COMPARING TWO ALTERNATIVE METHODS OF DETERMINING SAFETY STOCK LEVELS: THE DEMAND AND THE FORECAST SYSTEMS

by

Walter Zinn
University of Miami (Florida)

and

Howard Marmorstein
University of Miami (Florida)

HEADNOTE

Two competing methods of determining safety stock levels are discussed in the logistics literature. The primary difference between the two methods is that one bases the computation of safety stock on the variance of demand, while the other utilizes the variance of demand forecast errors. Accordingly, this paper presents a simulation which quantifies the difference in safety stock required under the two methods.

INTRODUCTION

An important function of logistics management is to maintain an appropriate level of safety stock. When insufficient safety stock is held, the firm fails to meet its customer service objective. Conversely, inventory carrying costs are too high when excessive safety stock is maintained. This paper presents a simulation comparing two competing methods of setting the level of safety stock. The simulation enables us to estimate and explain the difference in safety stock required by the two methods.

In the first method, referred to as the Demand System, the level of safety stock depends upon the *variability of demand*. In the second method, called the Forecast System, the required level of safety stock depends upon the *variability of demand forecast errors*.

The difference between the two systems is substantial. Simulation results indicate that the Forecast System, albeit less prominent in the logistics literature, typically requires about 15% less safety stock to provide the same level of customer service.[1] Interestingly, these savings vary from almost zero to as much as 70% for the products included in this simulation. The simulation also identifies the variables which have the greatest impact on savings. In order of importance they are demand variability, average demand, forecast quality, and lead time variability. Finally, a counterintuitive relationship was uncovered in the simulation. Specifically, the reduction in safety stock achieved by switching to the Forecast System is *negatively* related to the variability of the lead time of orders from suppliers.

The paper begins by presenting a detailed explanation of how safety stock is determined under each system. This is followed by the simulation which quantifies the difference in safety stock required under the two systems. The managerial implications of the simulation results are discussed in the final section.

THE DEMAND AND THE FORECAST SYSTEMS

The two methods of setting safety stock that have received the most attention are the Demand System (DS) and the Forecast System (FS). Logistics textbooks typically emphasize the DS. The DS has also been utilized in articles by Buffa and Reynolds,[2] Ronen,[3] Howard,[4] Mentzer and Krishnan,[5] Blumenfeld et. al.,[6] and Buffa.[7] The two major advantages of the DS are its relatively low cost and ease of application.

The FS is presented in Brown.[8] It has been less frequently employed than the DS. Mohanty and Chandrashekhar[9] adopted a variation of the FS wherein the safety stock is a product of the sales forecast and an empirically determined constant. An approximation of the FS is presented in Dror and Trudeau.[10] Another approximation of the FS is discussed in Mentzer and Krishnan.[11] Herron[12] incorporated the FS in a model which improves the cost-effectiveness of providing a predetermined aggregate level of customer service.

Under the DS, the safety stock needed to absorb demand and lead time uncertainties is computed as follows:[13]

$$SS_d = k\sqrt{t.s_d^2 + d^2.s_t^2} \qquad (1)$$

where:

SS_d	=	safety stock under the DS (units)
k	=	a dimensionless safety factor relating the investment in safety stock to measures of customer service
t	=	average lead time in periods (e.g. days or weeks)
S_d	=	standard deviation of past demand
S_t	=	standard deviation of lead time
d	=	average demand per period (e.g. days or weeks)

Under the FS, safety stock is determined in a similar fashion:

$$SS_f = k\sqrt{t.s_f^2 = d^2.s_t^2} \qquad (2)$$

where:

SS_f	=	safety stock under the FS (units)
S_f	=	standard deviation of the forecast errors

Computationally, the only difference is the substitution of S_f (standard deviation of the forecast errors) for S_d (standard deviation of demand). The origin and implications of this difference are as follows.

Under the DS the implicit forecast for each period is the average demand. As a result, the standard deviation of demand equals the standard deviation of the "forecast" errors. Use of the FS is more complex and costly than the DS but potentially more efficient.[14] The FS entails using a model-based approach (e.g., regression or time series) to forecast the demand for each upcoming period. To the extent that these forecasts are more accurate than the naive model implicit in the DS, S_f is less than S_d and the required safety stock is reduced commensurately. (See Equations 1 and 2 above.)

As long as the forecast errors are approximately normally distributed, the probability of a stockout is a function of the number of standard deviations of the forecast error held in safety stock. The fact that the standard deviation of the forecast errors is smaller than the standard deviation of average demand does not impact customer service. This happens because the base stock in the FS is adjusted in accord with the forecast for each period. Note that the *average base stock* is the same under the two systems, because the average forecast error equals zero.

With this background in mind, the paper now examines the factors which affect the potential for reducing the safety stock associated with a given product. The simulation is presented in order to quantify these relationships.

THE SIMULATION

The goals of the simulation are:

1. to quantify the safety stock savings for a broad range of products that can be expected by a firm switching from the DS to the FS.

2. to determine the extent to which the reduction in safety stock achieved by adopting the FS varies by product.

3. to identify the product-specific variables that explain the variation in savings.

The variables included in the simulation are presented first. This is followed by a multiple regression analysis which identifies the independent variables that have the greatest effect on savings in safety stock. A sensitivity analysis is presented next to examine the extent to which the results depend upon the specific range of input values employed. Finally, an interpretation of the results is presented and the managerial implications are discussed.

Simulation Input

The dependent variable in the simulation is "unit savings" (US), defined as safety stock under the DS (SS_d) minus safety stock under the FS (SS_f). In the discussion of managerial implications "percent savings" (PS) -- defined as US/SS_d -- is also considered.

The independent variables are the factors affecting safety stock under each of the two systems. The relationships between the independent and dependent variables are based on Equations 1 and 2. Table 1 displays the independent variables in the simulation and the specific input values which they assume. Additional relationships among variables in the simulation are presented in Table 2.

TABLE 1
SIMULATION INPUT

Variable	Definition	Input Values
CV_t	lead time variability	.2, .5, .8
t	average lead time	1, 3, 5 weeks
CV_d	demand variability	.2, .5, .8
d	average demand per period	100, 300, 500 units
F	forecast quality	.2, .5, .8
k	customer service safety factor	1.28, 1.64, 2.33

TABLE 2
ADDITIONAL RELATIONSHIPS

$$s_d = CV_d * d$$
$$s_t = CV_t * t$$
$$s_f = s_d * (1\text{-}F)$$

Four of the independent variables and corresponding input values require additional clarification. The forecast quality (F) input is related to the portion of the demand variance explained by the demand forecast. (See last equation in Table 2.) A forecast quality level of .5, for instance, means that the standard deviation of the forecast error is equal to half of the standard deviation of demand.

The customer service safety factor (k) is defined in terms of the probability of a stockout during a lead time. Under the assumption that demand is distributed normally, a safety factor equal to 2.33 standard deviations of demand results in a customer service level of 99%. The input values appearing in Table 1 are the safety factors corresponding to customer service levels of 90%, 95%, and 99% respectively.

Finally, variability of both demand and lead time are operationalized in terms of the coefficient of variation of each of these variables. This is preferred over using the standard deviation of demand (or lead time) for two reasons:

1. It preserves the orthogonality of the independent variables in the simulation. Specifically, demand (lead time) is orthogonal to variability of demand (variability of lead time). This facilitates an unambiguous interpretation of the beta weights obtained in the regression analysis.

2. It preserves the validity of the results by incorporating ecologically valid relationships among the inputs. Specifically, previous research indicates that the variance of demand (or lead time) is strongly positively correlated with the absolute level of demand (or lead time).[15] The input to this simulation preserves these relationships.

To analyze the impact of the individual independent variables, a data matrix was created by combining all possible input values. There are six independent variables with three levels each, for a total of 729 observations. A multiple regression analysis is used to assess the relative importance of the independent variables. To determine whether the results are very sensitive to the input values chosen for this simulation, a sensitivity analysis was then performed using a wider range for each input. Both the multiple regression and the sensitivity analysis are briefly presented before interpreting the results.

Multiple Regression Analysis

Results from the multiple regression using unit savings as the dependent variable are presented in Table 3. Each of the independent variables is significantly related to unit savings, albeit with differing levels of importance. The analysis of variance indicates the significance of the full model -- the independent variables collectively explain 64% of the variance in unit savings ($p<.001$). The effect of each independent variable is discussed separately after the sensitivity analysis is presented.

TABLE 3
MULTIPLE REGRESSION ANALYSIS

Dependent Variable: Unit Savings

Independent Variable	Parameter Estimate	Standardized (b=0)	t-value Level	Probability
Intercept	-286.8	0.0000	-15.60	0.0000
Lead time Variability (CV_t)	-147.8	-0.2632	-11.88	0.0000
Lead time (t)	8.8917	0.1055	4.76	0.0000
Demand variability (CV_d)	296.19	0.5272	23.80	0.0000
Demand (d)	.34170	0.4055	18.30	0.0000
Forecast quality (F)	167.04	0.2973	13.42	0.0000
Service level (k)	58.579	0.1855	8.37	0.0000

Analysis of Variance

Source	df	Sum of Squares	Mean Square	F-Ratio	Probability Level
Constant	1	7660965	7660965		
Model	6	8912317	1485386	219.18	0.000
Error	722	4892914	6776.89		
Total	728	1.380E+07	18963.2		
R Squared			0.6456		
Adjusted R Squared			0.6426		

Sensitivity Analysis

The purpose of the sensitivity analysis is to determine whether the results of the multiple regression analysis are significantly affected if a different range is chosen for the simulation inputs. Accordingly, the simulation and the multiple regression analysis were repeated six times -- once for each independent variable. In each iteration the range of a single independent variable was enlarged. This wider range was obtained by increasing the highest input value by 50% and reducing the lowest input value by 50%.[16] (See Table 4 below.)

TABLE 4
WIDER INPUT RANGES FOR THE SENSITIVITY ANALYSIS

Input	Definition	Original Range	Wider Range
CV_t	lead time variability	.2, .5, .8	.1, .5, 1.2
t	average lead time	1, 3, 5	.5, 3, 7.5
CV_d	demand variability	.2, .5, .8	.1, .5, 1.2
d	average demand	100, 300, 500	50, 300, 750
F	forecast quality	.2, .5, .8	.1, .5, .9
k	customer service	1.28, 1.64, 2.33	0.67, 1.64, 2.33

Table 5 presents the results of the sensitivity analysis. The first column in the table contains the standardized estimates obtained with the original inputs. Succeeding columns contain the standardized estimates for all of the independent variables when only the input indicated in the column heading is changed.

Results from the sensitivity analysis are favorable. As expected, when the input range of each variable is raised, the standardized estimate of that variable increases. The following two checks for sensitivity pertain to the estimates of the remaining predictors in each equation.

TABLE 5
SENSITIVITY ANALYSIS WITH UNIT SAVINGS AS THE DEPENDENT VARIABLE: STANDARDIZED ESTIMATES

	Original	Wider Input Range*					
Input	Range	CV_t	t	CV_d	d	F	k
CV_t	-.263	-.350	-.255‘	-.169	-.223	-.243	-.240
t	.105	.122	.162	.111	.089	.093	.096
CV_d	.527	.448	.504	.609	.446	.480	.482
d	.405	.369	.387	.339	.498	.368	.371
F	.297	.285	.283	.258	.251	.372	.272
k	.185	.169	.177	.155	.157	.168	.300

*all estimates are significant for $p<.001$

First, the estimates of the remaining variables showed very little change in their relative magnitudes as the target input range was increased. That is, the independent variables that were most important in the original simulation remained so in the sensitivity analysis. Second, there were no changes in the directionality of the relationships. Variables that were directly (inversely) related to unit savings in the original simulation remained so in the sensitivity analysis. Together, these results suggest that the simulation results are reasonably insensitive to the specific input ranges and merit further interpretation.

Interpretation of Results

Let us now systematically examine the effect of each independent variable on unit savings. Results provide the strength as well as the directionality of each independent variable with respect to unit savings. Although the directionality is predictable for some variables, some surprising results were also found.

Recall that unit savings is defined as the unit reduction in safety stock obtained by changing from the DS to the FS, while holding customer service constant. The variables are discussed in the order of their importance in terms of their impact on unit savings.

Variability of Demand. The effect of the variability of demand on unit savings is intuitive and easily interpretable. Clearly, the greater the variability of demand, the greater the opportunity to reduce safety stock by forecasting demand more precisely. Accordingly, Table 3 shows a very strong, positive relationship between demand variability and unit savings ($t = +23.80$; $p < .001$).

Demand. One is also unsurprised to find that the absolute level of demand is strongly related to the potential for savings in safety stock ($t = +18.30$; $p < .001$). (See Table 3.) As the number of units demanded increase, the potential for unit savings should also increase. Also, it was noted earlier that the variance of demand is usually highly correlated with the level of demand; thus, products for which average demand is greater offer greater potential for unit savings in safety stock.

Forecast Quality. The impact of forecast quality on unit savings is also quite predictable. Forecast quality translates directly into a reduction of the unexplained variability of demand during each lead time. Required safety stock is reduced commensurately. Table 3 shows a strong positive effect of forecast quality on savings in safety stock ($t = \$13.42$; $p < .001$).

Variability of Lead Time. One of the most interesting but less obvious findings of the simulation is that the variability of lead time is negatively related to unit savings ($t = -11.88$; $p < .001$). This may appear to be counterintuitive because the absolute level of safety stock clearly must be increased when the variability of the lead time increases. (See Equations 1 and 2, and the example below.) However, the portion of the safety stock that is attributable to *demand variability* (as opposed to the portion attributable to lead time variability) actually declines, reducing the potential for unit savings. Therefore, a negative relationship between lead time variability and unit savings is obtained because a good forecast of demand is less valuable when it is the lead time variability that is driving the need for safety stock. A simple empirical example should help to clarify this point.

Assume the following parameters for a product:

$CV_t = .2$

$t = 3$ weeks

$CV_d = .2$

$d = 300$ units

F = .5

k = 1

Then, from the relationships in Table 2:

$S_t = CV_t * t = .6$

$S_d = CV_d * d = 60$

$S_f = S_d * (1\text{-}F) = 30$

Table 6 presents the effect on unit savings when the variability of the lead time (CV_t) is increased from .2 to .5. Values in Table 6 were computed with Equations 1 and 2.

TABLE 6
ILLUSTRATING THE EFFECT OF LEAD TIME VARIABILITY ON UNIT SAVINGS

	Safety Stock		Unit
CVt	DS	FS	Savings
.2	208	187	21
.5	462	453	9

Note that the absolute level of safety stock increases along with the variability of lead time. Unit savings, however, decline for the reason given above. Although the example in Table 6 is provided as an illustration only, it is not a special case. The relationship between lead time variability and safety stock savings holds for the full range of input values adopted in this paper.

Customer Service. The positive relationship between customer service level and unit savings is straightforward ($t = +8.37$; $p < 001$). Safety stock requirements are exponentially related to the desired level of customer service under either system. Therefore, a higher level of customer service magnifies the benefit of switching to the FS and produces greater unit savings in safety stock.

An apparent surprising aspect of this simulation result is that customer service level did not emerge as a more important predictor variable (that is 5th out of 6). This is explained by the relatively narrow range of input values for this variable (90% to 99% probability of a stockout in a lead time). This range was chosen because provision of customer service below this level is seldom feasible in most industries.

Lead Time. Lead time has a positive effect on unit savings ($t = +4.76$; $p < 001$). The explanation for this result is that the variance of demand during a lead time is greater when the lead time period is longer. As noted above, the major factor affecting unit savings is demand variability. Thus, one would expect lead time to be directly related to unit savings.

Interestingly, lead time also exerts a negative effect on unit savings by another route. As noted previously, variability of lead time affects unit savings negatively. Since a longer lead time is generally accompanied by greater variability of lead time, increased lead time also produces a negative, indirect effect on unit savings.

The results of the simulation indicate that the overall effect of lead time on unit savings is positive (i.e., the positive effect described above more than offsets the negative effect). The presence of the two opposing effects helps explain why lead time is the least important of the independent variables.

MANAGERIAL IMPLICATIONS

Recall that the three major objectives of the simulation were:

1. to quantify the safety stock savings across the broad range of products that can be expected by a firm switching from the DS to the FS.

2.. to determine whether the reduction in safety stock achieved by adopting the FS varies by product.

3. to identify the product-specific variables that explain the variation in savings.

Accordingly, the discussion of the study's results and implications is organized around these three objectives.

Expected Savings

Since the unit savings in safety stock for a given product are highly dependent upon the absolute level of demand for that product, this section presents the percentage savings that can be expected. Recall that percentage savings equal unit savings divided by the safety stock required under the Demand System -- (PS = US/SSd).

The mean savings in safety stock for the 729 product configurations examined in the simulation was 14.6%. From a managerial perspective, it seems likely that the cost savings associated with this level of inventory reduction are worth pursuing.

Inter-Product Variation in Potential Savings

While the high mean level of savings achieved by changing to the FS is noteworthy, it is important to examine the dispersion of possible outcomes. An analysis of the distribution of outcomes could confirm that the mean percentage savings (14.6%) is "representative" of the savings that a company can expect. Alternatively, it could show that a small number of products, for which extraordinarily high savings are achieved, are distorting the results.

In fact, for the top 10% of the products studied, the mean savings was 50%; for the bottom 10% of the products, mean savings were less than 1%. Table 7 presents the distribution of percentage savings for the 729 products analyzed. Note that savings are 5% or less for almost a third of the products analyzed. Therefore, it is important to know which variables are good predictors of potential savings in order to enable a manager to decide whether or not to adopt the FS.

Factors Affecting the Potential Savings in Safety Stock

Demand variability, average demand, and forecast quality are the most important variables explaining unit savings. Customer service level and lead time variability are of moderate importance. It is interesting to note that forecast quality is not the most important explanatory variable. This suggests that a reasonably good forecast can produce major savings as long as other independent variables are favorable.

The implication of the negative relationship between unit savings and lead time variability merits emphasis. In a situation where lead time variability is very high, the benefits from adopting the FS are reduced. Under these circumstances, managers should devote more attention to improving the consistency of supply. Only when that situation is remedied can the full gains from the forecast system be realized.

TABLE 7
FREQUENCY DISTRIBUTION OF PERCENTAGE SAVINGS

Percentage Savings	Count	%
<=1	72	9.9
<=2	45	6.2
<=5	126	17.3
<=10	162	22.2
<=25	171	23.5
<=40	81	11.1
<=55	45	6.2
<=70	27	3.7
Total	729	100.0

Finally, all independent variables in this simulation are uncorrelated. This leads to the additional implication that the effect of the independent variables on unit savings is additive. Thus, conditions in which two or more of these variables are favorable provide a larger incentive to adopt the FS.

ENDNOTES

[1]Customer service level is defined in terms of the percentage of lead times in which a stockout occurs. The authors recognize that alternative measures of customer service are equally valid and that the results in this paper apply strictly to the definition adopted herein.

[2]Frank P. Buffa and John I. Reynolds, "A Graphical Total Cost Model for Inventory-Transport Decisions," *Journal of Business Logistics*, Vol. 1, No. 2 (1979), p. 131.

[3]David Ronen, "Measures of Inventory Availability," *Journal of Business Logistics*, Vol. 3, No. 1 (1982), p. 47.

[4]Keith Howard, "Inventory Management in Practice," *International Journal of Physical Distribution and Materials Management*, Vol. 14, No. 2 (1984), p. 23.

[5]John T. Mentzer and R. Krishnan, "The Effect of the Assumption of Normality on Inventory Control/Customer Service," *Journal of Business Logistics,* Vol. 6, No. 1 (1985), p. 103.

[6]Dennis E. Blumenfeld, Randolph W. Hall, and William C. Jordan," Trade-Off Between Freight Expediting and Safety Stock Inventory Costs," *Journal of Business Logistics*, Vol. 6, No. 1 (1985), p. 94.

[7]Frank P. Buffa, "Restocking Inventory in Groups: A Transport Inventory Case," *International Journal of Physical Distribution and Materials Management*, Vol. 16, No. 3 (1986), p. 31.

[8]Robert G. Brown, *Materials Management Review Systems* (New York: John Wiley & Sons, 1977), p. 136.

[9]R. P. Mohanty and V. Chandrashekhar, "Computer Simulation Study for a Production-Distribution System," *International Journal of Physical Distribution and Materials Management*, Vol. 13, No. 3 (1983), p. 57.

[10]Moshe Dorr and Pierre Trudeau, "Inventory Routing: Operational Design," *Journal of Business Logistics,* Vol. 9, No. 2 (1988), p. 174.

[11]Same Reference as Endnote 6, p. 104.

[12]David P. Herron, "Integrated Inventory Management," *Journal of Business Logistics*, Vol. 8, No. 1 (1987), p. 101.

[13]Robert B. Fetter and Winston C. Dalleck, *Decision Models for Inventory Management* (Homewood, Ill.: Richard D. Irwin, Inc., 1961), pp. 105-108.

[14]Generally, whenever the sales forecast is able to explain a non-zero portion of the variance in demand, the forecast system will improve upon the demand system. In this case, the standard deviation of forecast errors will be less than the standard deviation of demand and safety stock will be reduced. The demand system will require less safety stock than the forecast system only in the very special case when the forecast of demand is poorer than the naive forecast implicit in the demand system.

[15]Same reference as Endnote 8, p. 152.

[16]In several cases it was necessary to modify this rule slightly to keep the input values meaningful. For example, it was not possible to increase the maximum level of customer service by 50% since a customer service level above 100% is not meaningful.

ABOUT THE AUTHORS

Walter Zinn is an Assistant Professor in the Department of Marketing at the University of Miami (Florida). Dr. Zinn received his M.B.A. and Ph.D. degrees from Michigan State University. His current research interests are in the areas of postponement, inventory management, and the interface between marketing and physical distribution. Dr. Zinn's publications have appeared or are forthcoming in the *Journal of Business Logistics*, *International Journal of Logistics Management*, *International Journal of Physical Distribution and Logistics Management*, *Journal of Global Marketing*, and *Business Horizons*.

Howard Marmorstein is an Assistant Professor in the Department of Marketing at the University of Miami (Florida). Dr. Marmorstein received his M.B.A. from the University of Pennsylvania and his Ph.D. from the University of Florida. His current research interests are in the areas of consumer information processing, sales management, and logistics. Dr. Marmorstein's publications have appeared or are forthcoming in the *Journal of Consumer Research*, *Journal of Personal Selling and Sales Management*, and *Industrial Marketing Management*.