the calculation again. For example $I^*_6 = \sqrt{31.4428} \times \Phi^{-1}\left(0.9\right) = 7.19$.

4. Determine the *safety stock increase* between periods. Starting with $k = 2$, this is done by the subtraction $I^*_{t+\tau} - I^*_{t+\tau-1}$. The first change in safety stock, occurring at $\tau = 6$ is $I^*_6 - I^*_5 = 7.19 - 6.12 = 1.07$. The remaining safety stock changes are obtained by incrementing τ.

5. Calculate the *first receipt* according to the standard OUT policy. We order the lead-time forecast of demand, plus the safety stock, minus the inventory position, according to $\hat{D}_{0,5} + I_5^* - (I_0 + W_{0,5}) = 47.5 + 6.12 - 46.5 = 7.12$.

6. Calculate the *remaining receipts* using a simpler formula. The second receipt of the cycle, with $\tau = 6$, takes the single-period forecast, plus the increase in safety stock, $R_6 = \hat{D}_{0,6} + I_6^* - I_5^* = 9.85 + 1.07 = 10.92$. The remaining receipts of the cycle are calculated in the same way, with τ incremented.

Table 1 presents the optimal order quantities for the entire cycle, as well as the intermediate results. Contrary to the worked example, the table has been calculated with machine precision, so the last decimal of the calculations may vary.

4.2. Impact of the heteroskedastic inventory levels on costs and service

Now consider the case where a plan made at time t will generate its first receipt in time for it to affect I_{t+1} (that is, $L = 0$). The previous setting $L = 4$ has been substituted for $L = 0$ to highlight the effects of inventory heteroskedasticity.

Figure 3 illustrates the inventory standard deviation for a range of AR(1) demands. The configurations where $|\phi| < 1$ reflect stationary demand, while other configurations reflect nonstationary demand. In either case, the inventory level is stationary. As we can see from (15) the inventory standard deviation is increasing in τ. The consequences of this are clear: staggering increases the total inventory cost, particularly when there is significant autocorrelation. Staggering is least harmful when demand is negatively autocorrelated and stationary $(-1 < \phi < 0)$.

Corollary 7. Some special cases of the first-order autoregressive inventory variance can be identified.

(a) When $\phi \to 0$ the inventory variance is a linear function of τ,

$$\mathrm{Var}(I_t) = \sigma_\varepsilon^2 \tau. \tag{27}$$

(b) When $\phi \to 1$ demand is a random walk in discrete time (Box and Jenkins, 1976, p. 123), and the inventory variance is a cubic function, increasing in τ,

$$\mathrm{Var}(I_t) = \sigma_\varepsilon^2 \frac{\tau}{6}(1+\tau)(1+2\tau). \tag{28}$$

Table 1: Calculating the planned receipts for the numerical example.

Period (t)	0	1	2	3	4	5	6	7	8	9	10	11
k	3	4	5	6	7	1	2	3	4	5	6	7
τ	7	8	9	10	11	5	6	7	8	9	10	11
Inventory level (I)	5.2	-	-	-	-	-	-	-	-	-	-	-
Work-in-progress (W)	41.3	-	-	-	-	-	-	-	-	-	-	-
Demand (D)	8.71	-	-	-	-	-	-	-	-	-	-	-
Lead-time forecast[a] ($\hat{F}$)	-	-	-	-	-	47.50	-	-	-	-	-	-
Single-period forecast ($\hat{D}$)	-	-	-	-	-	-	9.85	9.89	9.93	9.95	9.96	9.97
Safety stock required (I^*)	-	-	-	-	-	6.12	7.19	8.19	9.12	10.00	10.83	11.61
Change in safety stock	-	-	-	-	-	-	1.07	1.00	0.94	0.88	0.83	0.78
Planned receipts (R)	-	-	-	-	-	7.12	10.92	10.89	10.86	10.83	10.79	10.76

Dashes (-) refer to values not needed for calculating the order quantities in this cycle.

[a] Forecasted demand over the lead time.

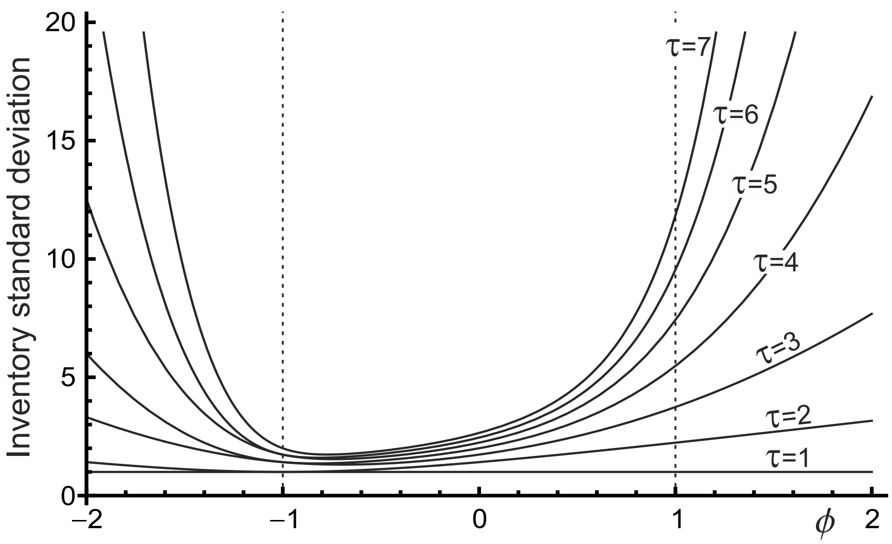

Figure 3: The inventory standard deviation for AR(1) demand is increasing in τ.

This expression is the variance of the error term, multiplied with a square pyramidal number.

(c) When $\phi \to (-1)$ the inventory variance is an increasing affine function for odd or even values of τ,

$$\mathrm{Var}(I_t) = \sigma_\varepsilon^2 \left[\frac{1-(-1)^\tau}{4} + \frac{\tau}{2}\right]. \tag{29}$$

From this equation, we see that the inventory variance only increases strictly for odd values of τ when $\phi = (-1)$. Consequently, when the lead time, L, is even, $\mathrm{Var}(I_{t+L+1}) = \mathrm{Var}(I_{t+L+2})$, $\mathrm{Var}(I_{t+L+3}) = \mathrm{Var}(I_{t+L+4})$ and so forth. For odd L the pattern starts with $\mathrm{Var}(I_{t+L+2}) = \mathrm{Var}(I_{t+L+3})$.

Equation (27) reveals that when demand lacks autocorrelation, the variance increases linearly, meaning that the inventory standard deviation is proportional to the square root of τ. This is a fundamental result, demonstrating that the staggered policy behaves like an iterated OUT policy with increasing lead times for each order in the cycle. The inventory variance increases step-wise in (29). This is noticeable in the inventory standard deviation of figure 3, where for $\phi = (-1)$ every curve coincides with another curve (except for $\tau = 7$, which would coincide with $\tau = 8$ if $P \geq 8$).

Consider the costs $B = 1$ and $H = 9$. They imply that the inventory costs are minimized when the availability, for every τ, is 90%. As we know, the inventory variance changes with τ, and hence time-varying safety stocks are optimal. If we insist on using constant safety stocks, the availability will change over the cycle. For example, a constant safety stock can be based on the worst-case inventory variance, obtained at the end of the cycle, providing $I^*_{t+k+L} = \sigma_{i,P}\Phi^{-1}(0.9)$. This is not cost-optimal, but it simplifies the order quantity calculations. The results of this alternative strategy can be seen in figure 4, where the availability, α, is given by

$$\alpha = \Phi\left(\frac{I^*_t}{\sigma_{i,k}}\right). \tag{30}$$

For any constant safety stock setting, availability degrades as τ increases. This is due to $\sigma_{i,k}$ being increasing in k. For the safety stock setting under discussion, $\sigma_{i,P} \geq \sigma_{i,k}$, making the availability lower-bounded at 90%.

A more sophisticated constant safety stock setting could be based on the average inventory variance,

$$I^*_t = \Phi^{-1}\left(\frac{B}{B+H}\right)\sqrt{P^{-1}\sum\nolimits_{n=1}^{P}\sigma^2_{i,n}}. \tag{31}$$

This safety stock setting is obtained if one ignores the cyclical heteroskedasticity and takes the variance of the inventory process as a whole. The resulting availability is shown in figure 5, illustrating that the target availability of 90% is no longer a lower bound. On average, however, the availability is above the target of 90%. This always results when $H < B$ and when $I^*_t \geq \Phi^{-1}\left\{\bar{\sigma}_{i,P}\Phi\left[B/(B+H)\right]\right\}$ (which for this strategy is true, due to Jensen's inequality) as a consequence of lemma 8.

Lemma 8. For fixed safety stocks $I^*_t \geq 0$,

$$\Phi\left(\frac{I^*_t}{\bar{\sigma}_{i,P}}\right) \leq \frac{1}{P}\sum_{k=1}^{P}\Phi\left(\frac{I^*_t}{\sigma_{i,k}}\right); \tag{32}$$

for $I^*_t \leq 0$ the inequality is reversed.

Proof. Observe that $\Phi(x)$ is concave for $x \geq 0$. Then (32) is an immediate result of Jensen's inequality. On the domain $x \leq 0$, $\Phi(x)$ is convex, and the inequality in (32) is reversed. This completes the proof. ■

The takeaway from this lemma is that the availability estimate $\hat{\alpha} = \Phi\left(I_t^*/\bar{\sigma}_{i,P}\right)$ is less than the realized availability, i.e. $\hat{\alpha} \leq \alpha$, when $\hat{\alpha} \geq 0.5$, or equivalently when $H \leq B$. Conversely, when $\hat{\alpha} < 0.5$, $\hat{\alpha}$ overestimates α.

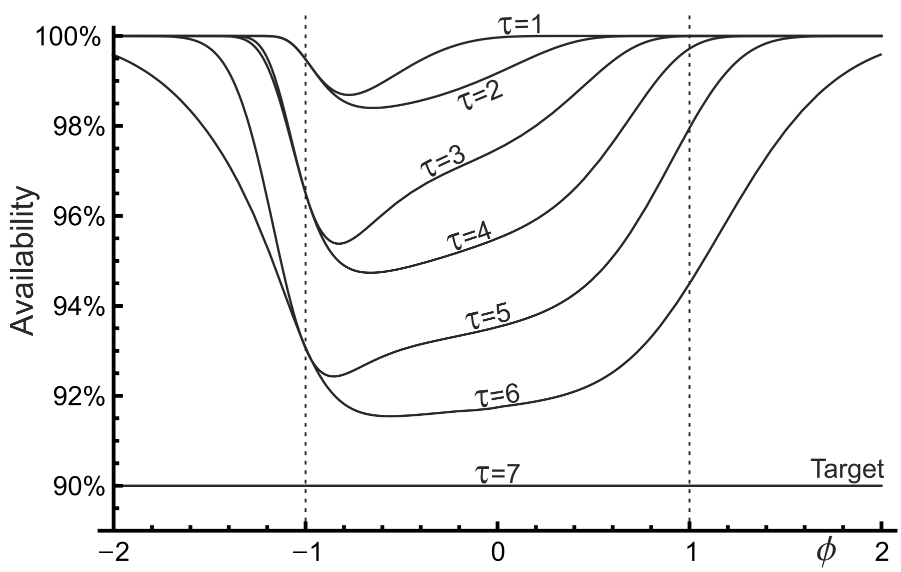

Figure 4: Availability for a fixed safety stock based on the inventory variance at the end of the order cycle.

The cost differential between these three strategies is worth considering. Figure 6 verifies that the optimal time-varying safety stock outperforms all constant settings. The worst economic performance results from the constant safety stock setting based on the end-of-cycle inventory variance. This is clear for nonstationary demand, but when there is little autocorrelation, the two fixed safety stock strategies are nearly cost-equal.

To see if these observations hold under different cost settings, we may consult figure 7, where H and B assume different values, but in all cases $H + B = 10$. These settings imply an optimal availability of 80%, 90%, 95%, or 99%. Regardless of the cost configuration, we notice that autocorrelation drives the cost differential between the constant and the time-varying safety safety stock settings. This effect appears for all of the cost settings, particularly

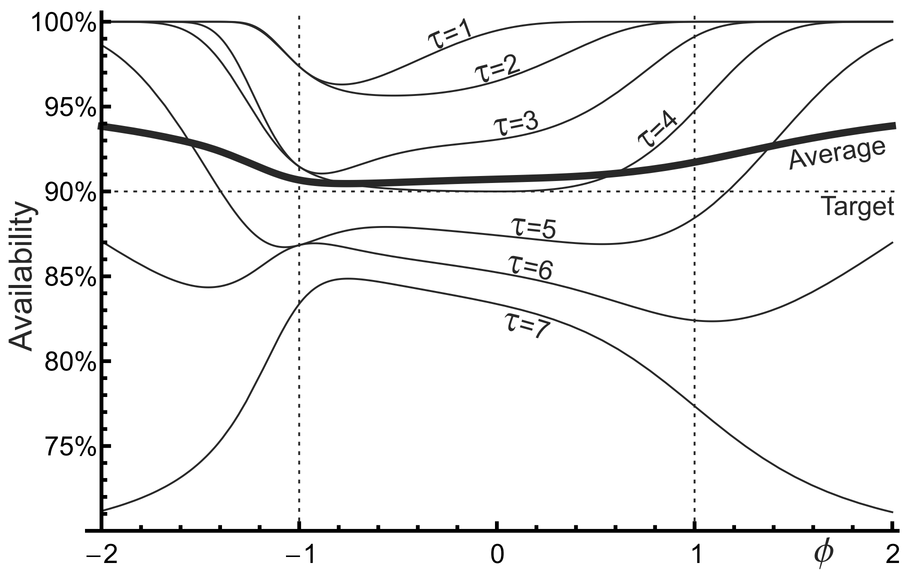

Figure 5: Availability for a fixed safety stock based on the average inventory variance.

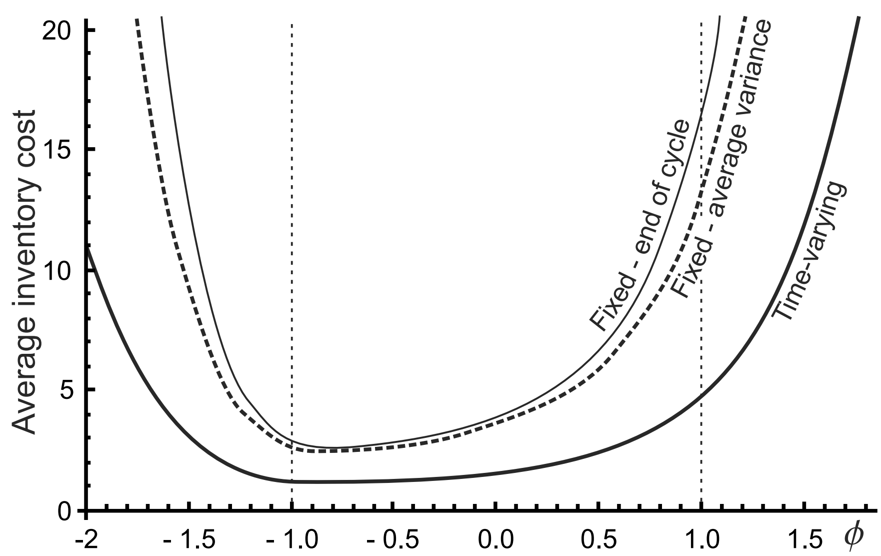

Figure 6: Average inventory cost over a seven-day cycle for three different safety stock settings with $H = 1, B = 9$.

when demand is nonstationary. For stationary demand, the cost varies less between the three settings when H and B are close, as with the $H = 2, B = 8$ case, and it varies more when the difference between H and B is large. In the $H = 0.1, B = 9.9$ setting, corresponding to an optimal availability of 99%, the superior performance of the time-varying safety stock becomes clear, leading to a fundamental insight: time-varying safety stocks are most important when demand exhibits strong autocorrelation, and when high service levels are required.

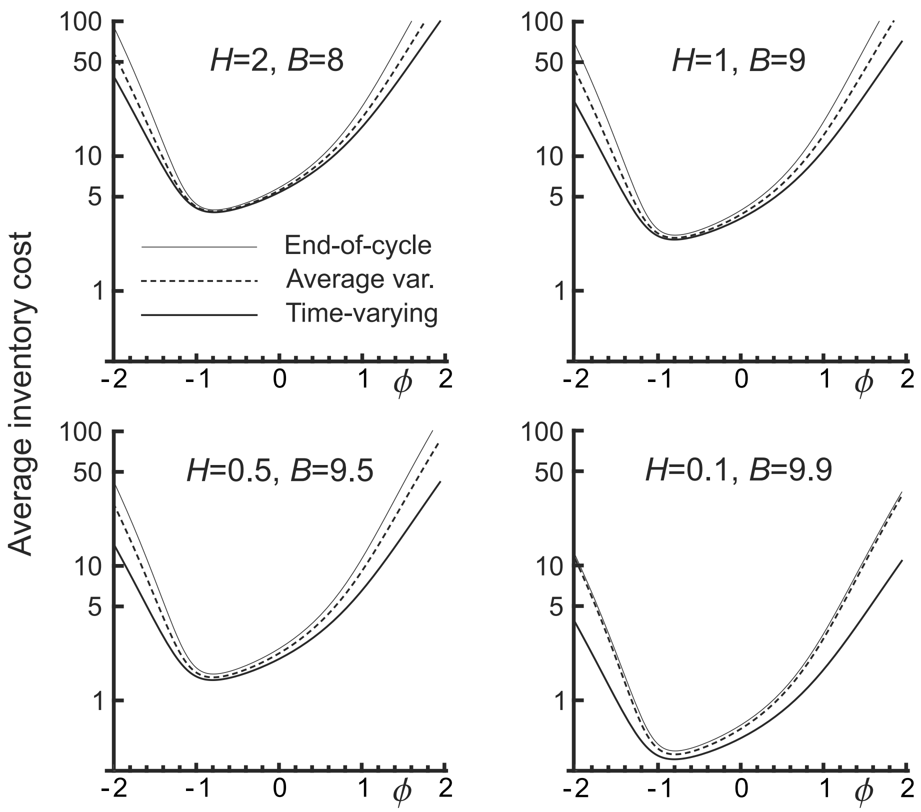

Figure 7: Average inventory cost over a seven-day cycle for various settings of H and B.

As the variable safety stock is the strategy of choice – providing the required availability at the lowest cost – we may wish to understand how the fill rate develops over a cycle, when time-varying safety stocks are in place. Though the availability remains constant, we see in figure 8 (restricted to $|\phi| < 1$ as the fill rate can only be defined for stationary demand) that the fill rate fluctuates over τ, and that it depends on ϕ. Furthermore, figure 8 indicates that the fill rate degrades as τ increases, particularly when demand is positively autocorrelated. Although the fill rate is undefined for $|\phi| \geq 1$, the numerical experiments consistently show that the fill rate approaches 100% as $|\phi| \to 1$.

From the definition of the fill rate (10) and the knowledge that the inventory variance is increasing in τ, we can make some observations about the fill rate under constant safety stock settings. For the end-of-cycle constant safety stock setting, the fill rate at the end of the cycle, with $\tau = 7$, is identical to the fill rate of the optimal time-varying safety stock. For $\tau < 7$, the fill rate is higher. The other constant safety stock setting, based on the average inventory variance, does not have the fill rate of the optimal time-varying safety stock at $\tau = 7$ as a lower bound.

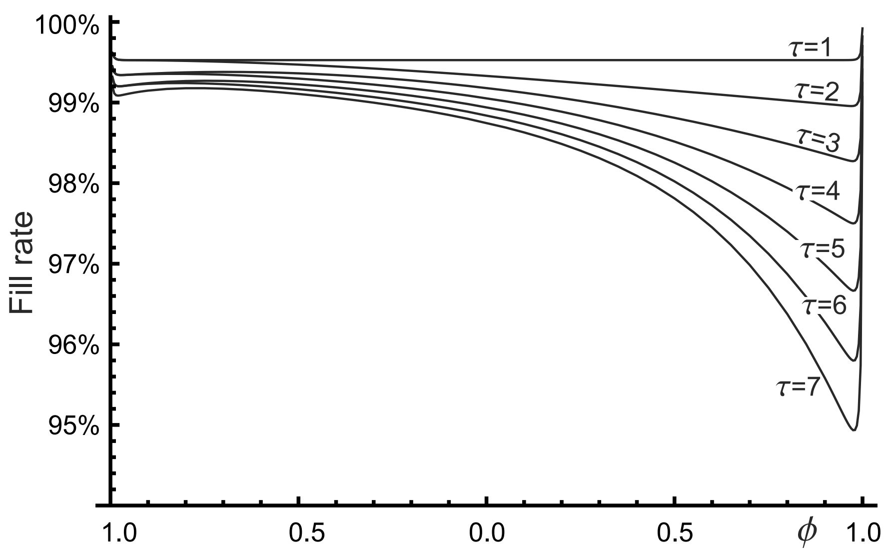

Figure 8: Fill rates are affected by staggering and by autocorrelation.

4.3. Determining the optimal length of the planning cycle

Figure 9 shows P^* equation for AR(1) demand under six different lead times, using (22) and (24). Each area between the contour lines indicates that a particular P^* is optimal; P^* is increasing in λ (every time we cross a contour in figure 9 from below, P^* increases by one). For the setting $\phi = (-1)$, $L + P^*$ is always even, as a consequence of (22) in conjuction with the odd-even effect in (29). Therefore, with an even lead time, P^* is also even, and vice versa.

The area a is figure 9 denotes the case when $P^* \geq 20$ but we have not drawn the contours as they become indistinguishable from each other. The following numerical example describes the optimization procedure.

Example 2. To determine the optimal planning cycle length, start by identifying the auditing, inventory holding, and backlog costs. Then use (20) to determine ψ for this set of costs, and more importantly $\lambda = V/\psi$. Finally, exploit theorem 5 to find P^*, either by inspecting figure 9, or by finding two reorder cycle lengths P_1 and P_2, such that $\lambda_{P_1} \leq \lambda \leq \lambda_{P_2}$, and then performing a binary search for λ between P_1 and P_2, until a P is found such that $\lambda_{P-1} \leq \lambda \leq \lambda_P$. Then $P^* = P$ denotes the optimum.

Suppose $V = 10, H = 1, B = 9$. This leads to $\lambda = 0.695$. With zero lead time and i.i.d demand ($\phi = 0$), the open circle in figure 9 shows that $P^* = 4$. Were demand instead positively correlated with $\phi = 0.9$, then $P^* = 2$, illustrated by the closed circle in figure 9. Were $L = 4$, then $P^* = 5$ with $\phi = 0$, and $P^* = 2$ with $\phi = 0.9$. This illustrates that positive autocorrelation favours short planning cycles, and also that the physical production lead time influences P^*.

5. Conclusion

5.1. Theoretical contribution

We have identified the inventory-optimal policy under staggered deliveries and autocorrelated demand. The strategy is to correct all inventory errors for the first order of the cycle, and then to order only the forecasted demand for the period in question and the required change in the safety stock, according to (6) and (7). This makes real the optimal policy identified by the Flynn and Chiang papers by applying the OUT policy to autocorrelated demand.

Flynn and Garstka (1990) and Chiang (2009) present staggered models with fixed audit costs, assuming i.i.d. demand. Our policy differs from Flynn and Garstka (1990), as $O_{t,k}$ may be negative. Conservative settings of μ, and σ_ε render the effect of negative orders negligible, making our policy computationally consistent with Flynn and Garstka (1990). Our policy also differs from the two policies in Chiang (2009), as we have a time-varying production quantity in each period. Our policy ensures that the target availability is maintained consistently. However, Chiang's policies may be practical for small P. Indeed, when $P \leq 2$, Chiang's simplified policy is

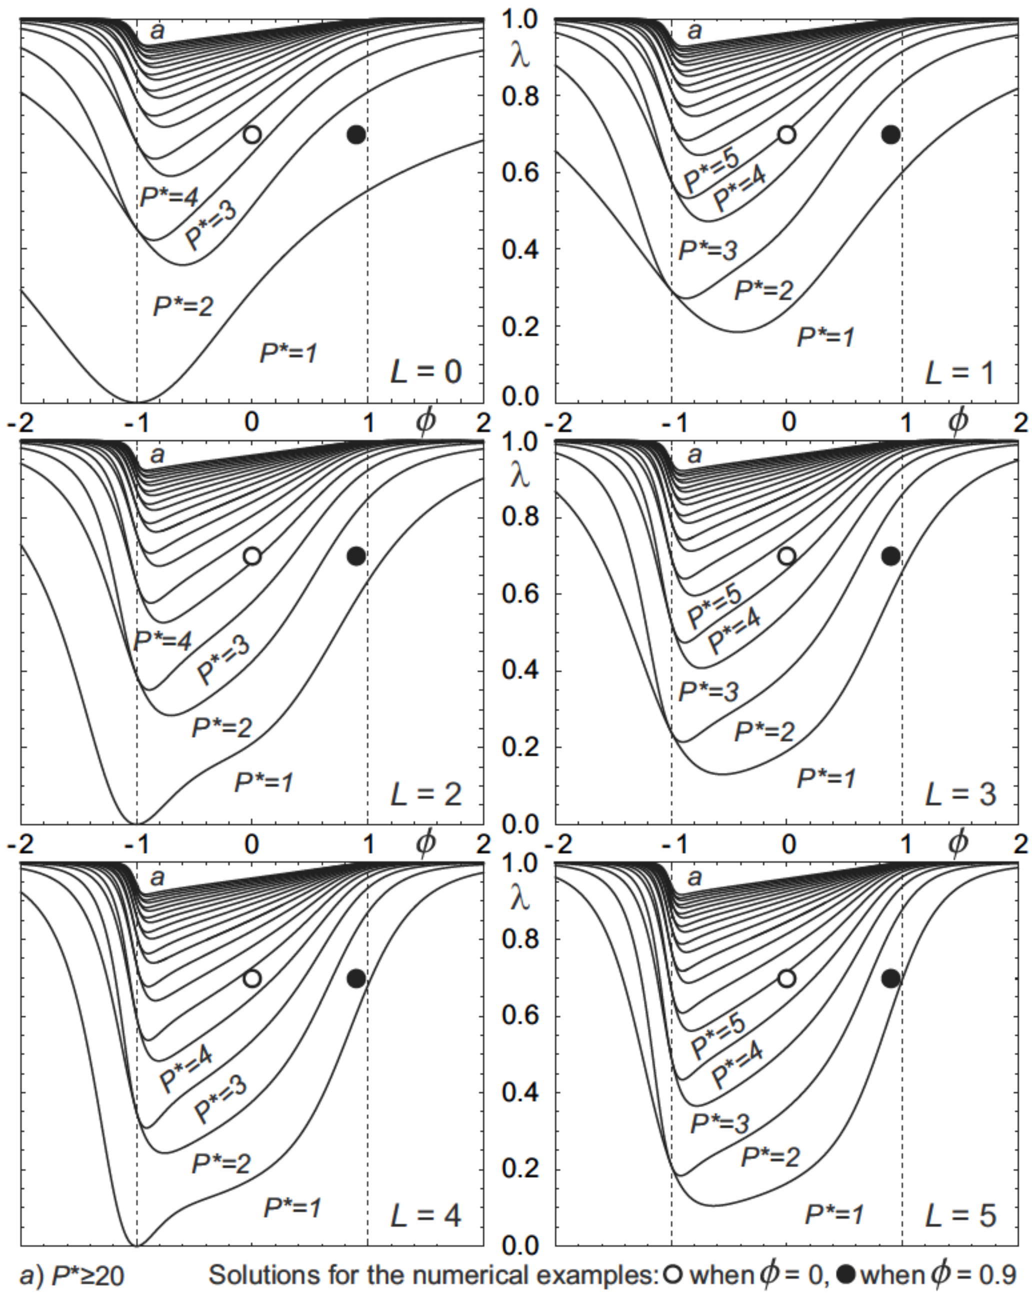

Figure 9: Optimal order cycle lengths P^* for some values of ϕ, λ, and L when $P^* < 20$.

optimal. The regular policy in Chiang (2009) has an added non-negativity constraint, which in Chiang's numerical examples caused an increased cost over the simplified policy.

Our model allows for constant safety stocks if desired, but we find that not only are time-varying safety stocks more economical, they also ensure that the target availability is achieved consistently. The overall safety stock is affected by the autocorrelation of demand and increases with the order cycle length. In the special case of an AR(1) demand process with $\phi = (-1)$, the safety stock only needs to be changed for periods when τ is odd. Our model and analysis also captures the nonstaggered case when $P = 1$, which results in a regular OUT policy.

The inventory variance is increasing over the order cycle for any demand autocorrelation function, and the heteroskedasticity affects fill rates, even when the availability is kept constant. This causes the fill rate to fluctuate cyclically. We have also provided an exact approach for determining P^*, the optimal length of the planning cycle, when auditing, holding, and backlog costs are present. The optimization procedure reveals that P^* is an increasing function of the audit cost V, and a decreasing function of B, H, L, and σ_ε.

5.2. Managerial insights

If a production system requires consistent availability, it is necessary to take into account the time-varying inventory variance. Ignoring the heteroskedasticity of inventory will result in either excessive service levels and unnecessary costs, or poor service on predictable days of the planning cycle. However, if time-varying safety stocks are deemed impractical or too complicated, we recommend a safety stock setting based on the average inventory variance over the cycle. Then the availability will fluctuate over the cycle, but on average it will exceed the critical fractile $B/(B + H)$ when $B \geq H$.

Even with optimal time-varying safety stocks, fill rates may degrade over the cycle, particularly when demand is positively autocorrelated. Reducing the length of the planning cycle provides an opportunity to reduce inventory costs, and is especially attractive when the initial planning cycle is long, and demand exhibits strong autocorrelation. However, short planning cycles require frequent audits, incurring a cost. The balance between inventory and audit costs must be regulated carefully via the order cycle length.

5.3. Limitations and further work

This investigation focused only on inventory performance and the optimal order cycle length under autocorrelated demand. Our numerical investigation of the fill rate reveals a rich behaviour that calls for further study. One might also consider the impact of staggered deliveries on bullwhip or capacity costs.

Furthermore, we have assumed perfect knowledge of the autocorrelation function (ACF) of demand, that the ACF does not change over time, and that we can observe past demand from the beginning of time. In a real setting we must estimate the ACF from a limited set of past observations. This may introduce specification errors, and robustness tests could be considered, perhaps along the lines set out in Hosoda and Disney (2009). Further consideration could also be given to the mis-specification of the demand distribution as in Akcay et al. (2011) and Lee (2014).

6. Acknowledgements

We thank Dr Xun Wang of Cardiff Business School and the anonymous referees, whose advice has led to substantial improvements in this paper.

7. References

Akcay, A., Biller, B., and Tayur, S., 2011. Improved inventory targets in the presence of limited historical demand data. *Manufacturing & Service Operations Management*, 13(3):297–309.

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

Bertsekas, D. P., 2005. *Dynamic Programming and Optimal Control*, vol. 1. Athena Scientific, Belmont, MA.

Bicheno, J. and Holweg, M., 2009. *The Lean Toolbox: The Essential Guide to Lean Transformation*. Picsie Books, Buckingham.

Box, G. E. P. and Jenkins, G. M., 1976. *Time Series Analysis: Forecasting and Control*. Holden-Day, San Francisco.

Burbidge, J., 1983. Five golden rules to avoid bankruptcy. *Production Engineer*, 62(10):13–14.

Cachon, G. and Terwiesch, C., 2009. *Matching Supply with Demand*. McGraw-Hill, Singapore.

Chiang, C., 2001. Order splitting under periodic review inventory systems. *International Journal of Production Economics*, 70(1):67–76.

Chiang, C., 2007. Optimal ordering policies for periodic-review systems with a refined intra-cycle time scale. *European Journal of Operational Research*, 177(2):872–881.

Chiang, C., 2009. A periodic review replenishment model with a refined delivery scenario. *International Journal of Production Economics*, 118(1):253–259.

Disney, S. M., Gaalman, G. J. C., Hedenstierna, C. P. T., and Hosoda, T., 2015. Fill rate in a periodic review order-up-to policy under auto-correlated normally distributed, possibly negative, demand. *International Journal of Production Economics*. In press, DOI: 10.1016/j.ijpe.2015.07.019.

Erkip, N., Hausman, W. H., and Nahmias, S., 1990. Optimal Centralized Ordering Policies in Multi-Echelon Inventory Systems with Correlated Demands. *Management Science*, 36(3):381–392.

Flynn, J., 2000. Selecting T for a periodic review inventory model with staggered deliveries. *Naval Research Logistics (NRL)*, 47(4):329–352.

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

Flynn, J., 2008. An effective heuristic for the review period in an inventory model with staggered deliveries and normal demands. *European Journal of Operational Research*, 186(2):671–680.

Flynn, J. and Garstka, S., 1990. A dynamic inventory model with periodic auditing. *Operations Research*, 38(6):1089–1103.

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

Heyman, D. P. and Sobel, M. J., 1984. *Stochastic Models in Operations Research, Vol. II: Stochastic Optimization*. McGraw-Hill, New York.

Hopp, W. J. and Spearman, M. L., 2008. *Factory Physics*. McGraw-Hill, New York.

Hosoda, T. and Disney, S. M., 2009. Impact of market demand mis-specification on a two-level supply chain. *International Journal of Production Economics*, 121(2):739–751.

Johnson, M. E., Lee, H. L., Davis, T., and Hall, R., 1995. Expressions for item fill rates in periodic inventory systems. *Naval Research Logistics*, 42(1):57–80.

Jonsson, P. and Mattsson, S.-A., 2013. *Lagerstyrning i svensk industri: 2013 års användning, användningssätt och trender*. Technical report, Chalmers University of Technology, Gothenburg. In Swedish.

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):627.

Lee, Y. S., 2014. A semi-parametric approach for estimating critical fractiles under autocorrelated demand. *European Journal of Operational Research*, 234(1):163–173.

Lian, Z., Deshmukh, A., and Wang, J., 2006. The optimal frozen period in a dynamic production model. *International Journal of Production Economics*, 103(2):648–655.

Pinedo, M. L., 2009. *Planning and Scheduling in Manufacturing and Services*. Springer, New York.

Prak, D., Teunter, R., and Riezebos, J., 2015. Periodic review and continuous ordering. *European Journal of Operational Research*, 242(3):820–827.

Silver, E. A., Pyke, D. F., and Peterson, R., 1998. *Inventory Management and Production Planning and Scheduling*. Wiley, New York.

Silver, E. A. and Robb, D. J., 2008. Some insights regarding the optimal reorder period in periodic review inventory systems. *International Journal of Production Economics*, 112(1):354–366.

Zhang, X., 2007. Inventory control under temporal demand heteroscedasticity. *European Journal of Operational Research*, 182(1):127–144.

Appendix A. Proof for theorem 2

Proof.

(a) Assume that we are to place an order at time t, to be received in period $t+\tau$. By setting the conditional expectation $E\left[I_{t+\tau}|\{D_t, D_{t-1}, \cdots\}\right] = I^*_{t+\tau}$, and solving for $R_{t+\tau}$, we obtain the policy that minimizes the inventory cost for any τ. Finally we set $\tau = L+1$, to obtain the solution for the first order in a cycle.

To begin, we use induction on the inventory balance equation (2) to obtain $I_{t+1} = I_{t-1} + R_t + R_{t+1} - D_t - D_{t+1}$. Extending this to $I_{t+\tau}$ yields

$$I_{t+\tau} = I_t + \sum_{n=1}^{\tau} (R_{t+\tau} - D_{t+\tau}) = I_t + W_{t,\tau} + R_{t+\tau} - F_{t,\tau}, \tag{A.1}$$

The inventory costs are convex, and we seek to minimize them for an arbitrary period by setting the expected inventory level to $I^*_{t+\tau}$. Therefore, let $I^*_{t+\tau}$ equal the expectation of (A.1), conditional on our observations of demand up to time t, when the order is determined. We obtain

$$E\left[I_{t+\tau}|\{D_t, D_{t-1}, \cdots\}\right] = I^*_{t+\tau} = I_t + R_{t+\tau} + W_{t,\tau} - \hat{F}_{t,\tau}. \tag{A.2}$$

The receipt rate $R_{t+\tau}$ can be found by rearranging (A.2),

$$R_{t+\tau} = \hat{F}_{t,\tau} + I^*_{t+\tau} - I_t - W_{t,\tau}. \tag{A.3}$$

Finally, we let $k=1$, so that $\tau = L+1$ and obtain

$$O_{t,1} = R_{t+L+1} = \hat{F}_{t,L+1} + I^*_{t+L+1} - I_t - W_{t,L+1}, \tag{A.4}$$

This concludes the first part of the proof.

(b) Assume that $L+1 < \tau \leq L+P$, so that $\tau+1$ corresponds to $k>1$. Inserting the receipts (A.3) back into the inventory equation (A.1) gives

$$I_{t+\tau} = I^*_{t+\tau} + \hat{F}_{t,\tau} - F_{t,\tau}. \tag{A.5}$$

Rearranging the inventory balance equation (2) yields $R_{t+\tau} = I_{t+\tau} - I_{t+\tau-1} + D_{t+\tau}$. Replacing both inventory levels with their equivalent form in (A.5), we obtain

$$R_{t+\tau} = O_{t,k} = I^*_{t+\tau} - I^*_{t+\tau-1} + \hat{D}_{t,\tau}, \tag{A.6}$$

where $\hat{D}_{t,\tau+1} = \hat{F}_{t,\tau+1} - \hat{F}_{t,\tau}$ is the single-period forecast for $D_{t+\tau+1}$ made at time t. This completes the proof. ■

Appendix B. Proof for theorem 4

Proof. We shall express I_t and $I_t + D_t$ as a weighted sum of independent error terms, and then take the variance or covariance.

(a) Recall the inventory equation (A.5)

$$I_t = I_t^* + \hat{F}_{t-\tau,\tau} - \sum_{n=0}^{\tau-1} D_{t-n}. \tag{B.1}$$

Let us express the lead-time demand as a weighted sum of error terms,

$$\sum_{a=0}^{\tau-1} D_{t-a} = \left(\sum_{m=0}^{\tau-1} \varepsilon_{t-m} \sum_{n=0}^{m} \theta_n\right) + \left(\sum_{x=\tau}^{\infty} \varepsilon_{t-x} \sum_{y=0}^{\tau-1} \theta_{x-y}\right), \tag{B.2}$$

and express the corresponding forecast as a weighted sum of error terms,

$$\hat{F}_{t-\tau,\tau} = \sum_{x=\tau}^{\infty} \varepsilon_{t-x} \sum_{y=0}^{\tau-1} \theta_{x-y}. \tag{B.3}$$

Substituting the lead-time demand (B.2) and the forecast (B.3) into the inventory equation, we obtain

$$I_t = I_t^* - \sum_{m=0}^{\tau-1} \varepsilon_{t-m} \sum_{n=0}^{m} \theta_n. \tag{B.4}$$

Clearly, $E[I_t] = I_t^*$. Taking the variance, we obtain

$$\mathrm{Var}(I_t) = \sigma_\varepsilon^2 \sum_{m=0}^{\tau-1} \left(\sum_{n=0}^{m} \theta_n\right)^2, \tag{B.5}$$

completing the first part of the proof.

(b) Without loss of generality, assume $I_t^* = 0$. We can then characterize $I_t + D_t$ as

$$\begin{aligned} I_t + D_t &= \left(\sum_{x=0}^{\infty} \varepsilon_{t-x}\, \theta_x\right) - \left(\sum_{m=0}^{\tau-1} \varepsilon_{t-m} \sum_{n=0}^{m} \theta_n\right) \\ &= \left(\sum_{x=0}^{\infty} \varepsilon_{t-x}\, \theta_x\right) - \left[\varepsilon_t\, \theta_0 + \sum_{m=1}^{\tau-1} \varepsilon_{t-m} \left(\theta_m + \sum_{n=0}^{m-1} \theta_n\right)\right] \\ &= \left(\sum_{x=\tau}^{\infty} \varepsilon_{t-x}\, \theta_x\right) - \left\{\sum_{m=1}^{\tau-1} \varepsilon_{t-m} \left[\left(\sum_{n=0}^{m} \theta_n\right) - \theta_m\right]\right\}. \end{aligned} \tag{B.6}$$

With $I_t + D_t$ of this form, we take the variance

$$\mathrm{Var}(I_t + D_t) = \sigma_\varepsilon^2 \left\{ \left[\sum_{m=1}^{\tau-1} \left(\sum_{n=0}^{m-1} \theta_n \right)^2 \right] + \sum_{x=\tau}^{\infty} \theta_x^2 \right\}, \tag{B.7}$$

completing this part of the proof.

(c) To obtain $\mathrm{Cov}(D_t,\, I_t + D_t)$, we exploit that (B.4) and (B.6) are already of the required form. Taking the covariance gives (17), completing the proof. ■