THE EFFECT OF AUTOCORRELATED DEMAND ON CUSTOMER SERVICE

by

Walter Zinn
Howard Marmorstein
and
John Charnes
University of Miami - Florida

The concept of a correlation between two variables is well known. It measures the extent to which changes in one variable are accompanied by changes in a second variable. Autocorrelation refers to a *single* variable. It measures the extent to which values for a variable are correlated over time. The demand for sweaters, for instance, tends to be autocorrelated. A spell of cold weather may cause sales to be above average for several successive days. With the return of warm weather, sales may then remain below average for successive days. Figures 1 and 2 illustrate the pattern exhibited by autocorrelated demand data in a time plot of 100 daily observations. The daily demands in Figures 1 and 2 have equal mean and variance, but only the data in Figure 2 are autocorrelated. The daily demands in Figure 1 are independent over time.

A standard procedure used for determining the safety stock necessary to achieve a managerially specified level of customer service assumes that daily demands from customers are not autocorrelated. While autocorrelation of daily demands is often present, it frequently goes undetected or is ignored by inventory managers. Therefore, the purpose of this paper is to explain and quantify the effect of autocorrelated demand on the level of customer service provided by a firm.

Results indicate that the impact of autocorrelated demand on customer service is threefold. First, observed stockouts will be significantly more frequent and larger than expected. Second, the effect of autocorrelation on the number of stockouts observed is directly related to the variability of customer demand. Third, the effect of autocorrelated demand on the number of stockouts observed is inversely related to the variability of lead time from suppliers.

FIGURE 1

TIME PLOT OF INDEPENDENT DAILY DEMANDS WITH MEAN OF 100 AND VARIANCE OF 400

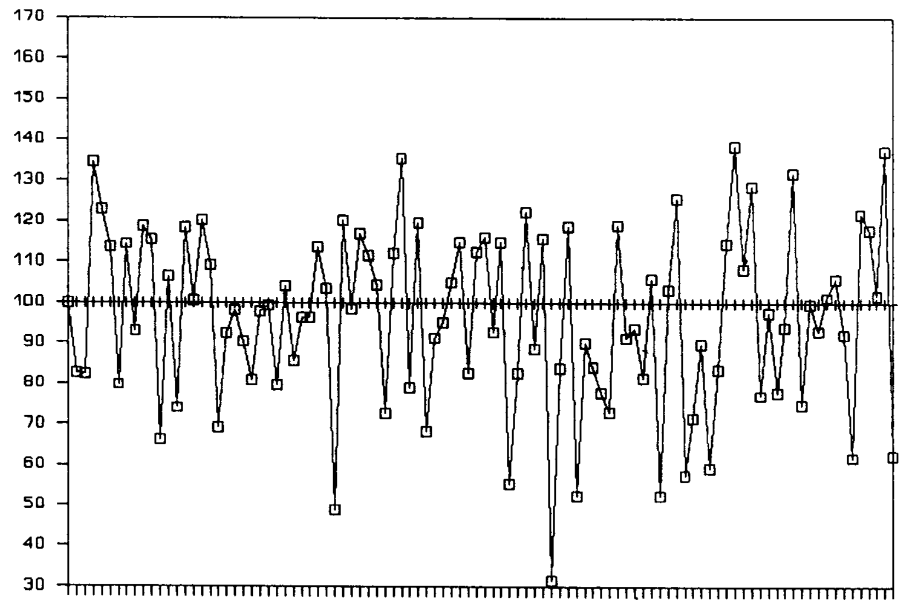

FIGURE 2

TIME PLOT OF AUTOCORRELATED DAILY DEMANDS WITH MEAN OF 100 AND VARIANCE OF 400

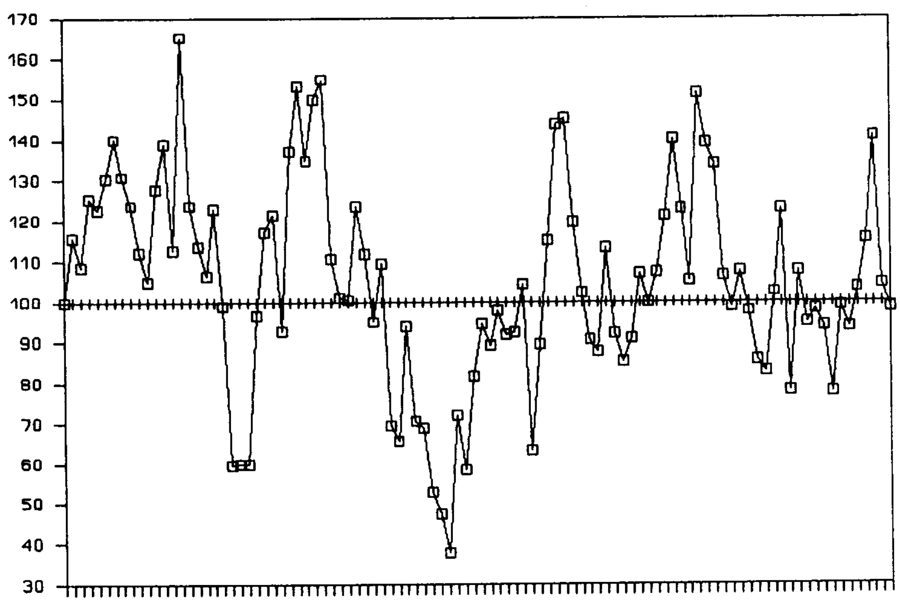

The paper begins with a brief review of the factors affecting safety stock requirements and a standard approach that incorporates these factors. The concept of autocorrelation is then introduced and its relevance to customer service provision is discussed in a qualitative fashion. Next, the results of a simulation study are described. These results enable us to quantify the effect of autocorrelated demand on both the number and magnitude of stockouts experienced by a firm's customers. Supplementary analyses of the results of the simulation identify the variables that moderate the effect of autocorrelated demand on customer service. These results are presented to clarify the circumstances under which the impact of autocorrelation will be more (or less) consequential. Methods available in the literature to mitigate the effect of autocorrelation on customer service are then suggested. Managerial implications are discussed in the last section of the paper.

BACKGROUD

In order to manage the opportunity cost of stockouts, firms must maintain a level of safety stock that balances the loss of sales and customer goodwill with the cost of carrying inventories.[1] The complexity associated with this decision stems from the fact that both daily demand from customers and lead time from suppliers can vary. The greater the variability, the larger the safety stock needed. The level of variability is measured by the variance of lead time demand. Lead time demand is defined as the total quantity demanded by customers during a lead time. The variance of lead time demand is obtained by combining the variance of daily demand and the variance of lead time.[2]

$$s^2_c = m_l \cdot s^2_d + m_d^2 \cdot s^2_l \tag{1}$$

where s^2_c is the variance of lead time demand, m_l is the average lead time in days and m_d is the average daily demand. s^2_d and s^2_l are the variances of daily demand and of lead time, respectively. The safety stock level (SS), in units, is obtained by:

$$SS = k \cdot s_c \tag{2}$$

where k is a dimensionless safety factor relating the investment in safety stock to measures of customer service.[3]

Under the simplifying assumption that lead time demand is distributed normally, the value for k can be taken directly from the cumulative normal distribution. For example, to achieve a service level of 95%, k is set to 1.645. Recent research, however, has shown that the assumption of normality is frequently violated and that the consequences of making this assumption erroneously are severe.[4] Subsequent research has also shown that departures from normality can be accommodated relatively easily.[5]

Autocorrelation, Safety Stock and Customer Service

Embedded in the aforementioned approach to setting safety stock is the assumption that the daily demands within a lead time are independent. This implies that the demand observed in any day is uncorrelated with the demand in previous days (i.e., autocorrelation = 0). When autocorrelation is present, this assumption is violated.

Autocorrelation may carry a positive or a negative sign. When autocorrelation is positive, there will be a sequence of days within a lead time where actual demand is continuously above (below) the expected demand. As a result, the variance of lead time demand increases systematically. When autocorrelation is negative, a day where demand is above average will tend to be followed by a day where demand is below average. While negative autocorrelation is theoretically possible, it is extremely rare in practice. Therefore, this research focuses on positive autocorrelation exclusively.

A more precise view of the effect of positive autocorrelation on the variance of lead time demand can be illustrated analytically. Equation 3 shows how the variance of lead time demand is computed from the variance of daily demand,[6] in the case where lead time is constant at m days. When daily demands are independent, the variance of demand within a lead time is exactly equal to the sum of the variance of daily demands. In contrast, when daily demands are autocorrelated, the variance of demand within a lead time equals the sum of the variance of the daily demands *plus* the covariance of the daily demands.

$$m_1 \,.\, s^2_d = \sum_{i=1}^{m} s^2 d_i + 2 \sum_{i<j} \sum COV\ (d_i,\ d_j) \qquad (3)$$

where $m_1 \cdot s^2_d$ is the variance of lead time demand when lead time is constant at m days, and $COV(d_i, d_j)$ is the covariance between demands on days i and j.

Each of the covariance terms is positive when daily demands are positively autocorrelated. Consequently, equations 1 and 2 (see above) will underestimate the variance of lead time demand and the required level of safety stock, respectively. The resulting level of customer service provided will necessarily be *less* than intended.

THE SIMULATION

This section begins by describing the method used to model autocorrelated demand. The goals and input to the simulation are described next. Several additional features of the simulator are then noted before examining the results.

Modelling Autocorrelated Demand

The model used to simulate the demand from customers is a first order autoregressive AR(1) process. The AR(1) process is described below:

$$d_t = a + bd_{t-1} + e_t, \quad |b| < 1, \tag{4}$$

where d_t is daily demand at time t, a is an intercept term, b is the autoregressive coefficient that describes the strength of the linear association between the demand in two successive days, and e_t is the non-autocorrelated random error term, with mean zero. This model was selected because it accurately describes many actual autocorrelated demand patterns.[7]

Simulation Input Variables

As noted in the introduction, the goals of the simulation are (1) to measure the effect of different levels of the autocorrelation parameter on the level of customer service provided and (2) to identify the variables that moderate the effect of autocorrelated demand on customer service. With these goals in mind, the variables included in the simulation are shown in Table 1 and described next.

In addition to the autocorrelation parameter, five other variables were incorporated in the simulation. The variables are average demand, demand variability, average lead time, lead time variability, and intended level of service. These variables were selected because equations 1 and 2 indicate that they affect safety stock when the standard approach to determining inventory is used.

Note in Table 1 that the variability of both the demand and the lead time were operationalized in terms of the coefficients of variation rather than their respective variances. This was preferred for two reasons. First, it preserves the othogonality of the independent variables in the simulation. This facilitates the interpretation of the results obtained in the subsequent regression analysis. Second, it enhances the external validity of the results by incorporating realistic relationships among the input variables. 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).[8]

TABLE 1

SIMULATION INPUT

Input	Definition	Values
b	autocorrelation parameter	0, .2, .4, .6, .8
MD	average daily demand in units	100, 200
CvD	coeff. of variation of demand	.20, .35
ML	average lead time in days	4, 9
CvL	coeff. of variation of lead time	.20, .35
SL	service level	90%, 99%

The Simulator

The simulator was run on an IBM 3090 mainframe.[9] The demands and the lead times were generated randomly. Demand is normally distributed, while the lead time distribution is discrete uniform. The range of each lead time distribution was determined by the coefficient of variation specified in the input section.

The simulation was comprised of 160 runs, where each run represents one of all possible combinations of input levels. The length of each run is 90,000 lead times. This sample size reprsents a 95% certainty that the random error in the observed number of stockouts stays within +/- .2% of the lead times.

As a check upon the validity of the simulator, one can examine the number of stockouts obtained for the cases when the autocorrelation parameter equals zero. In these cases, the number of stockouts observed should equal the theoretical values predicted by equations (1) and (2) plus or minus a random error term as noted above. Tables 2 and 3 provide these results, which confirm both the validity of the simulator and the adequacy of the sample size.

Three additional features of the simulator merit mention before reporting the simulation results. First, incoming orders can be used to satisfy demand only on the day after arrival. Second, all lead times are independent. Both of these characteristics are implicit in equation (1).

Third, the value of k used for each run was not obtained from the normal distribution. Instead, a more accurate value of k was calculated from the theoretical lead time distribution that pertained to each run. While the distribution of lead time demand was approximately normal, the assumption of normality would have produced substantial error in the computation of safety stocks regardless of the presence of autocorrelated demands. Again, the results in Table 2 and 3 indicate that the procedure for establishing k was sufficiently precise.

RESULTS

Recall that the primary objective of this research concerns the effect of autocorrelated demand on customer service when that fact goes undetected or is ignored by the inventory manager. A second objective is to identify the circumstances under which the effect of autocorrelation on customer service is of greatest importance to the firm. Results pertaining to these two objectives are therefore discussed sequentially in this section.

Effect of Autocorrelated Demand on Customer Service

Table 4 presents the mean percentage of lead times where stockouts occur as the autocorrelation parameter increases. Results are presented for two intended levels of customer service. It is apparent that the increases in stockouts are substantial, even in cases where the autocorrelation parameter is small and thus more likely to remain undetected or be ignored by managers.

TABLE 2

OBSERVED AND EXPECTED NUMBER OF LEAD TIMES WITH A STOCKOUT WHEN THE AUTOCORRELATION PARAMETER EQUALS ZERO AND SERVICE LEVEL EQUALS 90 PERCENT

Input	Value	Expected # Of Stockouts	Observed # Of Stockouts
MD	100	9000	8995
	200	9000	9010
CvD	.2	9000	9013
	.35	9000	8992
ML	4	9000	9002
	9	9000	9003
CvL	.2	9000	9012
	.35	9000	8993

TABLE 3

OBSERVED AND EXPECTED NUMBER OF LEAD TIMES WITH A STOCKOUT WHEN THE AUTOCORRELATION PARAMETER EQUALS ZERO AND SERVICE LEVEL EQUALS 99 PERCENT

Input	Value	Expected # Of Stockouts	Observed # Of Stockouts
MD	100	900	927
	200	900	898
CvD	.2	900	917
	.35	900	908
ML	4	900	916
	9	900	909
CvL	.2	900	899
	.35	900	926

TABLE 4

RELATIVE FREQUENCY OF LEAD TIMES WITH STOCKOUTS FOR DIFFERENT LEVELS OF AUTOCORRELATION

Autocorrelation Parameter	Percent Lead Times With Stockouts	
	SL=90%	SL=99%
0.00	10.0	1.0
0.20	10.9	1.8
0.40	12.3	3.2
0.60	14.6	5.7
0.80	19.5	11.1

Table 5 examines the effect of autocorrelation when an alternative measure of customer service is employed. These data show that the *average magnitude* of a stockout is directly related to the autocorrelation parameter and that this increase also is substantial. The average magnitude of a stockout is measured by the number of units short when a stockout occurred as a percentage of the expected lead time demand. Thus, when autocorrelation is present but goes undetected or is ignored by management, customers will experience larger and more frequent stockouts.

TABLE 5

AVERAGE MAGNITUDE OF STOCKOUTS AT FIVE LEVELS OF AUTOCORRELATION

Autocorrelation Parameter	Average Magnitude Of Stockouts	
	SL=90%	SL=99%
0.00	10.5	6.2
0.20	11.7	7.4
0.40	14.6	10.4
0.60	19.8	14.8
0.80	31.4	27.5

Factors Moderating the Effect of Autocorrelation on Customer Service

The second goal of the research is to identify the circumstances under which the manager should be more (or less) concerned with autocorrelation. This is accomplished by including five additional variables in the simulation (average demand, coefficient of variation of demand, average lead time, coefficient of variation of lead time, and intended service level) and testing whether they moderate the effect of autocorrelation on customer service.

Both the main effects of the supplemental variables and their respective interactions with the autocorrelation parameter are examined by way of a multiple regression analysis. The five variables mentioned above are included as independent variables. The square of the autocorrelation parameter is the sixth independent variable (b2). This transformation of the autocorrelation parameter is needed because its effect on customer service is non-linear, which violates an assumption of multiple regression analysis.

The dependent variable is the percent difference from the expected number of stockouts. This measure was selected to provide a dependent variable that is comparable across the 2 intended levels of service. Results are presented in Table 6.

The six possible main effects are considered first. The interactions of interest are then examined and explained.

Main Effects

Recall that the dependent measure in the regression is the percent difference from the expected number of lead times under which stockouts occur. As expected, this measure of customer service is not affected significantly by the average demand (MD), coefficient of variation of demand (CvD), average lead time (ML), or coefficient of variation of lead time (CvL). These null results are obtained because the standard approach to setting safety stock (see equations 1 and 2) is specifically designed to incorporate changes in the foregoing variables while maintaining a predetermined probability of a stockout. For instance, if MD increases, equation (1) shows that the standard deviation of lead time demand increases. Equation (2) then shows that the level safety stock needed to obtain a predetermined probability of a stockout increases commensurately. As a result, the increase in average demand does not affect the probability of a stockout.

TABLE 6
MULTIPLE REGRESSION ANALYSIS

Dependent Variable:
Percent Difference From Expected Number Of Stockouts

Independent Variable	Parameter Estimated	Standardized Estimated	t-value	Problem Level
Intercept	-33.420	3.2389	-10.32	0.0001
SL	.356	0.0312	11.38	0.0001
MD	-.000	0.0040	-0.01	0.9922
CvD	-.612	2.6762	-0.22	0.8192
ML	-0.25	0.0802	-0.31	0.7577
CvL	.207	2.6762	0.07	0.9384
b2	7.809	3.9399	1.98	0.0493
MDxb2	-.000	0.0119	-0.55	0.9566
CvDxb2	18.800	7.9513	2.36	0.0193
MLxb2	.017	0.2385	-0.07	0.9435
CvLxb2	-16.055	7.9513	-2.02	0.0453

Analysis of Variance

Source	DF	Sums of Squares	Mean Square	F-Ratio	Problem Level
Model	10	1089.981	108.998	34.40	0.000
Error	149	472.074	3.168		
Total	159	1562.055			
R Squared			0.6978		
Adjusted R Squared			0.6775		

There are two significant main effects, the intended level of service (SL) and the square of the autocorrelation parameter (b2). The reason for observing a significant effect on SL is statistical rather than theoretical. For the same reasons used to describe the results for the four variables above, there is no theoretical justification to expect a significant main effect of SL on the dependent variable. Statistically, however, the main effect is observed because the interaction of SLxb2 was excluded from the regression model. This interaction had to be eliminated because it is highly collinear with b2, which obscured the interpretation of results.

The other variable exerting a significant main effect on customer service is the square of the autocorrelation parameter (b2). This result was readily apparent in Table 4 and is merely confirmed by the regression results in Table 6 ($t = 1.98$; $p < .05$). As b2 increases, the percent difference between the expected and the observed number of lead times under which stockouts are observed increases significantly. This means that the higher the level of autocorrelation, the larger the difference between the level of customer service expected and observed by a manager. As noted earlier in the paper, the increase in the number of stockouts is explained by the increase in the variance of lead time demand caused by the autocorrelation parameter. The following discussion focuses upon the two significant interaction effects which were observed.

Interactions

The primary purpose of the regression analysis was to determine whether the effect of autocorrelation depends upon the level of any (or all) of the remaining four independent variables. Therefore, four interactions were included in the regression model. Two of the interactions were not significant and are not discussed further. There was no interaction of autocorrelation with mean demand or with mean lead time (see Table 6, $p > .20$ in both cases). The two significant interactions are now described and explained.

For each of the observed interactions, it is important to understand both the direction and the magnitude of the effect. Therefore, a conceptual explanation of each effect is offered first. This is followed by a small numerical illustration which conveys the magnitude of the effect in a more concrete form. This enables the reader to assess the substantive importance, in addition to the statistical significance, of these effects.

The interaction of b2 with the coefficient of variation of demand (CvD) is somewhat more intuitive and is therefore examined first. Recall that when daily demands are autocorrelated, the actual variance of demand in a lead time exceeds the sum of the variance of the daily demands. Thus b2 causes an increase in stockouts by magnifying the variability of lead time demand beyond its expected level. The significant interaction of b2 with CvD (Table 6, $t = 2.30$, $p < .05$) merely reflects the fact that the effect of b2 on the number of stockouts is multiplicative with the variability of daily demand (operationalized herein by CvD). The greater the variability of daily demand, the greater the effect of autocorrelation on the number of stockouts observed. This is illustrated numerically by the following example, which compares the effect of autocorrelation at two levels of CvD.

TABLE 7

EFFECT OF AUTOCORRELATION ON PERCENT LEAD TIMES WITH STOCKOUTS AT 2 LEVELS OF DEMAND COEFFICIENT OF VARIATION

Autocorrelation Parameter (b)	Squared B (b2)	Demand Coefficient of Variation (CvD) 20	35
0	0	10.0	10.0
.8	.64	17.0	22.1

When CvD is low, the effect of b is depicted by the difference in customer service provided when $b=0$ versus $b=.8$. The percentage of lead times with stockouts increases from 10.0% to 17.0% or a difference of 7%. In contrast, when CvD is high, the corresponding difference in customer service resulting from increasing autocorrelation is 12.1%. This example shows that the effect of autocorrelation is significantly greater when the level of demand variability is high. The implications of this result are discussed in the concluding section of the paper.

The negative interaction of b2 with the coefficient of variation of lead time (CvL) (Table 6, $t = -2.02$, $p < .05$) is examined next. The effect of b2 on the number of lead times with stockouts becomes smaller as CvL increases. This result is explained by the impact exerted by lead time variability (CvL) upon the absolute amount of safety stock on hand.

When CvL is high, the percent difference between the safety stock on hand and the safety stock that would be needed to compensate for the presence of autocorrelation is relatively small. In contrast, when CvL is low, a smaller absolute level of safety stock must be maintained. Consequently, the percent difference between the safety stock on hand and the safety stock that would be needed to compensate for the presence of autocorrelation is relatively greater. Consider the following illustration, which compares the effect of autocorrelation on the percentage of lead times with stockouts at two levels of lead time variability.

TABLE 8

EFFECT OF AUTOCORRELATION ON PERCENT LEAD TIMES WITH STOCKOUTS AT 2 LEVELS OF LEAD TIME COEFFICIENT OF VARIATION

Autocorrelation Parameter (b)	Squared B (b2)	Lead Time Coefficient of Variation (CvL) .20	.35
0	0	10.0	10.0
.8	.64	21.7	17.4

When CvL is low, the effect of increasing autocorrelation is shown by the difference in customer service provided at $b=0$ versus $b=.8$. The percentage of lead times with stockouts increases from 10.0% to 21.7%, an absolute increase of 11.7%. In contrast, when CvL is high, a much smaller decrease in customer service results. The percentage of lead times with stockouts still increases, but the change is only 7.4%.

In summary, this research has revealed three results pertaining to the impact of autocorrelation on customer service. First, when autocorrelated demand is present but goes undetected or is ignored by management, the number of stockouts observed will be significantly larger than expected. Furthermore, the magnitude of each stockout observed also will increase substantially. Second, the greater the variability of demand, the greater the effect of autocorrelation on the number of stockouts observed. Finally, the number of unexpected stockouts that result from autocorrelation increases as the variability of lead time decreases.

Note that these results are subject to a number of limitations. First, while customer demand was modelled with an autoregressive model AR(1), there are other models available. Examples include the first order moving average MA(1), and mixed models such as the autoregressive moving average (ARMA). Second, in accordance with the assumption implicit in equation (1), daily demands were maintained independent across lead times in the simulation. The removal of such limitations are feasible suggestions for further research. Managerial implications for each result are discussed next.

MANAGERIAL IMPLICATIONS

The first result is that autocorrelation has a significant impact on the number and magnitude of stockouts. Thus, managers should detect the presence of autocorrelation. One possible method of detection is visual inspection of time plots of daily demands. While this method is simple, it also is imprecise and may fail to recognize moderate levels of autocorrelation. Alternatively, first order autocorrelation can be detected with a well known statistical test, the Durbin-Watson statistic.[10] A third method of detection is to calculate the sample autocorrelation function (ACF).[11] This procedure is more complex than the Durbin-Watson statistic, but it also is more comprehensive in that it allows the manager to detect second and higher order autocorrelation.

If autocorrelation is detected, the manager may either increase safety stock to reflect the actual variance of lead time demand, or attempt to correct for the effect of autocorrelation. An estimate of the required increase in safety stock is presented in the Appendix.

There are at least three methods available in the literature[12] to treat demand data in order to correct for the effect of autocorrelation on customer service. All three methods involve demand forecasting. First, the autocorrelation pattern in demand may be explained and reduced by including additional variables in a regression model. A second possibility is to utilize a time series technique such as exponential smoothing to forecast demand. Finally, an appropriate variable transformation procedure may be employed to moderate the autocorrelation.[13]

The procedures indicated will often mitigate but not necessarily eliminate the effect of autocorrelation on customer service.[14] It has been shown, for example, that some autocorrelation typically remains in the error terms of the demand forecast when exponential smoothing is used to forecast demand.[15] Moreover, the choice of forecasting technique also will affect the level of autocorrelation in the residuals.[16] As a result, this research reinforces the importance of using demand forecasts in the management of inventories, with the proviso that firms using packaged forecasting techniques evaluate how well their chosen techniques correct for autocorrelation. This evaluation can be accomplished by applying one of the three autocorrelation detection techniques given at the beginning of this section to the forecasting model residuals.

The second result is the moderating effect of demand variability on the impact of autocorrelation on customer service. This implies that managers must assign priority in detecting autocorrelation to those items displaying the highest variability of demand.

The final result shows that the impact of autocorrelation on customer service is moderated by lead time variability. Specifically, lower levels of variability in lead time increase the effect of autocorrelation on customer service. This result is particularly significant because the variability of lead times is steadily being reduced by logistics advancements such as Electronic Data Interchange (EDI) and Just-in-Time (JIT). Thus, the need to recognize and respond to autocorrelated demand will become increasingly important in the future.

Authors' note: Funding for this project was provided to the second author by a generous grant from Francine and Herbert Tobin.

NOTES

[1]See, for instance, K. Howard, "Inventory Management in Practice," *International Journal of Physical Distribution and Materials Management* 14, no. 2 (1984): 3-36; D. P. Herron, "Integrated Inventory Management," *Journal of Business Logistics* 8, no. 1 (1987): 96-116; or J. F. Campbell, "Designing Logistics Systems by Analyzing Transportation, Inventory and Terminal Cost Tradeoffs," *Journal of Business Logistics* 11, no. 1 (1990): 159-179.

[2]R. B. Fetter and W. C. Dalleck, *Decision Models for Inventory Management* (Homewood, Ill.: Richard D. Irwin, Inc., 1961).

[3]Two measures of service are adopted in this research, the probability of a stockout within a lead time and the average magnitude of a stockout. The authors recognize that additional measures of service are equally valid and that the results in this paper apply strictly to the measures adopted herein.

[4]J. T. Mentzer and R. Krishnan, "The Effect of the Assumption of Normality on Inventory Control/Customer Service," *Journal of Business Logistics* 6, no. 1 (1985): 101-120.

[5]H. Lau, "Toward an Inventory Control System Under Non-Normal Demand and Lead-Time Uncertainty," *Journal of Business Logistics* 10, no. 1 (1989): 88-103.

[6]A. M. Mood, F. A. Graybill, and D. C. Boes, *Introduction to the Theory of Statistics* (New York: McGraw-Hill, 1974): 178-179.

[7]S. Fotopoulos, M. C. Wang, and S. S. Rao, "Safety Stock Determination with Correlated Demands and Arbitrary Lead Time," *European Journal of Operational Research* 35 (1988): 172-181.

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

[9]The computer program is available to anyone interested.

[10]On the application of this test, see for instance, J. Neter, W. Wasserman, and M. H. Kutner, *Applied Linear Regression Models* (Homewood, Ill.: Richard D. Irwin, Inc., 1989).

[11]G. E. P. Box and G. M. Jenkins, *Time Series Analysis* (San Francisco, Calif.: Holden-Day, 1976).

[12]S. Makridakis and S. C. Wheelwright, *Forecasting Methods for Management* (New York: John Wiley & Sons, 1989).

[13]Same reference as Note 10.

[14]R. Ballou, *Business Logistics Management* (Englewood Cliffs, N.J.: Prentice Hall Inc., 1985), pp. 101-102.

[15]R. G. Brown, *Statistical Forecasting for Inventory Control* (New York: McGraw Hill, 1959), pp. 54-59.

[16]H. V. Roberts, *Data Analysis for Managers* (San Francisco, Calif.: The Scientific Press, 1991), p. 358.

APPENDIX

INCREASE IN SAFETY STOCK REQUIRED TO COMPENSATE FOR AUTOCORRELATED DEMAND

This appendix is designed to quantify the increase in safety stock required to compensate for autocorrelated demand. The results in the table below correspond to the simulation parameters adopted in this research. The values presented represent the average percentage increase in safety stock needed for maintaining a predetermined level of customer service.

Note that the incremental safety stock required is substantial and is directly related to the size of the autocorrelation parameter. Moreover, the average increase across all conditions is 18.1 percent.

TABLE 9

AVERAGE PERCENTAGE INCREASE IN SAFETY STOCK NEEDED TO COMPENSATE FOR AUTOCORRELATION

Autocorrelation Parameter	Average Percentage Increase in Safety Stock
0.00	0.0
0.20	7.9
0.40	15.0
0.60	21.7
0.80	28.0

ABOUT THE AUTHORS

Walter Zinn is associate professor in the department of marketing at the University of Miami (Florida). He 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.

Howard Marmorstein is assistant professor in the department of marketing at the University of Miami (Florida). He 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.

John M. Charnes is assistant professor in the management science department at the University of Miami in Coral Gables, Florida. He holds B. Civil Eng., M.B.A., and Ph.D. degrees from the University of Minnesota. His research interests are in the areas of multivariate time series analysis and multivariate simulation output analysis.