The Effect of Lead Time Uncertainty On Safety Stocks.

Sunil Chopra
Kellogg Graduate School of Management
Northwestern University
Evanston, Illinois, 60201
(847) 491-8169
(847) 467-1220 (fax)
s-chopra@kellogg.northwestern.edu

Gilles Reinhardt(*)
Department of Management
DePaul University
Chicago, Illinois, USA, 60604
(312) 362-8849
(312) 362-6973 (fax)
greinhar@condor.depaul.edu

Maqbool Dada
Krannert Graduate School of Management
Purdue University
West Lafayette, Indiana, 47907
(765) 494-4490
(765) 496-1778 (fax)
dada@mgmt.purdue.edu

(*) Corresponding Author

The Effect of Lead-time Uncertainty on Safety Stocks

The pressure to reduce inventory investments in the supply chain have increased as competition expands and product variety grows. Managers are looking for areas they can improve to reduce inventories required without hurting the level of service provided to customers. Two important areas that managers focus on are the reduction of the replenishment lead time from suppliers and the variability of this lead time. The normal approximation of lead time demand distribution indicates that both actions reduce inventories for cycle service levels above 50 percent. The normal approximation also indicates that reducing lead time variability tends to have a greater impact than reducing lead times, especially when lead time variability is large. We build on the work of Eppen and Martin to show that the conclusions from the normal approximation are flawed, especially in the range of service levels where most companies operate. We show the existence of a service level threshold greater than 50 percent below which re-order points increase with a decrease in lead time variability. Thus, for a firm operating just below this threshold, reducing lead times decreases reorder points, whereas reducing lead time variability increases reorder points. For firms operating at these service levels, decreasing lead time is the right lever if they want to cut inventories, not reducing lead time variability.

1. Introduction and Framework

Managers have been under increasing pressure to decrease inventories as supply chains attempt to become leaner. The goal, however, is to reduce inventories without hurting the level of service provided to customers. Safety stock is a function of the cycle service level, demand uncertainty, the replenishment lead time, and the lead time uncertainty. For a fixed cycle service level, a manager thus has three levers that affect the safety stock - demand uncertainty, replenishment lead time, and lead time uncertainty. In this paper we focus on the relationship between lead time uncertainty and safety stock and the resulting implications for management.

Traditionally, a normal approximation has been used to estimate the relationship between safety stock and demand uncertainty, replenishment lead time, and lead time uncertainty. According to Eppen and Martin (1988), this approximation is often justified by using an argument based on the central limit theorem but in reality, they say, "the normality assumption is unwarranted in general and this procedure can produce a probability of stocking out that is egregiously in error". Silver and Peterson (1985), however, argue that trying to correct this effect with a more accurate representation of demand during lead time may be ineffectual because the gain in precision may only induce minimal improvement in the cost. Tyworth and O'Neill (1997) also address this issue in a detailed empirical study for *fast-moving* finished goods (demand per unit time have c.v.'s below 40%) in seven major industries. Their investigations reveal that "the normal approximation method can lead to large errors in contingency stock – say, greater than 25%. Such errors have relatively little influence on the optimal solutions, however, because contingency stock holding cost comprises a small portion of the total logistics system cost.” They further conjecture that reducing the *fill rate*, the proportion of orders filled from stock, “makes total system costs less sensitive to normal theory misspecifications” since this will in turn reduce the required safety stock level and thus make the total holding costs a smaller percentage of total system cost.

In this paper our focus is not on the size of the error resulting from using the normal approximation (that has been captured very well by Eppen and Martin) but on the flaws in the managerial prescriptions implied by the normal approximation. In particular, we focus on two prescriptions of the normal approximation:

1. For cycle service levels above 50 percent, reducing lead time variability reduces the reorder point and safety stock.
2. For cycle service levels above 50 percent, reducing lead time variability is more effective than reducing lead times because it decreases the safety stock by a larger amount.

In this paper we show that for cycle service levels that are commonly used in industry both prescriptions are false if we consider the exact demand during the lead time. Using the exact demand during the lead time instead of the normal approximation we infer the following:

1. For cycle service levels above 50 percent but below a threshold, reducing lead time variability **increases** the reorder point and safety stock.
2. For cycle service levels above 50 percent but below a threshold, reducing the lead time variability **increases** the reorder point and safety stock, whereas reducing the lead time **decreases** the reorder point and safety stock.

Both effects are more pronounced when the coefficient of variation of demand is high and less pronounced when the coefficient of variation of demand is low. This is consistent with the conclusion of Tyworth and O'Neill (1997) that the normal approximation is quite effective for low c.v. Our inference also support the results in Jong (1994) who assumes a periodic demand that follows a compound Poisson process and derives a threshold value underneath which base stocks increase with a reduction in lead-time uncertainty. It is easy to see that under the normal approximation, this threshold equals 0.5. For the distributions we analyze, assuming a normal period demand, we show that this threshold lies in a range where most firms operate (between 0.5 and 0.7). The comparison of the prescriptions is illustrated using Table 1.

Row	**Lead Time Process**	**CSL (α)**	**Safety Stock (Normal Approximation)**	**Safety Stock (Exact value)**
1	Gamma, L=10, s_L = 5	0.6	28	20
2	Gamma, L=10, s_L=4	0.6	23	22
3	Gamma, L=8, s_L=5	0.6	27	15
4	Gamma, L=10, s_L = 5	0.95	182	218
5	Gamma, L=10, s_L=4	0.95	153	181
6	Gamma, L=8, s_L=5	0.95	179	218

Table 1: Safety Stocks for Gamma Lead Times and Different Service Levels

Consider rows 1-3 of Table 1. For a cycle service level of 0.6, the normal approximation predicts that reducing the standard deviation of the lead time from 5 to 4 should decrease the safety stock from 28 to 23. The exact calculation, however, shows that reducing the standard deviation of lead time **increases** the required safety stock from 20 to 22. The normal approximation predicts that reducing the standard deviation of lead time by 20 per cent (5 to 4) is much more effective at reducing the safety stock than reducing the lead time by 20 percent (10 to 8). The exact calculation, however, shows that for a cycle service level of 0.6, decreasing lead time is more effective (safety stock decreases from 20 to 15) than reducing the standard deviation of lead time (safety stock increases from 20 to 22).

For a cycle service level of 95 percent, however, the prescriptions of the normal approximation are correct (see rows 4-6). At this cycle service level both the normal approximation and the exact calculation show that reducing the standard deviation of lead time decreases the safety stock and is more effective than decreasing the lead time itself.

We next argue that many firms in practice operate at cycle service levels in the 50-70 percent range rather than the 95-99 percent that is often assumed. In practice, managers often focus on the *fill rate* as a service quality measure (Aiginger, 1987, Lee and Billington, 1992, Byrne and Markham, 1991), rather than the cycle service level (*CSL* or α). The fill rate measures the proportion of demand that is met from stock, whereas the cycle service level measures the proportion of replenishment cycles where a stockout does not occur. Table 2 considers the cycle service level and fill rate for different reorder points for a product that has a weekly demand of 2,500, standard deviation of weekly demand of 500, lead time of 2 weeks, and a reorder quantity of 10,000. All calculations for fill rate are as detailed in Chapter 11 of Chopra and Meindl, 2003.

Reorder Point	Safety Stock	Cycle Service Level	Fill Rate
5000	0	0.500	0.9718
5040	40	0.523	0.9738
5080	80	0.545	0.9756
5120	120	0.567	0.9774
5160	160	0.590	0.9791
5200	200	0.611	0.9807

5240	240	0.633	0.9822
5280	280	0.654	0.9836
5320	320	0.675	0.9850
5360	360	0.695	0.9862
5400	400	0.714	0.9874

Table 2: Cycle Service Level and Fill Rate as a Function of Safety Stock

In this example, Table 2 illustrates that fill rates of between 97-99 percent are achieved for cycle service levels between 50-70 percent. Most firms aim for fill rates of between 97-99 percent (and not cycle service levels). This implies cycle service levels of between 50-70 percent. As we show in this paper, it is for cycle service levels between 50-70 percent that the prescriptions of the normal approximation to managers are most distorted and lead to managers pushing the wrong levers to reduce inventories. Our main point is that for cycle service levels where most firms operate, the normal approximation erroneously encourages managers to focus on reducing the variability of lead times when they would be better of reducing the lead time itself.

Our general results range from specific theoretical outcomes when the lead time follows a uniform distribution (Section 3) to numerical observations when the lead time follows a uniform, gamma, or normal distribution (Section 4). In the next section, we formalize our model and re-examine the response to reducing uncertainty when the normal approximation is used rather than an exact characterization. We conclude with the scope and managerial implications of our findings in Section 6.

2. Effect of Lead-time Uncertainty: The Normal Approximation

For a given cycle service level, determining the required safety stock levels is predicated on characterizing the distribution of demand during the lead time. We assume that there is an indivisible period of analysis; for example, a day. Demand during day *i*,

X_i, are independent and identically distributed random variables drawn from a normal distribution with mean μ_X and standard deviation σ_X. For a generic lead time distribution of mean L and standard deviation s_L, the demand during lead time under the normal approximation has mean $M = L\mu_X$ and standard deviation $S = \sqrt{L\sigma_X^2 + \mu_X^2 s_L^2}$ (Silver and Peterson, 1985).

Let F(•) represent the cumulative distribution function (cdf) of the standard normal distribution with mean 0 and standard deviation 1. Define z as the solution to F(z) = α and ROP_N as the re-order point for a cycle service level of α. Under the normal approximation, we have $ROP_N = M + zS$ with zS as the safety stock. Observe that under the normal approximation S increases (decreases) as s_L increases (decreases). Thus, whether the safety stock S and the reorder point ROP_N, rise or fall as s_L increases depends only on the sign of z. For a given mean lead-time L, Figure 1 depicts the relationship between the lead-time uncertainty (represented by s_L) and the re-order points predicted by the normal approximation for three α's. As can be seen from Figure 1, the safety stock S and reorder point ROP_N rise with an increase in s_L for a CSL above 0.5 (since $z > 0$) and drop with an increase in s_L for CSL below 0.5 (since $z < 0$). For a CSL of 0.5, the re-order point remains at the level of the deterministic case ($s_L = 0$) and does not change with an increase in s_L (since $z = 0$, $ROP_N = M$). These observations lead to the following conclusion:

Theorem 2.1: *Suppose that the lead-time uncertainty represented by s_L increases. Then for a given CSL = α, the following is true:*

(i) *If $\alpha < .5$, then ROP_N falls;*

(ii) *if $\alpha = .5$, then ROP_N is invariant; and*

(iii) *if* $\alpha > .5$*, then* ROP_N *rises.*

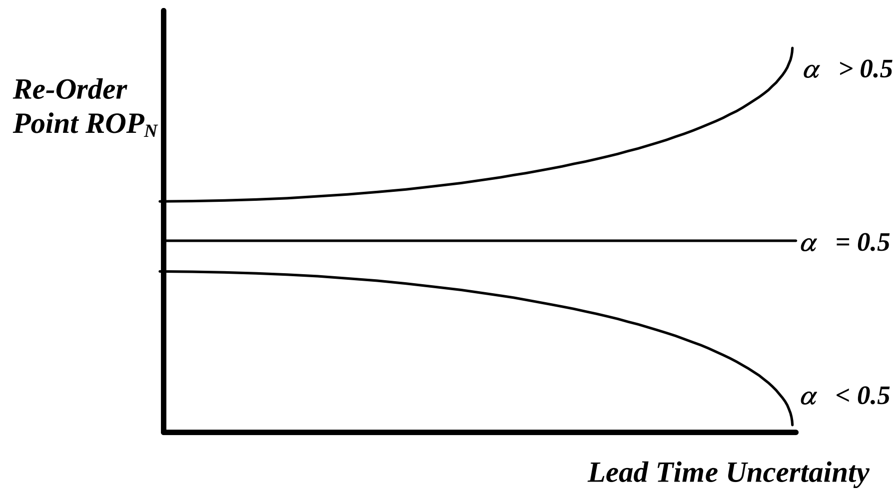

Figure 1

Theorem 2.1 indicates that, as management works on the reduction of lead-time uncertainty (reduction of s_L), the re-order point drops for CSLs above 0.5. Unfortunately, as demonstrated in Eppen and Martin (1988), this neat prescription is a consequence of the normal approximation. In the next section, we show the existence of a threshold $\bar{\alpha} > 0.5$ such that for CSLs in (0.5, $\bar{\alpha}$], the reorder point and safety stock actually increase as s_L decreases.

3. Effect of Lead-time Uncertainty: The Exact Distribution

In this section, we show how the prescriptions of the normal approximation in Theorem 2.1 are flawed for the case when periodic demand follows the normal distribution and the lead time has a discrete uniform distribution with a mean of Y and a range of $Y \pm y$. Denote the re-order point by R and let $G_y(R)$ be the (unconditional) probability that demand during the lead time is less than or equal to R when the lead-time is uniformly distributed between $Y \pm y$. If μ_x is the expected demand per period and σ_x is the standard deviation of demand per period, we define

$$z_Y(R) = {(R - Y\mu_x)} \big/ {(\sigma_x \sqrt{Y})} . \tag{3.1}$$

It is clear from the definition of $z_Y(R)$ that it represents the number of standard deviations R is away from the expected value of demand given that the lead time is Y. Let $F(z_Y(R))$ represent the probability that the standard normal is less than or equal to $z_Y(R)$. As in Eppen and Martin (1988) it then follows that

$$G_y(R) = (\frac{1}{2y+1}) \sum_{W=Y-y}^{Y+y} F(z_W(R)) . \tag{3.2}$$

From (3.1) and (3.2), it thus follows that

$$G_y(R_1) > G_y(R_2) \text{ if and only if } R_1 > R_2 . \tag{3.3}$$

Observe that the case when $y=0$ corresponds to the case of deterministic lead-time. We are interested in examining how $G_y(R)$ behaves as the lead time uncertainty represented by y changes. We begin by examining the effect of increasing uncertainty by increasing y by one period. Then, simple algebra yields:

$$G_{y+1}(R) - G_y(R) = (\frac{1}{2y+3})\{F(z_{Y+y+1}(R)) + F(z_{Y-y-1}(R)) - 2G_y(R)\} , \tag{3.4}$$

and

$$G_{y+1}(R) = (\frac{2y+1}{2y+3}) G_y(R) + (\frac{1}{2y+3})[F(z_{Y+y+1}(R)) + F(z_{Y-y-1}(R))] . \tag{3.5}$$

Since $Y + y + 1 > Y - y - 1 \geq 0$, it readily follows that $z_{Y-y-1}(R) \geq z_{Y+y+1}(R)$, so that

$$1 > F(z_{Y-y-1}(R)) > F(z_{Y+y+1}(R)) > 0 . \tag{3.6}$$

Our objective for the rest of the section is to try and come up with an analogue to Figure for the case when we use the exact distribution of demand during the lead time (shown in (3.2)). To proceed we need the following lemma.

Lemma 3.1: *Let* R_y *and* R_{y+1} *be such that* $G_y(R_y) = G_{y+1}(R_{y+1}) = \alpha$. *We have* $R_{y+1} >$ *(<)* R_y *if and only if* $F(z_{Y+y+1}(R_y)) + F(z_{Y-y-1}(R_y)) < (>) 2G_y(R_y)$.

Proof: Observe that if $F(z_{Y+y+1}(R_y)) + F(z_{Y-y-1}(R_y)) < (>) 2G_y(R_y)$, we have $G_{y+1}(R_y) < (>) G_y(R_y) = \alpha$ by (3.4). Since $G_{y+1}(R_{y+1}) = \alpha$, using (3.3) we thus have $R_{y+1} > (<)\ R_y$.

On the other hand, if $R_{y+1} > (<)\ R_y$, (3.3) implies that $\alpha = G_{y+1}(R_{y+1}) > (<) G_{y+1}(R_y)$. Since $\alpha = G_y(R_y)$, we have $G_y(R_y) > (<) G_{y+1}(R_y)$. From (3.4) we thus have $F(z_{Y+y+1}(R_y)) + F(z_{Y-y-1}(R_y)) < (>) 2G_y(R_y)$. The result thus follows.

●●●●

Another result needed is presented below. The proof follows from the definition of the standard normal distribution.

Lemma 3.2: *Let F(.) be the standard normal cumulative distribution function. If* $z_1 < 0 < z_2$, *then* $1 < F(z_1) + F(z_2)$ *if and only if* $-z_1 < z_2$.

We start by considering the reorder point changes as y increases for the case where the CSL is 0.5. For a given value of lead-time uncertainty y, let $R_y(0.5)$ be the re-order point such that $G_y(R_y(0.5)) = 0.5$. For the case $y = 0$, observe that $R_y(0.5)$ is the expected demand, $M = Y\mu_x$, during the lead time Y. We have $z_{Y+1}(R_0(0.5)) < 0 <$

$-z_{Y+1}(R_0(0.5)) < z_{Y-1}(R_0(0.5))$ from (3.1). From Lemma 3.2 it thus follows that $F(z_{Y+1}(R_0(0.5))) + F(z_{Y-1}(R_0(0.5))) > 2\,G_0(R_0(0.5)) = 1$. Using Lemma 3.1 it thus follows that $R_1(0.5) < R_0(0.5)$.

●●●●

In other words, the re-order point decreases as lead-time uncertainty increases from $y = 0$ to $y = 1$ for a cycle service level of 0.5. In the next result we prove that this pattern continues to hold as lead-time uncertainty (y) increases, i.e., the re-order point continues to drop as lead-time uncertainty (y) increases for a cycle service level of 0.5.

Theorem 3.3: *For a cycle service level α = 0.5, the re-order point $R_y(0.5)$ declines with an increase in lead-time uncertainty y, i.e., $R_{y+1}(0.5) < R_y(0.5)$.*

●●●●

Theorem 3.3 is equivalent to stating that the median of the distribution of demand during the lead time declines as lead-time uncertainty represented by y increases. In contrast, the median is invariant when the normal approximation is used. For the specific case of the median, Theorem 3.3 provides a complete characterization of the behavior of the reorder point as y increases. In general, the reorder point is the solution to

$$G_y(R) = \left(\frac{1}{2y+1}\right)\sum_{W=Y-y}^{Y+y} F(z_W(R)) = \alpha\,. \tag{3.7}$$

Let $R_y(\alpha)$ represent the unique solution to (3.7). To examine the effect of increasing the cycle service level α it is sufficient to specialize (3.4) to:

$$G_{y+1}(R_y(\alpha)) - G_y(R_y(\alpha)) = \left(\frac{1}{2y+3}\right)\{F(z_{Y+y}(R_y(\alpha))) + F(z_{Y-y}(R_y(\alpha))) - 2\alpha\} \tag{3.8}$$

Observe that Theorem 3.3 implicitly analyzes (3.8) for the special case $\alpha = .5$. By Lemma 3.1, for arbitrary α, determining whether the re-order point increases or decreases depends on the sign of the term $2\alpha - [F(z_{Y+y}(R_y(\alpha))) + F(z_{Y-y}(R_y(\alpha)))]$. We have been unable to fully characterize the sign of this term for all values of y. However, for the special case when y increases from 0 to 1 we can sign this term and be more definitive.

Theorem 3.4:

1. *If* $0 < \alpha \leq 0.5$ *then* $R_1(\alpha) < R_0(\alpha)$;
2. *There exists* $\overline{\alpha} \in (0.5, 1]$ *such that* $R_1(\alpha) < R_0(\alpha)$ *if* $\alpha \in (0.5, \overline{\alpha}]$

Part 1 of this theorem states that the optimal re-order point falls with increasing uncertainty if the CSL is less than 0.5. Part 2 indicates that when α > 0.5, but below a threshold, the re-order point *initially falls* with an increase in lead-time uncertainty. This contradicts the prediction for the normal approximation.

In the next section we numerically study the effect of decreasing lead time uncertainty on safety stocks for various lead time distributions.

4. Numerical Results and Analysis

Theorems 3.3 and 3.4 show that there is a range of cycle service levels above 50% where decreasing the lead time uncertainty increases the reorder point and safety stock when the lead time is uniformly distributed. In this section we present computational evidence to show that these claims are valid when lead times follow the gamma, uniform, or the normal distribution. For the gamma lead time distribution we show that for cycle service levels around 60 percent, decreasing lead time variability increases the reorder point. For the uniform lead time this effect is observed for cycle service levels close to, but above, 50 percentAs we have discussed earlier, most firms operate at cycle service levels in this range because they imply fill rates of around 98 percent. Using the computational results we also show that in this range, a manager is better off decreasing lead time rather than lead time variability if reducing inventories is their goal.

We first consider the effect of reducing lead time variability on reorder points and safety stock when the lead time follows the uniform or gamma distribution. In both case we try and keep the mean lead time fixed and vary the standard deviation. Demand per period is assumed to be normal with a mean μ=20 and standard deviation σ= 15 or 5. This allows us to analyze the effect for both high and low coefficient of variation of demand.

Figure 3 shows the effect of reducing lead time variability when periodic demand has a high coefficient of variation (15/20) and lead time is uniformly distributed with a mean of 10 and a range of 10 ± y, where y ranges from 0 to 10. We plot the change in reorder point as y changes using both the normal approximation and the exact calculations for cycle service levels of 0.5, 0.51, and 0.55. The safety stock is calculated as ss = ROP - 200 because 200 is the mean demand during the lead time. The curves marked Norm(.) represent the results of the normal approximation, whereas the others represent the exact calculation of reorder point.

From Figure 3 we conclude that for high coefficient of variation of periodic demand, if lead times are uniformly distributed, there is a range of cycle service levels above 0.5 (but close to 0.5), where reducing lead time uncertainty increases safety stocks, whereas the normal approximation predicts the opposite.

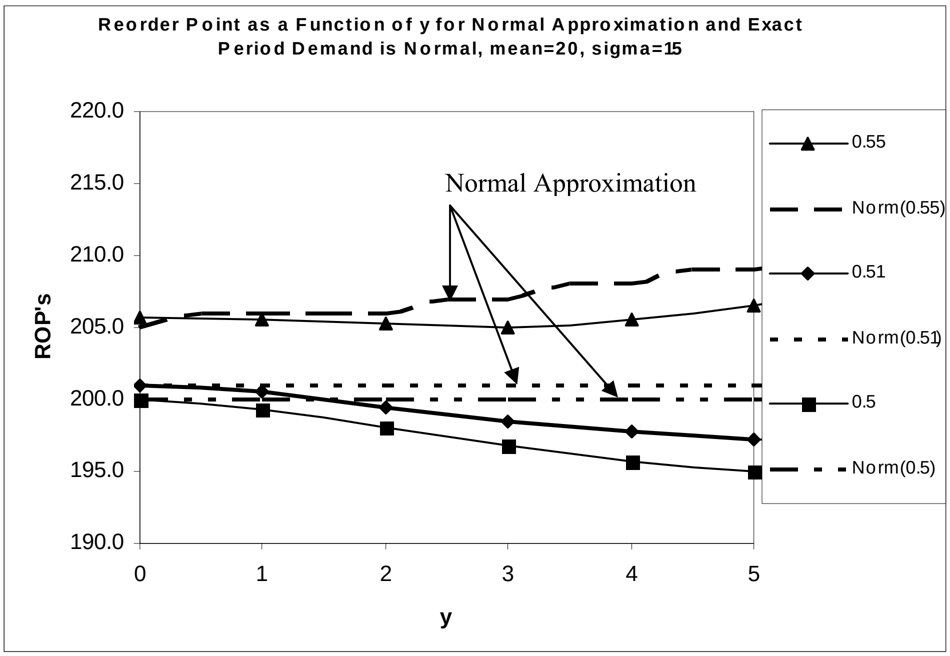

Figure 3: ROP as a Function of *y* for Uniform Lead Time

Figure 4 shows that the effect is even more pronounced when the lead time follows a gamma distribution. Once again we consider periodic demand to be normally distributed with a high coefficient of variation (15/20). Lead time is assumed to follow the gamma distribution with a mean of 10 and standard deviation varied from 6 to 1 and the reorder point calculated for cycle service levels of 0.50, 0.55, and 0.6. The plot compares the reorder points obtained using the normal approximation and the exact calculation.

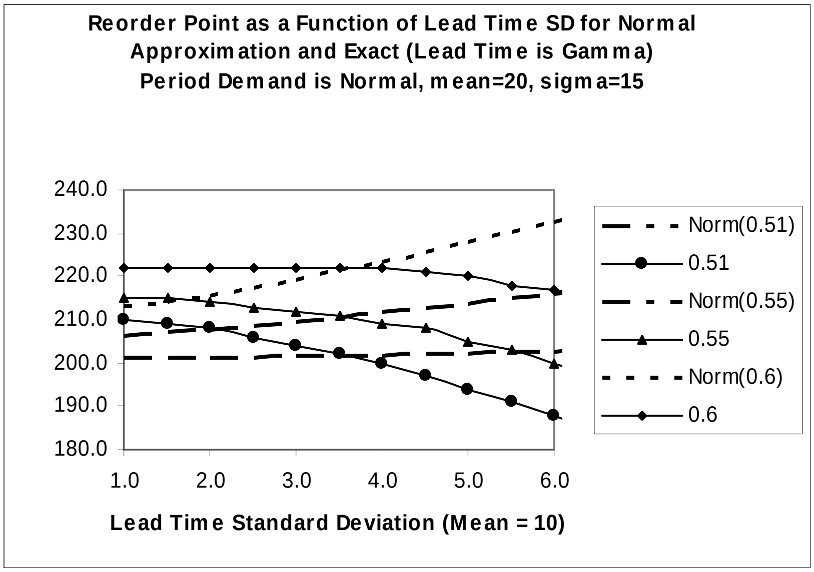

Figure 4: ROP as a Function of Lead Time Standard Deviation for Gamma Lead Time

Figure 4 shows that the normal approximation is even more erroneous when lead time follows the gamma distribution. Even for a cycle service level of 0.6, decreasing the lead time uncertainty increases safety stocks, whereas the normal approximation predicts just the opposite. When lead times are gamma distributed we thus conclude that there is a range of cycle service levels even beyond 0.6 when decreasing lead time variability increases the required safety stock.

Figure 5 repeats the results of Figures 3 and 4 but for low coefficient of variation (5/20) of periodic demand. Figure 5 shows that even with a low coefficient of variation of periodic demand, for cycle service levels between 50 percent and a threshold, the exact calculation shows that decreasing lead time variability increases the required safety stock, whereas the normal approximation predicts the opposite. The computational results show that the threshold value decreases as the coefficient of variation of periodic demand decreases. Thus, for very low coefficient of variation of period demand, the error of the normal approximation is less pronounced.

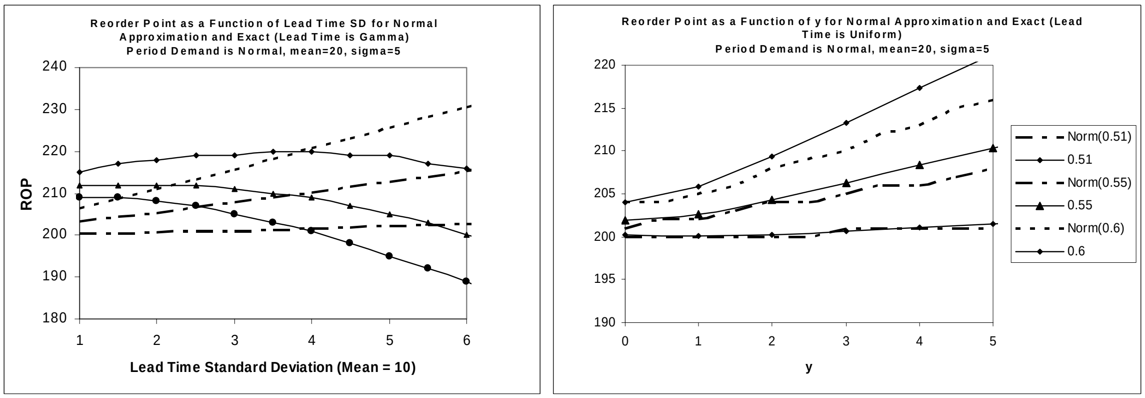

Figure 5: Reorder Point as a Function of Lead Time Uncertainty for Low c.v.

In Figures 6 and 7, we compare the impact of reducing lead time variability and lead time on safety stocks. In both cases we consider periodic demand to have a high coefficient of variation (15/20). In Figure 6 we consider lead time to be uniformly distributed with a mean of 10 and a range of 10 ± y. The chart on the left shows how the ROP changes as y is decreased from 5 to 0. The chart on the right shows how for $y = 5$, the ROP changes as the lead time decreases from 10 to 5. The results are shown for cycle service levels of 0.51, 0.55 and 0.6. In both cases the results show that the error of the normal approximation is less pronounce for low coefficient of variation of periodic demand.

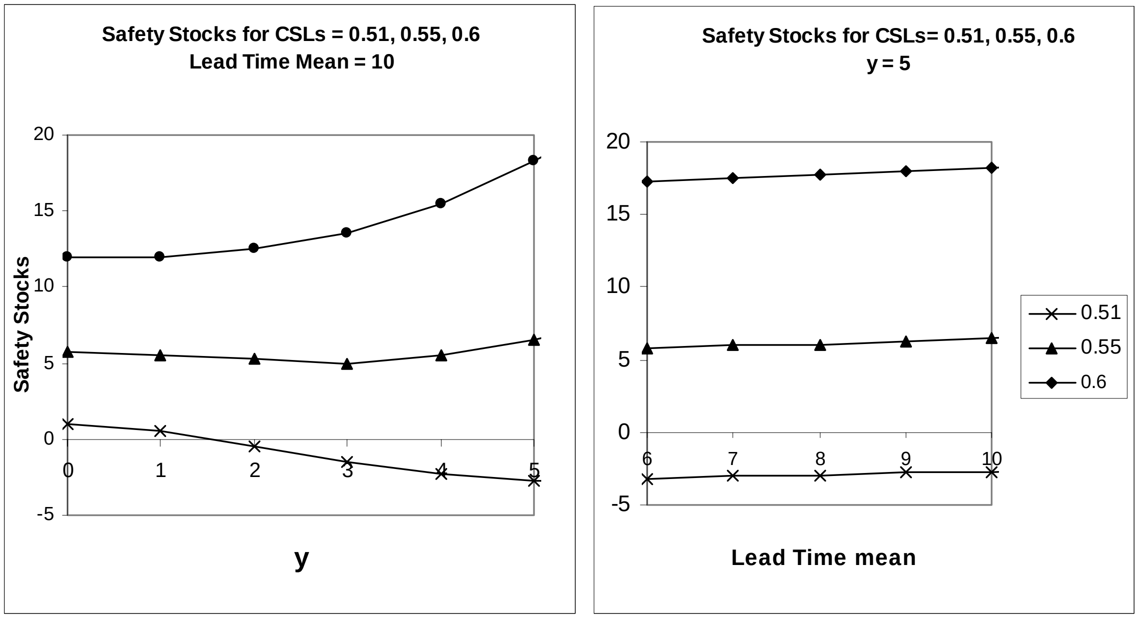

Figure 6: Safety Stock as a function of lead time uncertainty (left) and lead time mean (right) for Uniform Lead Times

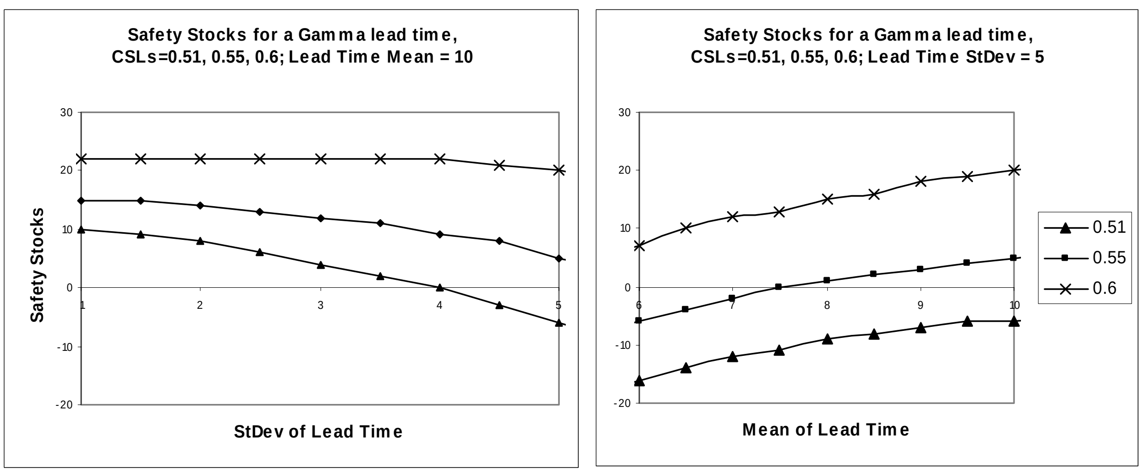

Figure 7: Safety Stock as a function of lead time uncertainty (left) and lead time mean (right) for Gamma Lead Times

5. Finding the Thresholds

We now show how the thresholds, below which the conclusions of the normal approximation are flawed, are obtained. These thresholds are obtained by studying the cumulative distribution function (CDF) of the demand during the lead-time. The CDF shows how the ROP changes as a function of the cycle service level. Recall that the safety stock *ss* = ROP - mean demand during lead time. In Figure 8 we represent two CDFs corresponding to the cases where the lead-time is uniformly distributed between $Y \pm y$ (represented by *y*) and $Y \pm (y+1)$ (represented by *y* + 1). The crossover $\overline{\alpha}$ represents the CSL for which both distributions require the same safety stock.

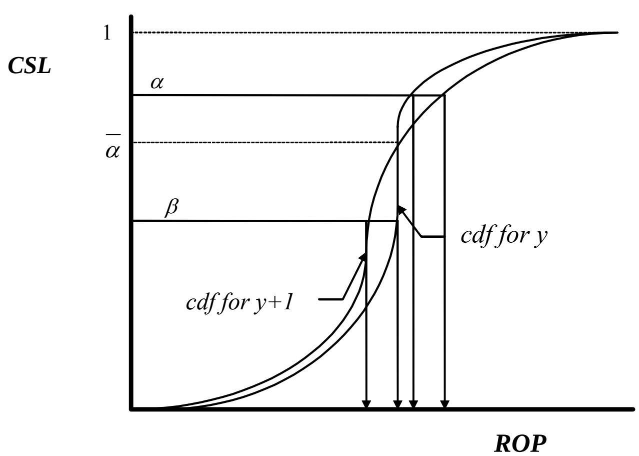

Figure 8: CDF for Demand During Lead time

For a cycle service level α, larger than the crossover point $\overline{\alpha}$, decreasing lead-time range from $Y \pm (y+1)$ to $Y \pm y$ results in a decrease in the safety stock. However, when the CSL is below $\overline{\alpha}$, say β, decreasing lead-time range from $Y \pm (y+1)$ to $Y \pm y$ results in an *increase* in the *ROP*. Thus, the *ROP* increases with a decrease in lead-time uncertainty for cycle service levels below the crossover $\overline{\alpha}$. The crossover point $\overline{\alpha}$ is the threshold below which decreasing the lead time variability increases the required safety stock. For cycle service levels between 50 percent and the crossover point, decreasing the lead-time uncertainty results in an increase of safety stock.

Next, we numerically describe the cumulative distribution functions for the case where the lead time distribution is uniform or Gamma and obtain crossover points to explain the results in Section 4. Figure 9 shows the CDF for the demand during the lead time when lead time is uniformly distributed between 10 ± *y* as *y* changes from 9 to 1. Observe that the crossover point between the CDF for *y* = 3 and *y* = 1 is at 0.564. This implies that for any cycle service level between 50 and 56.4 percent, decreasing *y* from 3 to 1 will increase the safety stock, whereas the normal approximation predicts otherwise. The crossover point thus establishes the threshold below which the normal approximation is directionally wrong when lead time is uniformly distributed.

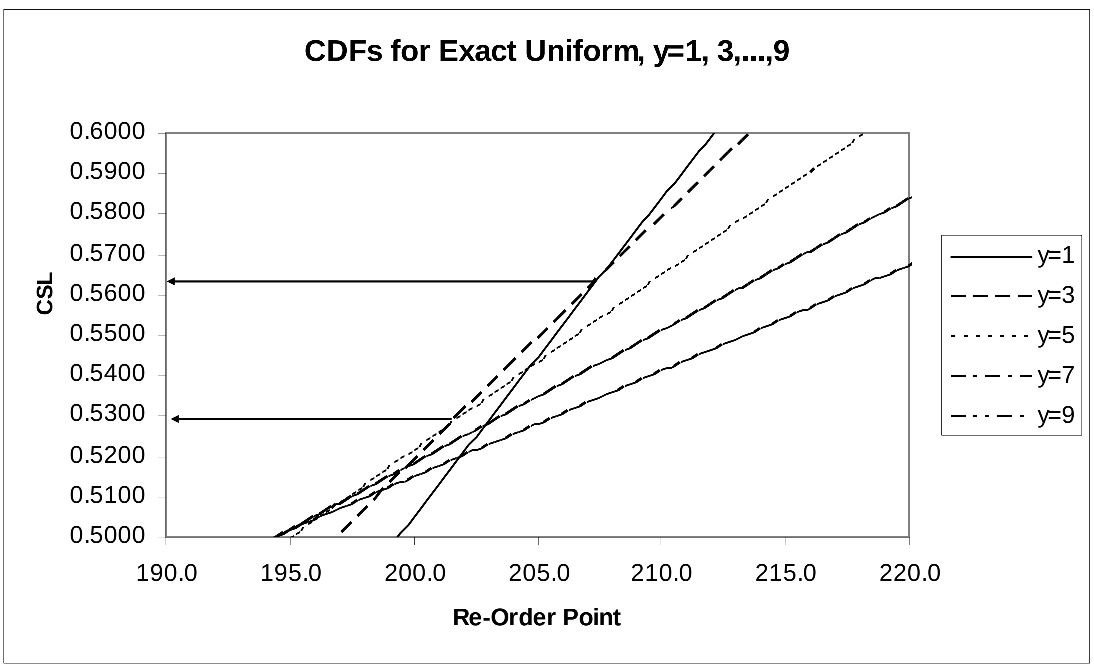

Figure 9: CDF of Lead Time Demand for different *y*'s when μ=20, σ=15 and L=10

Figure 10 shows how the reorder point ROP varies with the cycle service level when lead time follows a gamma distribution with a mean of 10 and a standard deviation that varies from 9 to 1. Once again the cross over point helps explain the range of cycle service levels over which the conclusions of the normal approximation are directionally flawed. For example, the crossover point for the CDF for a standard deviation of 5 and 3 is at 0.68. Thus, for cycle service levels between 50 and 68 percent, decreasing the standard deviation of lead time from 5 to 3 increases the safety stock required, whereas the normal approximation predicts just the opposite.

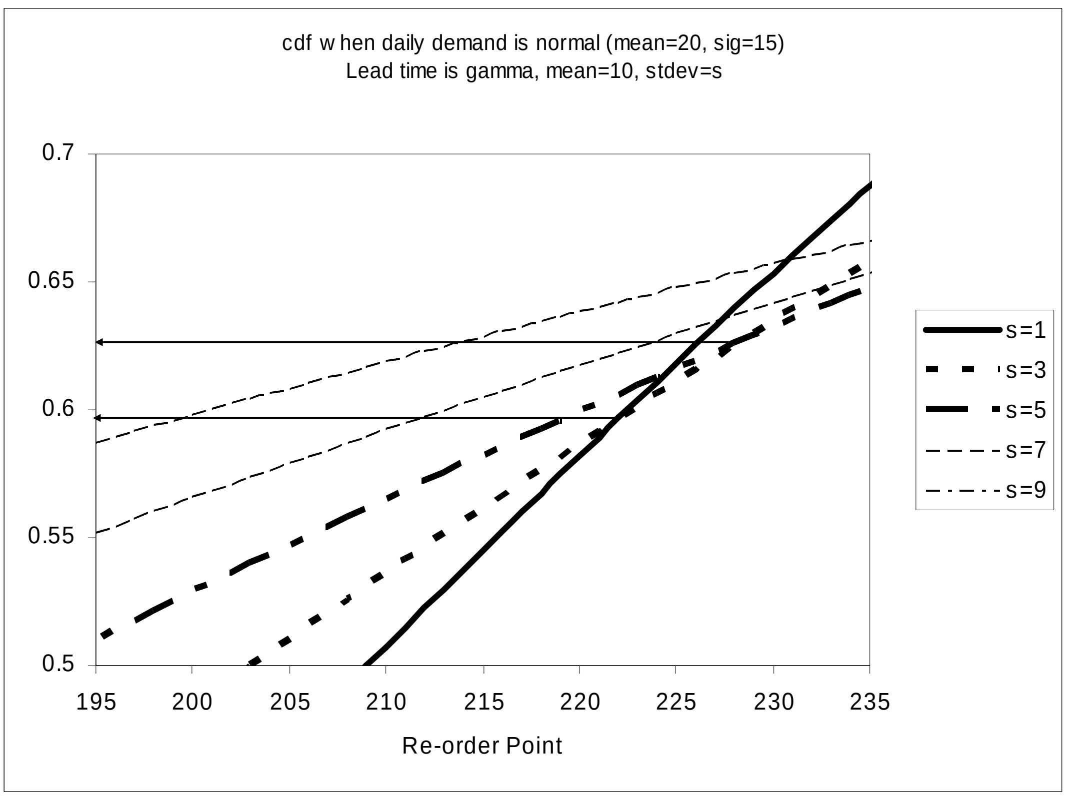

Figure 10: CDF of Lead Time Demand for Gamma lead time, μ=20, σ=15 and L=10

Figure 11 shows how the reorder point ROP varies with the cycle service level when lead time follows a truncated normal distribution (with mean of 10 and a discrete support from 0 to 20 days). Here we note that the cross over points are closer to 0.5 (0.54 when s=1 and s=3 intersect, and 0.51 when s=3 and s=5 intersect) illustrating that the width of the interval of service levels in which the re-order point decrease as variability increases is dependent on the shape of the lead time distribution.

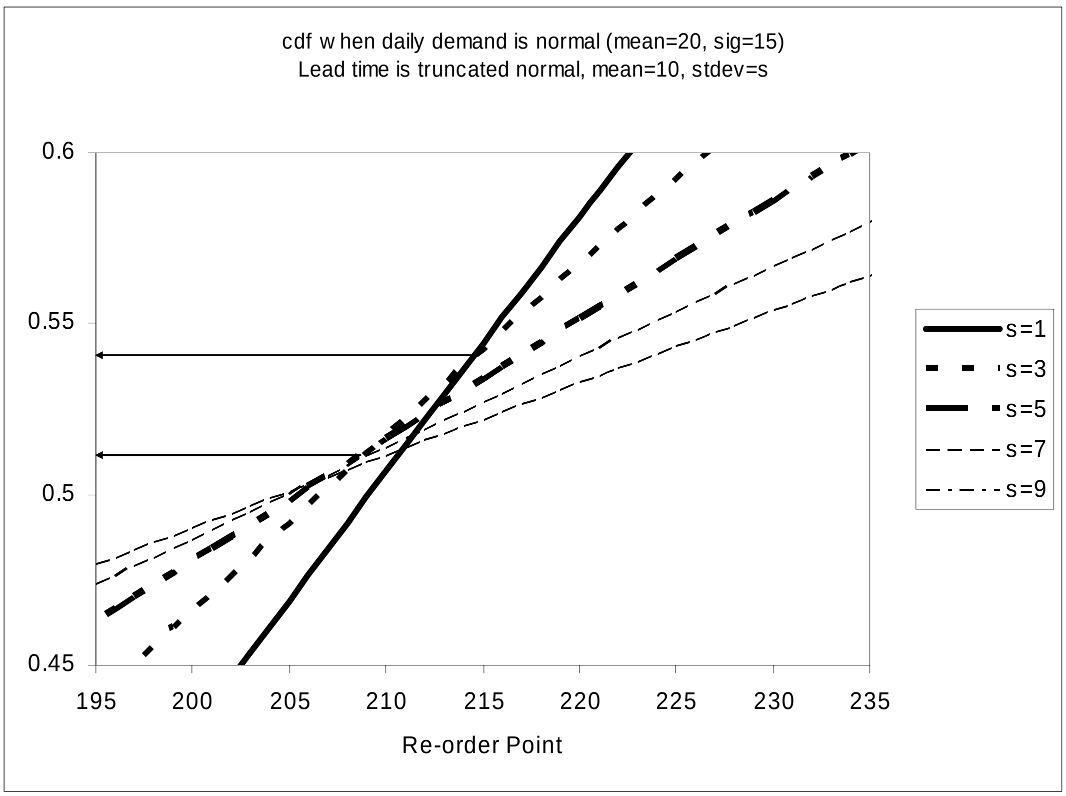

Figure 11: CDF of Lead Time Demand for Normal lead time, μ=20, σ=15 and L=10

The procedure detailed above for uniform, gamma, and normal lead times can be used for any lead time distribution to estimate the threshold below which decreasing lead time variability increases the safety stock required. To identify whether decreasing the lead time variability from σ_h to σ_l will increase or decrease the safety stock required, we plot the CDF of demand during the lead time for each lead time variability and identify the crossover point *X*. Cycle service levels between 50 percent and the crossover point represent the range over which the exact distribution predicts an increase in safety stock if lead time variability is decreased, whereas the normal approximation predicts the opposite.

6. Conclusion

For the most part, management's understanding of the effect on safety stocks of uncertainty in lead-time is based on an approximate characterization of demand during lead time using the normal distribution. For cycle service levels above 50 percent the normal approximation predicts that a manager can reduce safety stocks by decreasing lead time uncertainty. Our analytical results and numerical experiments, however, indicate that for cycle service levels between 50 percent and a threshold, the prescriptions of the normal approximation are flawed and decreasing the lead time uncertainty, in fact, *increases* the required safety stock. In this range of cycle service levels, a manager who wants to decrease inventories should focus on decreasing lead times rather than lead time variability. This contradicts the conclusion drawn using the normal approximation.

Our conclusion is more pronounced when demand has a high coefficient of variation. When the lead time follows a gamma distribution, the prescriptions of the normal approximation are flawed over a wide range of cycle service levels. This range is narrower when lead times are uniformly or normally distributed. Thus, using the normal approximation makes sense if lead times are normally distributed but would not make sense if lead times follow a distribution closer to the gamma.

Appendix: Proof of theorem 3.3

Theorem 3.3: ***For α = 0.5, the re-order point $R_y(0.5)$ declines with an increase in lead-time uncertainty y, i.e., $R_{y+1}(0.5) < R_y(0.5)$.***

Proof: The result is proved using induction. We first consider the case for $y = 0$, i.e., the lead time is fixed at Y. For a fixed lead time Y, the reorder point for a cycle service level of 0.5 is given by $R_o(0.5) = Y\mu_x$.

Now consider the lead time to be uniformly distributed with equal support on $\{Y+1, Y, Y-1\}$, i.e., $y = 1$. If the re-order point is kept at $R_0(0.5) = Y\mu_x$, the cycle service level is given by

$$\frac{1}{3}\left\{F\left(\frac{(Y-(Y-1))\mu_x}{\sigma_x\sqrt{Y-1}}\right)+F\left(\frac{(Y-(Y+1))\mu_x}{\sigma_x\sqrt{Y+1}}\right)+F\left(\frac{(Y-Y)\mu_x}{\sigma_x\sqrt{Y}}\right)\right\}$$

To lighten notation, we let $c = \sigma_x/\mu_x$. We claim that $F\left(\frac{1}{c\sqrt{Y-1}}\right)+F\left(\frac{-1}{c\sqrt{Y+1}}\right)>1$. By Lemma 3.2, this follows because $\frac{-1}{c\sqrt{Y+1}}<0<\frac{1}{c\sqrt{Y-1}}$ and $\frac{1}{c\sqrt{Y+1}}<\frac{1}{c\sqrt{Y-1}}$. This implies that if $y = 1$, the cycle service level for a reorder point of $R_0(0.5)$ is strictly greater than 0.5.Thus, $R_1(0.5) < R_0(0.5)$. Define $\Delta_1 = (R_0(0.5) - R_1(0.5)) / \mu_x$. The service level at $R_1(0.5)$ is 0.5 and is given by

$$\frac{1}{3}\left\{F\left(\frac{(Y-(Y-1))-\Delta_1}{c\sqrt{Y-1}}\right)+F\left(\frac{(Y-(Y+1))-\Delta_1}{c\sqrt{Y+1}}\right)+F\left(\frac{(Y-Y)-\Delta_1}{\sqrt{Y}}\right)\right\}=0.5. \qquad \text{(A1)}$$

Since $\Delta_1 > 0$, we have $F\left(\frac{-\Delta_1}{c\sqrt{Y}}\right)<0.5$.Thus, it must be the case that

$F\left(\frac{1-\Delta_1}{c\sqrt{Y-1}}\right)+F\left(\frac{-1-\Delta_1}{c\sqrt{Y+1}}\right)>1$ or by Lemma 3.2

$$\frac{1-\Delta_1}{c\sqrt{Y-1}}>\frac{1+\Delta_1}{c\sqrt{Y+1}} \qquad \text{(A2)}$$

Now consider raising the lead time uncertainty by assuming lead time to be uniformly distributed over $\{Y-2, Y-1, Y, Y+1, Y+2\}$. If the re-order point is kept at $R_1(0.5) = Y\mu_x - \Delta_1\mu_x$, the cycle service level is given by

$$\frac{1}{5}\left\{F\left(\frac{(Y-(Y-2))-\Delta_1}{c\sqrt{Y-2}}\right)+F\left(\frac{(Y-(Y+2))-\Delta_1}{c\sqrt{Y+2}}\right)+F\left(\frac{(Y-(Y-1))-\Delta_1}{c\sqrt{Y-1}}\right)+F\left(\frac{(Y-(Y+1))-\Delta_1}{c\sqrt{Y+1}}\right)+F\left(\frac{(Y-Y)-\Delta_1}{\sqrt{Y}}\right)\right\}$$

We now claim that $F\left(\frac{2-\Delta_1}{c\sqrt{Y-2}}\right)+F\left(\frac{-2-\Delta_1}{c\sqrt{Y+2}}\right)>1$

By lemma 3.2, this is equivalent to showing that

$$\frac{2-\Delta_1}{c\sqrt{Y-2}}=\frac{1}{c\sqrt{Y-2}}+\frac{1-\Delta_1}{c\sqrt{Y-2}}>\frac{2+\Delta_1}{c\sqrt{Y+2}}=\frac{1}{c\sqrt{Y+2}}+\frac{1+\Delta_1}{c\sqrt{Y+2}} \qquad \text{(A3)}$$

From (A2), we have
$\frac{1-\Delta_1}{c\sqrt{Y-2}}>\frac{1-\Delta_1}{c\sqrt{Y-1}}>\frac{1+\Delta_1}{c\sqrt{Y+1}}>\frac{1+\Delta_1}{c\sqrt{Y+2}}$.
Given the fact that $\frac{1}{c\sqrt{Y-2}}>\frac{1}{c\sqrt{Y+2}}$, (A3) thus follows. This implies that if $y = 2$, the cycle service level for a reorder point of $R_1(0.5)$ is strictly greater than 0.5. Thus, $R_2(0.5) < R_1(0.5)$. Define $\Delta_2 = (R_1(0.5) - R_2(0.5)) / \mu_x$.

We now use induction to complete the proof. Define $\Delta_y = (R_{y-1}(0.5) - R_y(0.5)) / \mu_x$. To show that $R_{y+1}(0.5) < R_y(0.5)$ we start with the induction assumption that $\frac{y-\Delta_y}{c\sqrt{Y-y}}>\frac{y+\Delta_y}{c\sqrt{Y+y}}$. We now need to prove that $F\left(\frac{y+1-\Delta_y}{c\sqrt{Y-(y+1)}}\right)+F\left(\frac{-(y+1)-\Delta_y}{c\sqrt{Y+(y+1)}}\right)>1$.

Observe that

$$\frac{(y+1)-\Delta_y}{c\sqrt{Y-(y+1)}}=$$

$$\frac{y}{c\sqrt{Y-(y+1)}}+\frac{1-\Delta_y}{c\sqrt{Y-(y+1)}}>$$

$$\frac{y}{c\sqrt{Y-(y+1)}}+\frac{1-\Delta_y}{c\sqrt{Y-y}}>$$

$$\frac{y}{c\sqrt{Y+(y+1)}}+\frac{1+\Delta_y}{c\sqrt{Y+y}}> \qquad \text{\{This follows from the induction hypothesis\}}$$

$$\frac{y}{c\sqrt{Y+(y+1)}}+\frac{1+\Delta_y}{c\sqrt{Y+(y+1)}}=$$

$$\frac{(y+1)+\Delta_y}{c\sqrt{Y+(y+1)}}$$

The result thus follows using Lemma 3.2. This implies that $R_{y+1}(0.5) < R_y(0.5)$.

Appendix: Technical Methodology

The Normal Approximation for the Exact Uniform Distribution

$ROP_N = F^{-1}$ {α, $\mu_y \mu_x$, $\sqrt{\mu_y \sigma_X^2 + y(y+1)\mu_X^2/3}$} where F^{-1} {•, •, •} is the inverse of the normal distribution (NORMINV) of given mean and standard deviation.

The Normal Approximation for the Gamma Distribution

$ROP_N = F^{-1}$ {α, $L\mu_x$, $\sqrt{L\sigma_X^2 + {s_L}^2\mu_X^2}$} where F^{-1} {•, •, •} is the inverse of the normal distribution of given mean and standard deviation and L and s_L are the mean and standard deviation of the Gamma distribution.

The Exact Uniform

Sheet 1. Generate Table indexed by (row) ROP and (column) y with CSL in the body of the Table.

Sheet 2. Use VLOOKUP function to extract ROP index which corresponds to given y and CSL

The Discrete Gamma Distribution

We seek the inverse (G^{-1}) of the cumulative distribution function of the demand during lead time. The lead time distribution has mean L and standard deviation s_L. We know

$ROP = G^{-1}\{P(D < X) = \sum_{l=1,\ldots,30} w_l * NORMDIST(X, l*\mu_X, \sigma_X*sqrt(l),1)\}$

Sheet 1. Generate weights (w_l , l=1,…,30) Table indexed by (row) support (between 0 and 30) and (column) standard deviation with, in row j and column s the value

$$GAMMADIST\left(j,\left(\frac{L}{s}\right)^2,\frac{s^2}{L},1\right)-\sum_{i=0}^{j-1}GAMMADIST\left(i,\left(\frac{L}{s}\right)^2,\frac{s^2}{L},1\right) \text{ for}$$

j=1,…,29 with 0 in row j=0 and, in row j=30,

$$GAMMADIST\left(30,\left(\frac{L}{s}\right)^2,\frac{s^2}{L},1\right)-\sum_{i=0}^{29}GAMMADIST\left(i,\left(\frac{L}{s}\right)^2,\frac{s^2}{L},1\right)+$$

$$1-\sum_{i=0}^{30}GAMMADIST\left(i,\left(\frac{L}{s}\right)^2,\frac{s^2}{L},1\right)$$

That is, we add to j=30 the mass of the tail to the right of 30.
In Sheet 1, we also generate a table of NORMDIST values as per the ROP formula above.

Sheet 2. Matrix multiply the Sheet 1's Normdist table to Sheet 1's weights table.

Sheet 3: Use a VLOOKUP(CSL) on Sheet 2 to find the ROP which yields the given CSL.

The resulting discrete gamma distributions for lead time are illustrated in Figure A1.

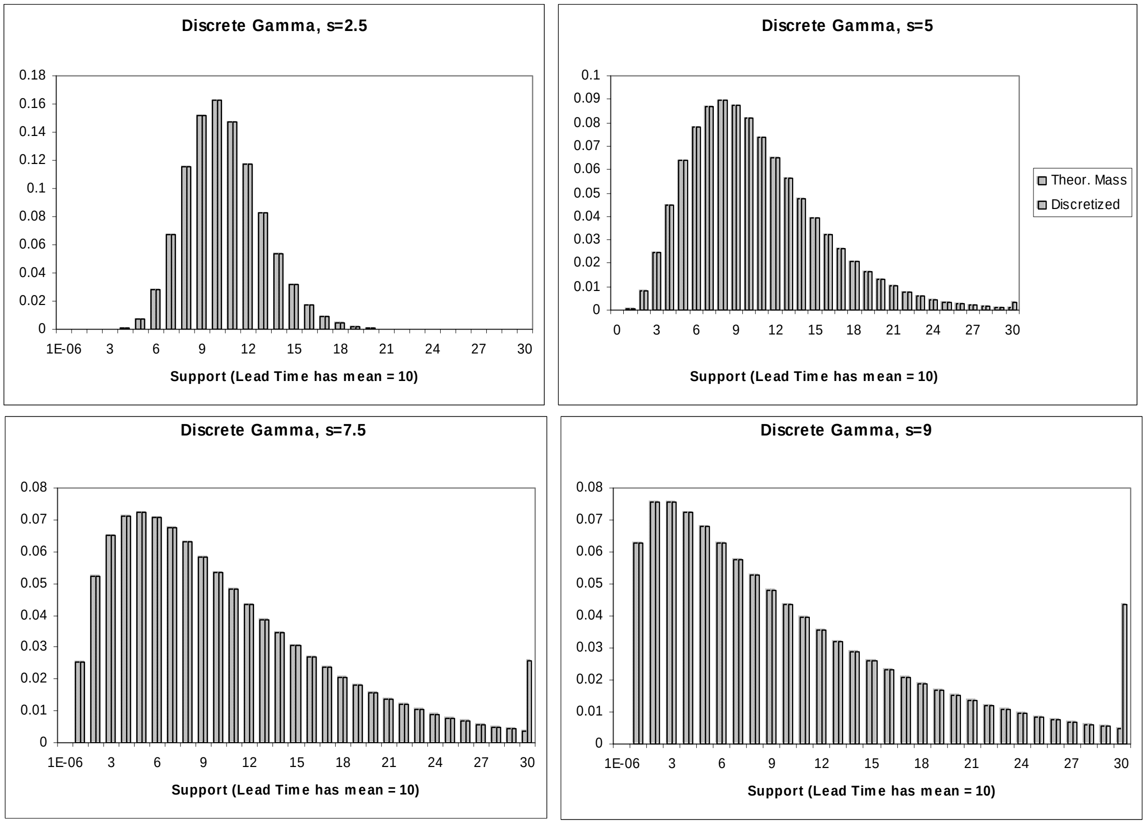

Figure A1: Gamma Lead Time Distributions for standard deviations of 2.5, 5, 7.5, and 9.

The Truncated Normal Distribution

We seek the inverse (G^{-1}) of the cumulative distribution function of the demand during lead time. The lead time distribution has mean L and standard deviation s_L. We know

$ROP = G^{-1}\{P(D < X) = \sum_{l = 1,\ldots,L} w_l * NORMDIST(X, l*\mu_X, \sigma_X*sqrt(l),1)\}$

Sheet 1. Generate weights Table indexed by (row) support (between 0 and 30) and (column) standard deviation with, in row j and column s the value

$NORMDIST(j,L,s,1)-\sum_{i=0}^{j-1} NORMDIST(i,L,s,1)$ for j=1,…,29 with

$NORMDIST(0,L,s,1)$ 0 in row j=0 and, in row j=30,

$$NORMDIST(30,L,s,1)-\sum_{i=0}^{29} NORMDIST\left(i,\left(\frac{L}{s}\right)^2,\frac{s^2}{L},1\right)+$$

$$1-\sum_{i=0}^{30} NORMDIST\left(i,\left(\frac{L}{s}\right)^2,\frac{s^2}{L},1\right)$$

That is, we add to j=0 and j=30, respectively, the mass of the tail to the left of 0 and to the right of 30. In Sheet 1, we also generate a table of NORMDIST values as per the ROP formula above.

Sheet 2. Matrix multiply the Sheet 1's Normdist table to Sheet 1's weights table.

Sheet 3: Use a VLOOKUP(CSL) on Sheet 2 to find the ROP which yields the given CSL.

The resulting truncated normal distributions for lead time are illustrated in Figure A2.

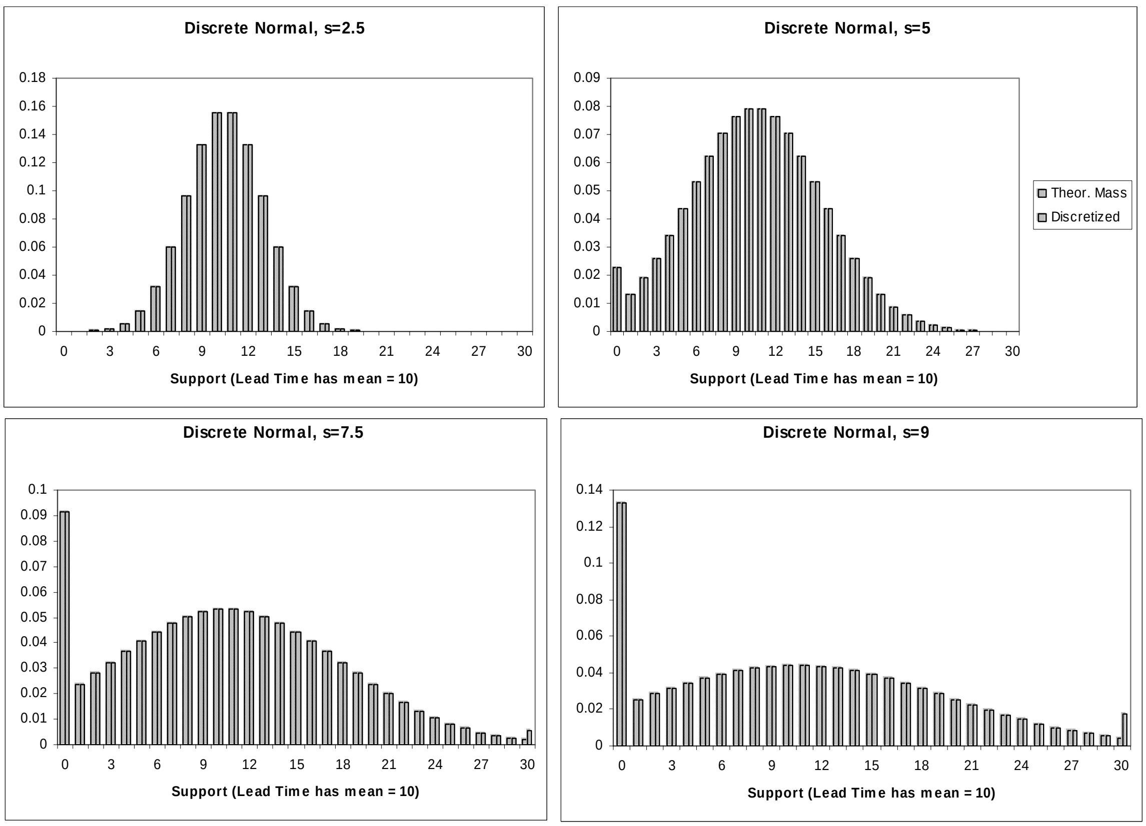

Figure A2: Normal Lead Time Distributions for standard deviations of 2.5, 5, 7.5, and 9.

References

Aiginger, K., 1987. Production and Decision Theory under Uncertainty, Blackwell, London

Byrne, P.M. and W.J. Markham, 1991. Improving Quality and Productivity in the Logistics Process, The Council of Logistics Management, Oak Brooks, IL.

Chopra, S. and P. Meindl, 2003. Supply Chain Management: Strategy, Planning, and Operations, 2nd Edition, Prentice-Hall, New York.

Eppen, G.D. and R.K. Martin, 1988. Determining Safety Stock In The Presence Of Stochastic Lead Time And Demand, Management Science, 34(11).

Hadley, G. and T.M. Whitin, 1963. Analysis of Inventory Systems. Prentice-Hall, Englewood Cliffs, NJ.

Gerchak, Y. and D. Mossman, 1992. On The Effect Of Demand Randomness On Inventories And Costs, Operations Research, 40(4)

Lee, H.L. and C. Billington, 1992. Managing Supply Chain Inventory: Pitfalls and Opportunities, Sloan Management Review 33, 65-73.

Silver, E.A. and R. Peterson, 1985. Decision Systems for Inventory Management and Production Planning, Wiley, New York, NY.

Song, J.-S., 1994. The Effect of Leadtime Uncertainty in a Simple Stochastic Inventory Model, Management Science 40(5).

Tyworth, J.E. and L. O'Neill, 1997. Robustness of the Normal Approximation of Lead-Time Demand in a Distribution Setting, Naval Research Logistics 44, 165-186.

Tyworth, J.E., Guo, Y. and R. Ganeshan, 1996. Inventory Control under Gamma Demand and Random Lead Times, Journal of Business Logistics, 17(1).