Determining the Reorder Point and Order-Up-To-Level in a Periodic Review System So As to Achieve a Desired Fill Rate and a Desired Average Time Between Replenishments

Edward A. Silver
Haskayne School of Business
University of Calgary
2500 University Drive NW
Calgary, AB T2N 1N4
Canada

Hussein Naseraldin
Joseph L. Rotman School of Management
105 St. George Street
Toronto, ON M5S 3E6
Canada

Diane P. Bischak
Haskayne School of Business
University of Calgary
2500 University Drive NW
Calgary, AB T2N 1N4
Canada

October 2007
Revised April 2008

Abstract

In this paper we consider a periodic review, reorder point, order-up-to-level system, a type commonly used in practice. Motivated by a specific practical context, we present a novel approach to determining the reorder point and order-up-to-level (for a given review interval) so as to target desired values of i) customer fill rate and ii) average time between consecutive replenishments. Specifically, by using a diffusion model (producing normally distributed demand) we convert a periodic review, constant lead time setting into one having continuous review and a random lead time. The method is simple to implement and produces quite reasonable results.

Keywords: Inventory control, Heuristics, Stochastic, Diffusion process, Supply chain management

Suggested running head: Determining *s* and *S* in a periodic system

1. Introduction

This paper is concerned with an inventory control system commonly used in practice. Specifically, the status of an item is examined at equi-spaced (review) intervals and, if the inventory position (on-hand plus on-order minus backorders) is at or below the reorder point (denoted by s), then a replenishment, that raises it to the order-up-to-level (denoted by S), is initiated. The review interval (denoted by R) is often preset at a convenient value (e.g., day, week), which is what will be assumed here.

The research leading to this paper was motivated by the actual context faced by a major international producer and distributor of food products. In particular, they were using a periodic (weekly) review control system and it was deemed crucial to determine s and S so as to approximately satisfy two practical constraints: i) a specified fill rate (fraction of demand satisfied without backordering), i.e., a marketing requirement, and ii) a specified average time between consecutive replenishments (e.g., three weeks), desirable from the perspective of the supplier (production department).

It is inherently difficult to find proper values of the two control parameters, the reorder point and order-up-to-level, primarily due to periodic review causing undershoots of the reorder point before replenishments are triggered. This is illustrated in Figure 1, where replenishments are placed at times 0 and $3R$ and a shortage occurs because the undershoot at time $3R$ plus the demand during the lead time L exceeds the reorder point s. The probability distribution of the undershoot is a complicated function of the distance $S - s$ and the distribution of demand during the review interval R. This complexity is not present in the simpler periodic review, order-up-to-level system where a replenishment is initiated at each review instant (see, for example, Robb and Silver, 1998, or Silver et al., 1998). However, this latter type of system is not appropriate when there is a non-negligible fixed cost associated with each replenishment.

Practical implementation would be facilitated by a relatively simple procedure for determining appropriate values of s and S. An early software package, IBM's IMPACT system (IBM, 1971), provided a simple, but overly conservative, choice of s by assuming that the inventory position is just above s at the review prior to the one at which a replenishment is initiated, hence s must provide protection over an interval of length $R + L$. This approach had been advocated even earlier (Brown, 1967).

In contrast, most of the literature presents rather complicated procedures. Moreover, none of these explicitly deals with both of the above-mentioned constraints. Schneider (1978, 1981) and Tijms and Groenevelt (1984) used asymptotic results from renewal theory (Roberts, 1962) to approximate the undershoot distribution. For a different type of service measure (fraction of demand being on backorder), Schneider and Ringuest (1990) developed power approximations for s and S in the spirit of the original power approximation work of Erhardt (1979). Bashyam and Fu (1998) advocated a simulation-based approach to minimize setup and holding costs subject to meeting a prescribed fill rate. Other authors (e.g., Erhardt and Mosier, 1984, and Zheng and Federgruen, 1991) considered shortage costs rather than a service constraint. More recently, Moors and Strijbosch (2002) developed an efficient descriptive method for determining the fill rate for given values of s and S under the assumption of gamma distributed demand. Unlike in the approach we propose, their method would have to be combined with a search on s and S to achieve a desired value of the fill rate.

In the next section the underlying assumptions are laid out. Then, in Section 3 we provide a relatively non-technical overview of the general approach for determining appropriate values of s and S, subject to the practical constraints mentioned earlier. This is followed in Section 4 by a listing of the detailed steps of the method as well as a numerical illustration of its use. Simulation testing (in terms of how closely the two constraints are met) is presented in Section 5. Summary comments are provided in Section 6 and supporting technical details, including mathematical derivations, are placed in appendices.

2. Assumptions underlying the approach

The assumptions underlying the proposed procedure include the following:

i) The inventory position is reviewed every R units of time, where R is prespecified, not controllable. (For convenience we set R=1, i.e. the review interval is redefined as unit time.)

ii) There is a constant replenishment lead time (L) from when a replenishment is triggered (at a review instant) until it is available in stock.

iii) Demands in disjoint intervals of time are independent, stationary, and normally distributed.

iv) There is complete backordering of any demand during a stockout situation.

v) The service measure is the fill rate, the fraction of demand to be routinely met from stock.

vi) A target average time (n units of time) between consecutive replenishments is specified rather than explicitly incorporating setup and carrying costs.

3. Overview of the suggested approach

The key idea in the approach is to recognize that the reorder point is reached at a random time between reviews. As shown in Figure 1, at the instant that the inventory position drops to the reorder point there is a time, denoted by τ, remaining until the next review instant. (Under periodic review, the actual value of τ would **not** be observed in practice.) Thus, we can think in terms of a continuous review model with effective lead time $\tau + L$, where τ is a random variable. Under the assumption of stationary, independent, normally distributed demands in non-overlapping time intervals, the behaviour of the inventory position can be modeled by a diffusion process which, in turn, permits us to develop estimates of the first two moments of τ as a function of the distance $S - s$ and a measure of variability (the coefficient of variation, or CV) of the demand process. These moments are used as building blocks in choosing *S* and *s* so as to meet the two practical constraints in the following manner.

First, we recognize that the expected time between consecutive replenishments in the (*R*, *s*, *S*) system is the average size of a replenishment divided by the demand rate. But the average size of a replenishment is $S - s$ plus the average size of the undershoot (see Figure 1). The latter, in turn, is easily developed from the expected value of τ. Thus we can select $S - s$ so as to target the desired average time between replenishments, hence satisfying the second practical constraint.

Second, the moments of τ also lead to expressions for the mean and variance of total demand (denoted by *X*) over the effective lead time $\tau + L$. Assuming that this total demand is approximately normally distributed, we can then determine a value of *s* appropriate to satisfy the fill rate constraint. Finally, from the previous paragraph we know the desired value of $S - s$, hence we now can compute the value of *S*.

4. Suggested procedure for selecting appropriate values of *s* and *S*

In this section the steps of the implementation method are laid out. The underlying mathematical derivations are provided in Appendix A. Recall that the given review interval (*R*) is used as the basic time period.

Step 1: Determination of the mean and standard deviation of τ

For given values of the parameters *n* (the desired integer average number of time intervals between replenishments) and CV (the standard deviation divided by the mean of the demand per unit time), the functions given in Tables 1 and 2 provide easily computed values of the mean $\mathrm{E}(\tau)$ and the variance $\mathrm{Var}(\tau)$.

Step 2: Determination of the (normalized) mean and standard deviation of X

Setting μ to be the average demand during *R*, the basic time interval, compute

$$\frac{\mathrm{E}(X)}{\mu} = \mathrm{E}(\tau) + L \tag{1}$$

and

$$\frac{\sigma_X}{\mu} = \sqrt{[\mathrm{E}(\tau) + L](CV)^2 + \mathrm{Var}(\tau)} \tag{2}$$

Step 3: Finding the appropriate value of the safety factor k

Compute the quantity $(1-P_2)n\mu / \sigma_X$, where P_2 is the desired (fractional) value of the fill rate. Then set

$$G_u(k) = \frac{(1-P_2)n\mu}{\sigma_X} \tag{3}$$

where $G_u(k)$ is the unit normal loss function. There are tables, spreadsheet functions (e.g., in Excel), and other accurate, rational approximations that permit finding *k* when the value of $G_u(k)$ is known.

Step 4: Determination of the (normalized) appropriate value of the reorder point, s

Now s/μ is given by

$$\frac{s}{\mu} = \frac{\mathrm{E}(X)}{\mu} + k\frac{\sigma_X}{\mu} \tag{4}$$

Step 5: Determination of the (normalized) appropriate value of the order-up-to-level S

Use

$$\frac{S}{\mu} = \frac{s}{\mu} + n - \mathrm{E}(\tau) \tag{5}$$

From equations (1) to (5) and the fact that $\mathrm{E}(\tau)$ and $\mathrm{Var}(\tau)$ do not depend on the value of μ, it can be seen that all of the key results are normalized with respect to μ, the average demand per unit time, i.e. the value of μ does not affect the choice of s/μ or S/μ.

Numerical illustration. Consider the following realistic set of parameters values:

n = 4 (i.e., on average, a replenishment is desired every 4 periods)

CV = 0.3

L = 2 periods (i.e. the lead time is equal to 2 review intervals)

P_2 = 0.9 (desired fill rate)

Step 1: From the fractional polynomial approximations in Tables 1 and 2, $\mathrm{E}(\tau)$ = 0.5033 and $\mathrm{Var}(\tau)$ = 0.0839.

Step 2: Equations (1) and (2) give $\mathrm{E}(X)/\mu = 0.5033 + 2 = 2.5033$ and σ_X / μ = $\sqrt{2.5033(0.3)^2 + 0.0839} = 0.5561$.

Step 3: From equation (3) we have $G_u(k) = (0.1)4 / 0.5561 = 0.7194$. Using a table lookup or Excel function (see Appendices B and C of Silver et al., 1998), $k = -0.5309$.

Step 4: Use of equation (4) produces $s/\mu = 2.5033 + (-0.5309)(0.5561) = 2.208$.

Step 5: From equation (5) there results $S/\mu = 5.705$. So, for example, if the average demand per review interval, μ = 100, then s = 220.8 and S = 570.5.

5. Simulation testing of the approach

Our approach has involved approximations including the assumption of a normal distribution of X. Thus, it is important to simulate the actual performance associated with using the prescribed values of s and S. The simulation replicates what would happen in actual application of (R, s, S) control. Specifically, the inventory position is only updated at each review instant. For the update

$$\begin{array}{c}\text{inventory position at a review instant} \quad = \\ \text{inventory position after a possible order at previous review} \quad - \\ \text{total demand in the review interval} \;. \end{array} \tag{6}$$

The total demand is randomly generated from a normal distribution with parameters μ and σ. Note that in the simulation the times at which the inventory position hits s (hence the τ's) are **not** observed, just like in practice.

If the updated inventory position (denoted by I) is less than or equal to s, then an order (of size $Q = S - I$) is placed. Otherwise nothing is done until at least the next review. Any order placed arrives a constant time L later. Rather than keeping track of any backorders just before replenishments arrive, we have used a far more efficient method of estimating the actual fill rate achieved. The technical details of the method, as well as how the actual average time between replenishments is estimated, are presented in Appendix B.

5.1 Selection of parameter values

As seen from the detailed steps of Section 4, there are four independent parameters, namely CV, n, L and P_2. As shown in Table 3, each of these was set at three reasonable values (from a practical standpoint) to produce a total of 81 experiments.

As mentioned earlier, in the practical context that motivated this work, integer values of n were desired. The lowest value of interest is 2 (in that $n = 1$ would imply (R, S) control, which is much easier to analyze as the protection period is always exactly $R + L$). Because demand is normally distributed we used an upper value of 0.5 for the CV. For higher values of CV a different distribution (e.g., gamma) would likely be more appropriate. The L values were selected so that the smallest was less than the review interval and the largest appreciably larger than the interval. Finally, the three fill rates span the range typically targeted in practice.

5.2 Results

The results are displayed in Table 4. Although the details are not shown, all half-widths of 95% confidence intervals were less than 0.007 (0.12%) for the estimates of the average time between orders and less than 0.0004 (0.05%) for the estimates of the fill rates.

The major finding, indicated by the low percent deviations from the two targets (n and P_2), is that the approach works very well overall. Deviations from n as high as just over 5% occur in the average time (AT) between replenishments. However, such deviations are relatively unimportant in practice because, even if the average time between replenishments was right at the targeted value, such as 2 periods, there would still be occasions where the observed times would differ, e.g., be 1 or 3 periods, simply due to the random nature of the demand, i.e., one could never assure that the time between consecutive replenishments was a constant.

All of the entries in the last column of the table are negative, indicating that the achieved fill rate is always somewhat below the target. The largest deviation of 3.03% occurs in the third line, where $n = 2$ and $CV = 0.5$. (This also produces one of the highest deviations of AT from n.) This combination of n and CV tends to produce a number of instances of small τ values (in that the distribution of first passage time from S to s has a significant mass below, but close to, 2). This occurs for all cases where $CV = 0.5$, but to a lesser degree as n, L and P_2 increase.

There are two possible causes of the inaccuracy. The more obvious reason is the inaccuracy in assuming that X (the total demand in $\tau + L$) has a normal distribution. We were able to empirically determine the actual distribution of X in the following way. The time (u) for the diffusion process to first move the inventory level from S down to s has an analytical distribution (see equation (9) in Appendix A), thus u values were randomly generated from that distribution. But each value of u implies an associated value of τ, the remaining time until the next review (e.g., with R=1 a u value of 3.4 implies a τ of 0.6). Thus we were able to randomly generate a large number of values of τ. Then, for each such τ we randomly generated a value of X from a normal distribution with mean $\mu(\tau + L)$ and variance $\sigma^2(\tau + L)$. We thus developed empirical distributions of X for different combinations of the parameters, comparing each with a normal distribution having the same moments. As expected, the empirical distribution became closer to the normal as the fixed component L increased. This helps explain why the fill rate deviation decreases as L increases.

The second cause of the slight inaccuracy of our method is more subtle. If the reorder point (s) is first reached at a distance τ from the next review instant, we have used an associated effective lead time of length $\tau + L$. However, recalling the diffusion representation of the demand process, which permits negative demands (returns), it is conceivable that by the time of the next review the inventory position could have come back up above s due to negative overall demand in the interval τ. It can be shown that this is far more likely to happen for low values of τ. In such a case, the position will very quickly drop back to the reorder point, resulting in a relatively large value of τ (occurring in the next R interval). Thus, low values of τ assumed by our model will tend to be replaced by actual larger values. Consequently, our approach underestimates both $\mathrm{E}(\tau)$ and $\mathrm{Var}(\tau)$, resulting in a fill rate slightly lower than the target as well as an average order size (which has a component proportional to $\mathrm{E}(\tau)$) larger than targeted, hence leading to an average time between replenishments somewhat higher than *n*. As mentioned earlier, it is precisely when *CV* takes on its highest value of 0.5 that there is a significant probability of small τ values occurring.

The authors (Silver et al, 2008) have developed a modified approach for adjusting the estimates of $\mathrm{E}(\tau)$ and $\mathrm{Var}(\tau)$ under such circumstances. This modified approach is available upon request from the authors. However, as discussed earlier, it is our opinion that, from a practical standpoint, even without the adjustment (which would complicate matters from an implementation perspective) the deviations from the desired target values of fill rate and average time between replenishments are small enough not to be of concern.

6. Summary

In this paper we have considered an (*R*, *s*, *S*) inventory control system, a type commonly used in practice. Motivated by a practical context, we have presented a novel approach to determining *s* and *S* (for a given review interval *R*) so as to target desired values of i) customer fill rate and ii) average time between consecutive replenishments. Specifically, by using a diffusion model we converted a periodic review, constant lead time situation into one having continuous review and a random lead time. The basic method is simple to implement and produces quite reasonable results.

There are a number of possible extensions of this work, including:

i) Dealing with a variable lead time. Except for the possibility of crossing of orders (which cannot happen in the context of the current paper because τ is, by definition, less than R), this should be a straightforward extension.

ii) Using a demand distribution other than the normal, particularly to encompass situations where the CV exceeds 0.5. One possibility would be to modify the normal so that it cannot take on negative values as suggested by Strijbosch and Moors (2006). Another would be to use the gamma distribution. A third option would be to consider a discrete distribution, such as the Compound Poisson, so as to be able to model large, individual demand transactions.

iii) Incorporating shortage costing rather than a service measure.

iv) Treating a different control system, specifically an (R, s, mQ) system where orders must be placed in an integer multiple m of a prescribed, basic replenishment quantity. This would be appealing where the supply comes from production, e.g., the production department may prefer to produce in integer multiples of a convenient lot size (Q). An equivalent context would be where supply is only delivered as an integer number of non-unit-sized (Q) packs.

Acknowledgements

The authors thank Daniel Costa, Supply Chain Director, Nestlé Italiana, for describing a practical situation that motivated the topic studied. Also, helpful comments from an anonymous referee have led to the paper being written in a form more appealing to practitioners. The research leading to this paper was supported by the Natural Sciences and Engineering Research Council of Canada under grants A1485 and 239153-05. It was carried out while the second author held a postdoctoral position at the University of Calgary. Associated financial support from the Haskayne School of Business is gratefully acknowledged.

References

Bashyam S and Fu MC (1998). Optimization of (*s*, *S*) inventory systems with random lead times and a service level constraint. *Mngt Sci* **44**: 243-256.

Brown, RG (1967). *Decision Rules for Inventory Management.* Holt, Rinehart and Winston: New York, p. 238.

Burgin TA (1969). Time to sell a fixed quantity of stock depleted by distributed demand. *Operational Research Quarterly* **20**: 421-428.

Carter G and Ignall E (1975). Virtual measures: a variance reduction technique. *Mngt Sci* **21**: 607-616.

Cox DR and Miller HD (1965). *The Theory of Stochastic Processes*. John Wiley & Sons: New York.

Ehrhardt R (1979). The power approximation for computing (*s, S*) policies. *Mngt Sci* **25**: 777-786.

Ehrhardt R and Mosier C (1984). Revision of the power approximation for computing (*s*, *S*) policies. *Mngt Sci* **30**: 618-622.

IBM (1971) System/360 Inventory Control (360 A-MF-04X). Program description manual GH20-0555-1. New York.

Law AM and Kelton WD (2000). *Simulation Modeling and Analysis*. Third Edn. McGraw-Hill: Boston.

Miltenburg GJ and Silver EA (1984). The diffusion process and residual stock in periodic review, coordinated control of families of items. *Int J Prod Res* **22**: 629-646.

Moors JJA and Strijbosch LWG (2002). Exact fill rates for (*R*, *s*, *S*) inventory control with gamma distributed demand. *J Opl Res Soc* **53**: 1268-1274.

Robb DJ and Silver EA (1998). Inventory management with periodic ordering and minimum order quantities. *J Opl Res Soc* **49**: 1085-1094.

Roberts DM (1962). Approximations to Optimal Policies in a Dynamic Inventory Model. In: Arrow K, Karlin S and Suppes P (eds.), *Studies in Applied Probability and Management Science.* Stanford University Press: Stanford, pp. 207-229.

Schneider H (1978). Methods for determining the reorder point of an (*s*, *S*) ordering policy when a service level is specified. *J Opl Res Soc* **29**: 1181-1193.

Schneider H (1981). Effect of service-levels on order-points or order-levels in inventory models. *Int J Prod Res* **19**: 615-531.

Schneider H and Ringuest JL (1990). Power approximation for computing (s, S) policies using service level. *Mngt Sci* **36**: 822-834.

Silver EA, Naseraldin H and Bischak DP (2008). Adjusting the estimates of $\mathrm{E}(\tau)$ and $\mathrm{Var}(\tau)$ in a diffusion model of an (R, s, S) control system. Working paper, Haskayne School of Business, University of Calgary.

Silver EA, Pyke D and Peterson R (1998). *Inventory Management and Production Planning and Scheduling.* Third Edn. John Wiley & Sons: New York.

Stata Corp. (2003). *Stata Statistical Software: Release 8.0*. Stata Corporation: College Station, TX.

Strijbosch LWG and Moors JJA (2006). Modified normal distributions in (R, S)-inventory control. *Eur J Opl Res* **172**: 201-212.

Tijms HC and Groenevelt H (1984). Simple approximations for the reorder point in periodic and continuous review (s, S) inventory systems with service level constraints. *Eur J Oper Res* **17**: 175-190.

Tyworth JE and O'Neill L (1997). Robustness of the normal approximation of lead-time demand in a distribution setting. *Nav Res Logist* **44**: 165-186.

Zheng YS and Federgruen A (1991). Finding optimal (s, S) policies is about as simple as evaluating a single policy. *Oper Res* **39**: 654-665.

Appendix A – Derivations Leading to the Procedure for Selecting Values of *s* and *S*

A.1 $E(\tau)$ *and* $Var(\tau)$ *for a given value of* $S - s$

As indicated in Figure 2 (which reflects setting $R = 1$), τ is the time from the instant that the inventory position first reaches *s* until the moment of the next review. The behaviour of cumulative normally distributed demand (hence the change in the inventory position away from the order-up-to-level, *S*) can be modeled as a continuous time diffusion process (Miltenburg and Silver, 1984). Moreover, the probability density function of the first passage time, *u*, for the inventory position to drop from *S* to *s* is given by (Cox and Miller, 1965)

$$f_u(u_0) = \frac{S-s}{\sigma\sqrt{2\pi u_0^3}} \exp\left[-\frac{(S-s-\mu u_0)^2}{2\sigma^2 u_0}\right] \qquad 0 < u_0$$

Setting

$$m = \frac{S-s}{\mu} \tag{7}$$

and

$$CV = \frac{\sigma}{\mu} \tag{8}$$

we obtain

$$f_u(u_0) = \frac{m}{CV\sqrt{2\pi u_0^3}} \exp\left[-\frac{(m-u_0)^2}{2(CV)^2 u_0}\right] \qquad 0 < u_0 \tag{9}$$

As an aside, Burgin (1969) developed an analytic expression for the expected value of *u*.

Now, multiple values of *u* can produce the same value of τ. Specifically $\tau = \tau_0$ will result from any of $u = 1 - \tau_0$, $2 - \tau_0$, etc. Hence, the density function of τ is given by

$$f_\tau(\tau_0) = \sum_{i=1}^{\infty} f_u(i-\tau_0)$$

Using (9) we have

$$f_\tau(\tau_0) = \frac{m}{CV\sqrt{2\pi}} \sum_{i=1}^{\infty} \frac{1}{\sqrt{(i-\tau_0)^3}} \exp\left[-\frac{(m+\tau_0-i)^2}{2(CV)^2(i-\tau_0)}\right] \qquad 0<\tau_0<1 \qquad (10)$$

Unfortunately, it is not possible to analytically develop expressions for the moments of τ. However, note that the density function, hence the moments, depend on only two parameters, namely *m* and *CV*. For given values of these parameters τ_0 can be discretized on a fine grid and accurate estimates of $\mathrm{E}(\tau)$ and $\mathrm{Var}(\tau)$ can be found by numerical integration.

A.2 Determining $S-s$ to target the desired average time between replenishments (also leading to Step 1 of the procedure)

For a given value of τ the expected size of the undershoot

$$\mathrm{E}(\text{undershoot} \mid \tau_0) = \mu\tau_0$$

Hence $\mathrm{E}(\text{undershoot}) = \mu\mathrm{E}(\tau)$. Let *Q* denote the order size. Then, the average order size

$$\mathrm{E}(Q) = S-s+\mathrm{E}(\text{undershoot}) = S-s+\mu\mathrm{E}(\tau)$$

The average time between replenishments

$$\mathrm{E}(t) = \frac{\mathrm{E}(Q)}{\mu}$$

Using (7), $\mathrm{E}(t) = m+\mathrm{E}(\tau)$. But we want to target an integer *n* for E(*t*). Thus we require

$$m+\mathrm{E}(\tau) = n \qquad (11)$$

Now $\mathrm{E}(\tau)$ is a function of both *m* and *CV*. Thus, for a given value of *CV* we have to find the *m* that satisfies (11). This was done using Mathematica where for each value of *m*, as mentioned earlier, numerical integration is needed to estimate $\mathrm{E}(\tau)$. Rather than providing detailed tables of $\mathrm{E}(\tau)$ vs. *CV*, for several practical values of *n* we have instead used Stata (StataCorp. 2003) to fit fractional polynomial functions of degree 3 (with very high R^2 values). The results of fitting functions to 51 values of $\mathrm{E}(\tau)$, resulting from equi-spaced values of *CV* between 0.1 and 0.5, are shown in Table 1. Note that the powers of *CV* have been adjusted to the average *CV* value (0.3) in the approximation expressions; see the Stata Base Reference Manual (vol. 1, pp. 402-3) for further information.

In the above, once the *m* is found that satisfies (11) for given values of *n* and *CV*, we then use numerical integration to estimate $\mathrm{Var}(\tau)$. Again, very accurate fractional polynomial functions were fit to $\mathrm{Var}(\tau)$ as a function of *CV* for selected values of *n*, as shown in Table 2.

A.3 Determining s to give a desired service level (Steps 2, 3 and 4 of the procedure)

As discussed in the previous subsection, for given values of *n* and *CV* we are able to determine $\mathrm{E}(\tau)$ and $\mathrm{Var}(\tau)$. Let *X* represent the total demand in the effective lead time $\tau + L$. As evident in Figure 1, a shortage will occur if *X* exceeds the reorder point *s*. We can model the situation as a continuous review, reorder point system where the effective lead time $\tau + L$ is a random variable with mean $\mathrm{E}(\tau) + L$ and variance $\mathrm{Var}(\tau)$. Under such circumstances *X* has moments (see Silver et al., 1998, p. 283)

$$\mathrm{E}(X) = [\mathrm{E}(\tau) + L]\mu \tag{12}$$

and

$$\mathrm{Var}(X) = [\mathrm{E}(\tau) + L]\sigma^2 + \mu^2 \mathrm{Var}(\tau) \tag{13}$$

Dividing (12) by μ and (13) by μ^2 we have

$$\frac{\mathrm{E}(X)}{\mu} = \mathrm{E}(\tau) + L \tag{14}$$

and

$$\frac{\mathrm{Var}(X)}{\mu^2} = [\mathrm{E}(\tau) + L](CV)^2 + \mathrm{Var}(\tau)$$

Taking the square root of the latter results in

$$\frac{\sigma_X}{\mu} = \sqrt{[\mathrm{E}(\tau) + L](CV)^2 + \mathrm{Var}(\tau)} \tag{15}$$

In the usual fashion, set $s = \mathrm{E}(X) + k\sigma_X$. It is again convenient to divide through by μ to obtain

$$\frac{s}{\mu} = \frac{\mathrm{E}(X)}{\mu} + k\frac{\sigma_X}{\mu} \tag{16}$$

If we assume that *X* is normally distributed (which we know is an approximation because τ is a random variable; however, Tyworth and O'Neill, 1997, and Silver et al., 1998, pp. 272-3, argue that as long as $\sigma_X / \mathrm{E}(X) \leq 0.5$, there is little risk in making the normality assumption), then use of (8) leads to the expected units short per replenishment cycle (see Silver et al., 1998)

$$\mathrm{EUSPRC} \simeq \sigma_X G_u(k) \tag{17}$$

where

$$G_u(k) = \int_k^\infty (z_0 - k)\phi(z_0)dz_0 \tag{18}$$

is the unit normal loss function and $\phi(z_0)$ is the unit normal density function. Note that there is a more precise version of (18) which turned out to not be required in our numerical experiments, specifically

$$\mathrm{EUSPRC} = \sigma_X \left[G_u(k) - G_u\left(k + \frac{n\mu}{\sigma_X}\right) \right] \tag{19}$$

The target allowed (average) units short per replenishment cycle

$$\text{AUSPRC} = (1-P_2)\text{E}(Q) \tag{20}$$

From (17) and (20) and noting that E(Q) = $n\mu$, we thus require

$$G_u(k) = \frac{(1-P_2)n\mu}{\sigma_X} \tag{21}$$

A.4 Determining the S value (Step 5)

Use of (7) and (11) gives

$$\frac{S-s}{\mu} + \text{E}(\tau) = n$$

Hence

$$\frac{S}{\mu} = \frac{s}{\mu} + n - \text{E}(\tau) \tag{22}$$

Appendix B – Some Technical Details of the Simulation

To estimate the fill rate achieved, we have used the "virtual measures" approach of variance reduction first suggested by Ignall and Carter (1975). Specifically, let $I_i (\leq s)$ be the inventory position when the ith order (of size Q_i) is placed. Then (for reference purposes see (19)) the expected units short at the end of the lead time conditional on I_i can be computed analytically as

$$\text{EUS}(I_i) = \sqrt{L}\sigma\left[G_u\left(\frac{I_i - L\mu}{\sqrt{L}\sigma}\right) - G_u\left(\frac{I_i + Q_i - L\mu}{\sqrt{L}\sigma}\right)\right]$$

Over a large number (N) of replenishments the estimated average number of units short per replenishment is

$$\text{EUS} = \frac{1}{N}\sum_{i=1}^{N}\text{EUS}(I_i)$$

Also, $I_i + Q_i = S$. Thus

$$\frac{\text{EUS}}{\sqrt{L}\sigma} = \frac{1}{N}\sum_{i=1}^{N} G_u\left(\frac{I_i - L\mu}{\sqrt{L}\sigma}\right) - G_u\left(\frac{S - L\mu}{\sqrt{L}\sigma}\right)$$

where the last term on the right side is a constant. Consequently, we only have to compute

$G_u\left(\frac{I_i - L\mu}{\sqrt{L}\sigma}\right)$ for each replenishment *i*.

Now an estimate of the average order size is given by

$$\text{E}(Q) = \frac{1}{N}\sum_{i=1}^{N} Q_i$$

and the expected fill rate is

$$FR = 1 - \frac{\text{EUS}}{\text{E}(Q)}$$

The other performance measure we need is the average time between replenishments

$$AT = \frac{\mu}{\text{E}(Q)}$$

Besides point estimates of each of *FR* and *AT*, confidence intervals were developed based upon the replication-deletion method for estimating a steady-state mean (see, e.g., Law and Kelton, 2000).

Figure 1 – Undershoots in a Periodic Review System

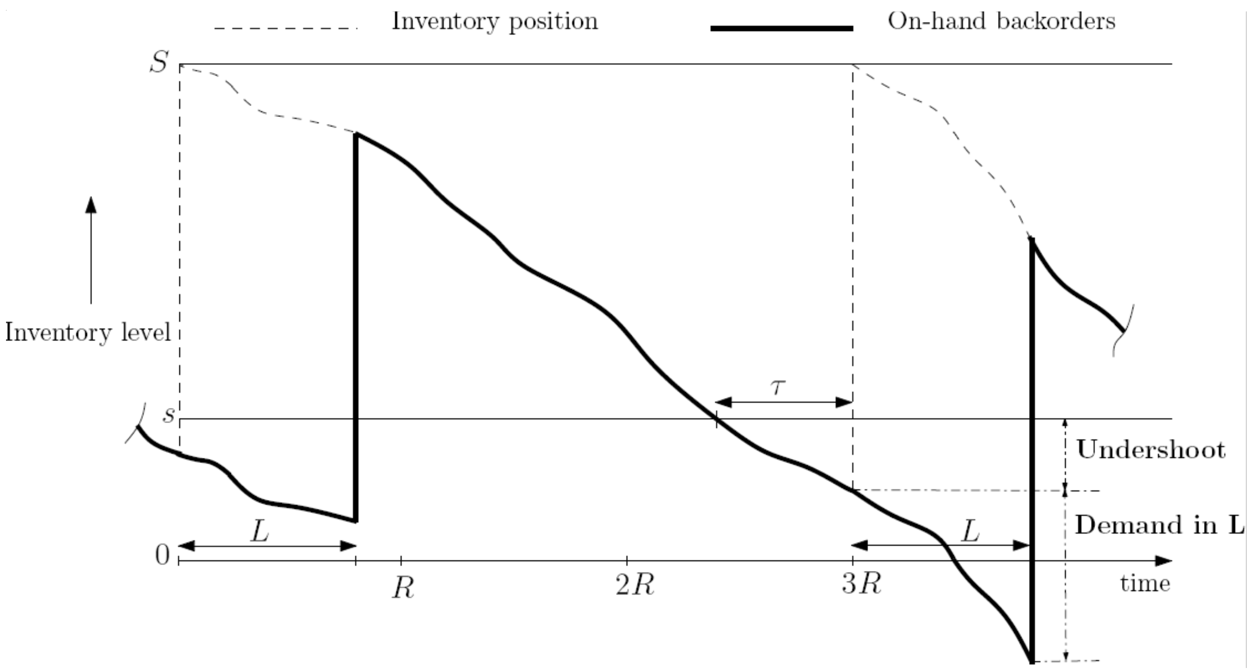

Figure 2 – Relation Between First Passage Time (u) and τ

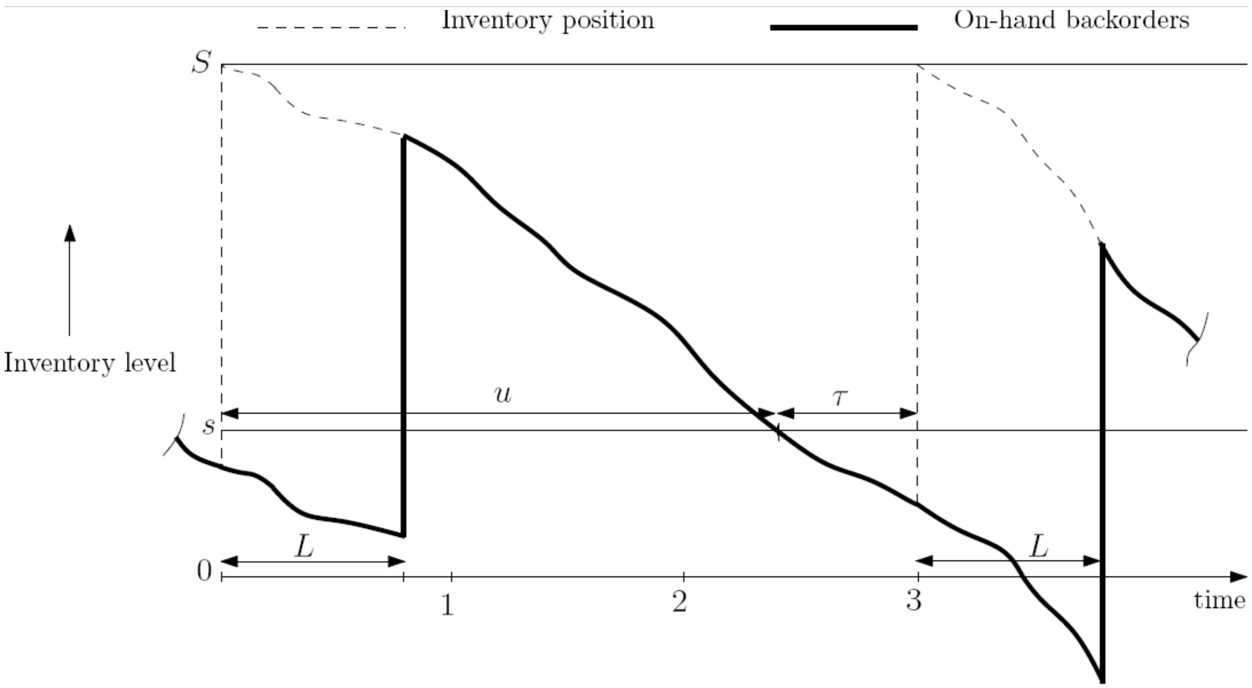

Table 1 – Fractional Polynomial Approximations of $\mathrm{E}(\tau)$ as a Function of *CV* for Selected Values of *n*

n	Approximation (note $C \equiv CV$)
2	$.53608 + .44271(C^{-1} - 3.333) + 1.7634(C^{-½} - 1.826) + 1.0508(C^{-½} \ln C + 2.198)$
3	$.51211 + 1.8652(C - .3) - 1.1430(C^{2} - .09) + 3.1367(C^{2} \ln C + .1084)$
4	$.50325 + 2.1455(C^{½} - .5477) - .65943(C^{½} \ln C + .6594) + .50973(C^{½} (\ln C)^{2} - .794)$
5	$.50079 - .13438(C^{-1} - 3.333) - .39946(C^{-½} - 1.826) - .28188(C^{-½} \ln C + 2.198)$
6	$.50004 - .00237(C^{-2} - 11.11) - .03307(C^{-1} - 3.333) - .02296 (C^{-1} \ln C + 4.013)$

Table 2 – Fractional Polynomial Approximations of $\mathrm{Var}(\tau)$ as a Function of *CV* for Selected Values of *n*

n	Approximation (note $C \equiv CV$)
2	$.07401 + .53380(C - .3) - .58217(C^{2} - .09)$
3	$.08283 + .36794(C^{½} - .5477) - .35809(C^{½} \ln C + .6594)$
4	$.08387 - .12888(C^{-½} - 1.826) - .10939(\ln C + 1.204)$
5	$.08371 + .01876(C^{-1} - 3.333) + .00887(C^{-1} \ln C + 4.013)$
6	$.08352 - .00078(C^{-2} - 11.11) + .00503(C^{-1} - 3.333)$

Table 3 – Settings of Independent Parameters

Parameter	Definition	Values Selected
n	target average number of review (unit) intervals between replenishments	2, 4, 6
CV	coefficient of variation of (the normally distributed) demand in a unit interval	0.1, 0.3, 0.5
L	lead time	0.5, 2, 4
P_2	target fill rate	0.8, 0.9, 0.99

Table 4 – Simulation Results

P_2	L	n	CV	Average Time (AT) Between Orders Observed	% Deviation from Target, n	Fill Rate (FR) Observed	% Deviation from Target, P_2
0.8	0.5	2	0.1	1.980	-0.99%	0.800	-0.01%
0.8	0.5	2	0.3	2.023	1.13%	0.790	-1.29%
0.8	0.5	2	0.5	2.102	5.11%	0.776	-3.03%
0.8	0.5	4	0.1	3.965	-0.88%	0.799	-0.11%
0.8	0.5	4	0.3	4.005	0.12%	0.792	-0.99%
0.8	0.5	4	0.5	4.079	1.97%	0.780	-2.45%
0.8	0.5	6	0.1	5.948	-0.87%	0.799	-0.11%
0.8	0.5	6	0.3	5.985	-0.25%	0.794	-0.74%
0.8	0.5	6	0.5	6.061	1.01%	0.785	-1.90%
0.8	2	2	0.1	1.980	-0.99%	0.800	-0.01%
0.8	2	2	0.3	2.023	1.13%	0.790	-1.23%
0.8	2	2	0.5	2.101	5.06%	0.779	-2.68%
0.8	2	4	0.1	3.964	-0.89%	0.799	-0.11%
0.8	2	4	0.3	4.005	0.14%	0.793	-0.94%
0.8	2	4	0.5	4.081	2.03%	0.783	-2.08%
0.8	2	6	0.1	5.947	-0.88%	0.799	-0.11%
0.8	2	6	0.3	5.987	-0.21%	0.794	-0.74%
0.8	2	6	0.5	6.063	1.05%	0.786	-1.71%
0.8	4	2	0.1	1.980	-0.99%	0.800	-0.02%
0.8	4	2	0.3	2.022	1.11%	0.791	-1.14%
0.8	4	2	0.5	2.101	5.06%	0.781	-2.39%
0.8	4	4	0.1	3.964	-0.89%	0.799	-0.09%
0.8	4	4	0.3	4.006	0.16%	0.793	-0.86%
0.8	4	4	0.5	4.079	1.96%	0.786	-1.81%
0.8	4	6	0.1	5.947	-0.88%	0.799	-0.11%
0.8	4	6	0.3	5.986	-0.24%	0.795	-0.67%
0.8	4	6	0.5	6.060	0.99%	0.788	-1.57%
0.9	0.5	2	0.1	1.980	-0.98%	0.900	-0.01%
0.9	0.5	2	0.3	2.023	1.14%	0.890	-1.11%
0.9	0.5	2	0.5	2.102	5.11%	0.879	-2.28%
0.9	0.5	4	0.1	3.964	-0.90%	0.899	-0.07%
0.9	0.5	4	0.3	4.007	0.18%	0.894	-0.70%
0.9	0.5	4	0.5	4.077	1.93%	0.885	-1.69%
0.9	0.5	6	0.1	5.948	-0.87%	0.899	-0.08%
0.9	0.5	6	0.3	5.985	-0.26%	0.895	-0.58%
0.9	0.5	6	0.5	6.061	1.01%	0.887	-1.45%
0.9	2	2	0.1	1.980	-0.99%	0.900	-0.01%
0.9	2	2	0.3	2.022	1.12%	0.892	-0.92%
0.9	2	2	0.5	2.102	5.12%	0.884	-1.74%
0.9	2	4	0.1	3.964	-0.89%	0.899	-0.08%

Table 4 – Simulation Results (continued)

				Average Time (AT) Between Orders		Fill Rate (FR)	
P_2	L	n	CV	Observed	% Deviation from Target, n	Observed	% Deviation from Target, P_2
0.9	2	4	0.3	4.007	0.17%	0.894	-0.70%
0.9	2	4	0.5	4.082	2.05%	0.887	-1.47%
0.9	2	6	0.1	5.947	-0.88%	0.899	-0.09%
0.9	2	6	0.3	5.988	-0.20%	0.895	-0.56%
0.9	2	6	0.5	6.067	1.12%	0.889	-1.27%
0.9	4	2	0.1	1.980	-0.99%	0.900	-0.00%
0.9	4	2	0.3	2.022	1.11%	0.893	-0.74%
0.9	4	2	0.5	2.102	5.09%	0.887	-1.42%
0.9	4	4	0.1	3.964	-0.90%	0.899	-0.07%
0.9	4	4	0.3	4.006	0.15%	0.894	-0.63%
0.9	4	4	0.5	4.080	1.99%	0.889	-1.20%
0.9	4	6	0.1	5.947	-0.89%	0.899	-0.08%
0.9	4	6	0.3	5.987	-0.21%	0.895	-0.53%
0.9	4	6	0.5	6.065	1.09%	0.890	-1.10%
0.99	0.5	2	0.1	1.980	-0.98%	0.990	-0.00%
0.99	0.5	2	0.3	2.021	1.07%	0.987	-0.36%
0.99	0.5	2	0.5	2.102	5.11%	0.983	-0.76%
0.99	0.5	4	0.1	3.965	-0.88%	0.990	-0.04%
0.99	0.5	4	0.3	4.006	0.15%	0.987	-0.35%
0.99	0.5	4	0.5	4.083	2.08%	0.984	-0.60%
0.99	0.5	6	0.1	5.947	-0.89%	0.989	-0.07%
0.99	0.5	6	0.3	5.986	-0.24%	0.987	-0.30%
0.99	0.5	6	0.5	6.062	1.03%	0.985	-0.49%
0.99	2	2	0.1	1.980	-0.98%	0.990	-0.00%
0.99	2	2	0.3	2.022	1.11%	0.988	-0.24%
0.99	2	2	0.5	2.100	5.01%	0.986	-0.42%
0.99	2	4	0.1	3.965	-0.89%	0.990	-0.04%
0.99	2	4	0.3	4.005	0.13%	0.988	-0.24%
0.99	2	4	0.5	4.083	2.07%	0.987	-0.36%
0.99	2	6	0.1	5.947	-0.88%	0.990	-0.05%
0.99	2	6	0.3	5.989	-0.19%	0.988	-0.20%
0.99	2	6	0.5	6.062	1.03%	0.987	-0.32%
0.99	4	2	0.1	1.980	-0.99%	0.990	0.00%
0.99	4	2	0.3	2.023	1.15%	0.988	-0.17%
0.99	4	2	0.5	2.102	5.09%	0.987	-0.28%
0.99	4	4	0.1	3.964	-0.90%	0.990	-0.03%
0.99	4	4	0.3	4.005	0.13%	0.988	-0.17%
0.99	4	4	0.5	4.076	1.90%	0.987	-0.26%
0.99	4	6	0.1	5.947	-0.88%	0.990	-0.04%
0.99	4	6	0.3	5.986	-0.23%	0.989	-0.15%
0.99	4	6	0.5	6.064	1.06%	0.988	-0.25%

Figure Captions

Figure 1 – Undershoots in a Periodic Review System

Figure 2 – Relation Between First Passage Time (u) and τ

Table Headings

Table 1 – Fractional Polynomial Approximations of $\mathrm{E}(\tau)$ as a Function of *CV* for Selected Values of *n*

Table 2 – Fractional Polynomial Approximations of $\mathrm{Var}(\tau)$ as a Function of *CV* for Selected Values of *n*

Table 3 – Settings of Independent Parameters

Table 4 – Simulation Results