
Evaluating Forecast Performance in an Inventory Control System

Author(s): Everette S. Gardner, Jr.

Source: *Management Science*, Vol. 36, No. 4 (Apr., 1990), pp. 490-499

Published by: **INFORMS**

Stable URL: <http://www.jstor.org/stable/2632012>

Accessed: 16/02/2011 15:18

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=informs>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



INFORMS is collaborating with JSTOR to digitize, preserve and extend access to *Management Science*.

EVALUATING FORECAST PERFORMANCE IN AN INVENTORY CONTROL SYSTEM*

EVERETTE S. GARDNER, JR.

College of Business Administration, University of Houston, Houston, Texas 77204-6282

This paper analyzes the impact of forecasting on inventory decisions in a large physical distribution system. Alternative forecasting models are evaluated by developing tradeoff curves between inventory investment and customer service. The results demonstrate that the choice of forecasting model is an important factor in determining the amount of investment needed to support any target level of customer service.

(FORECASTING—TIME SERIES, APPLICATIONS; INVENTORY/PRODUCTION—PARAMETRIC ANALYSIS; SIMULATION—APPLICATIONS; MILITARY—LOGISTICS)

1. Introduction

Forecasting is a prerequisite to inventory decisions in practice. Unfortunately most research in inventories ignores forecasting altogether and simply assumes that the distribution of demand and all its parameters are known. Only a few studies are available on the interactions between forecasting and inventory decisions. Lee and Adam (1986) show that the size of forecast errors influences the choice of lot-sizing rule in material requirements planning systems for manufacturing inventories. In distribution inventories, Croston (1972), Brown (1982), Watson (1987), and Eppen and Martin (1988) show that forecast errors can seriously distort projections of customer service.

From the research to date, it is not clear how managers should evaluate alternative forecasting models in the inventory context. This paper is a study of the impact of forecasting on inventory control in a large physical distribution system. We show that alternative forecasting models define unique tradeoff curves between aggregate inventory investment and customer service. The differences between the tradeoff curves are significant. Careful selection of the forecasting model for an inventory system can increase the customer service provided by a fixed investment. Another possibility is to reduce investment while maintaining previous levels of customer service.

The plan of this paper is as follows. §2 introduces the concept of tradeoff curves for inventory analysis. This is followed in §3 by a summary of the forecasting and inventory decision rules used in the physical distribution system. §4 analyzes the characteristics of the time series of inventory demands in order to identify alternative forecasting models. These models are reviewed in §5, while §6 develops the research design for testing the models. The savings available to management from improved forecasting are discussed in §7. Finally, conclusions are offered in §8. The results of this study are presently under implementation and should be useful in other inventory systems. The results also point to further research opportunities in forecasting for operational decisions.

2. Tradeoff Curves in Inventory Control

The setting for this study is a large military physical distribution system managed under centralized decision rules. Headquarters inventory managers determine system-wide requirements for about 50,000 inventory items, procure these items from industry, and make allocations of the items to a network of eleven supply centers. The supply

*Accepted by Vijay Mahajan; received May 4, 1988. The paper has been with the author 5 months for 2 revisions.

centers have little decision-making authority and act as warehousing and distribution agents to serve customers in their geographic areas.

Customer service in this system is measured by delay time, the average number of days that a customer requisition is backordered before it can be delivered. When an inventory item is not available for issue at one of the supply centers, headquarters managers contact manufacturers and attempt to expedite procurements. The chief performance goal of management is thus to minimize average delay time for customer requisitions.

Prior to this study, management believed that the only important influence on delay time was the amount of aggregate inventory investment, defined as the sum of average order quantity stocks and safety stocks for all items in the inventory. As the amount of aggregate investment is increased, more funds are available for safety stocks to protect against forecast errors and delay time is reduced. The tradeoff curves in Figure 1 illustrate the relationship between investment and delay time. Consider the curve labeled "simple smoothing." At an investment level of \$375 million, the curve shows that delay time to fill customer requisitions is 43 days. At the other end of the curve, an investment of \$425 million reduces delay time to 33 days. The curve has a negative slope as should be expected. At a relatively small investment level, safety stocks are small and delay time is high. Moving to the right on the investment axis, safety stocks are increased and delay time falls. Since it is very difficult to measure the cost of delay time in any inventory system, it is not clear exactly where management should operate on the tradeoff curve. The curve simply shows the range of possible operating positions. The correct position is a matter of managerial judgment.

Tradeoff curves between aggregate inventory investment and customer service are widely used as analytical tools in inventory control. However, it is generally accepted practice to consider only one forecasting model for an inventory and thus to develop only one tradeoff curve. See for example the analysis in Brown (1982) or Silver and Peterson (1983). This research shows that different forecasting models define a range of

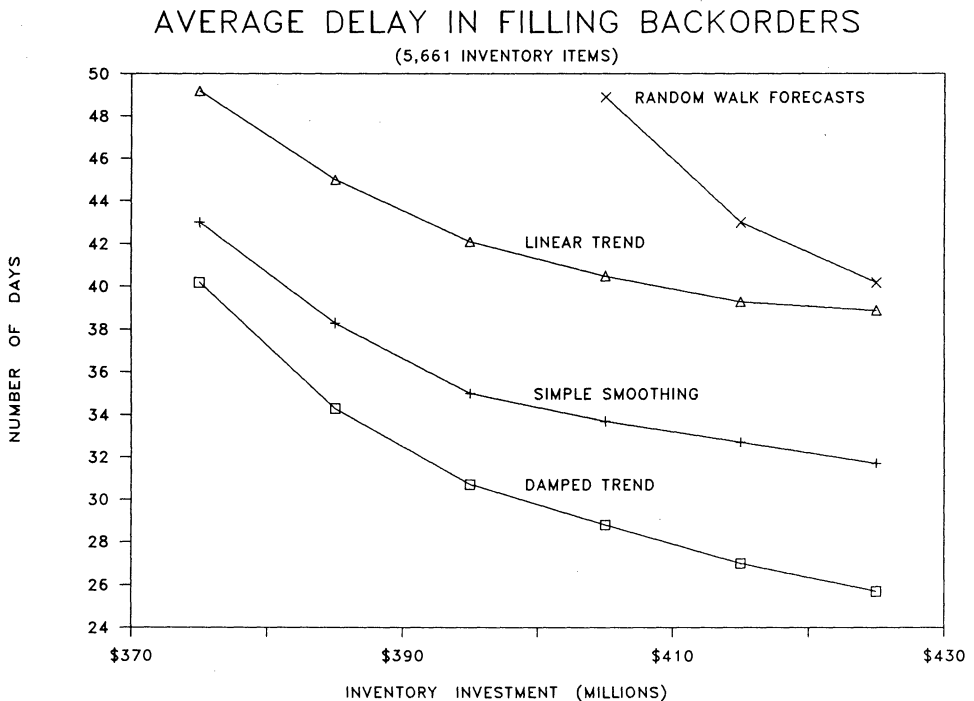


FIGURE 1.

different tradeoff curves between aggregate inventory investment and customer service. The reason is that forecast errors are the primary determinant of the safety stock component of inventory investment. In general, the better the forecast accuracy, the smaller the inventory investment needed to reach any particular target value for customer service. As forecast accuracy improves the best tradeoff curve shifts down and to the left. For example, in Figure 1 the random walk yields a poor tradeoff curve. Forecasting with a linear trend improves the tradeoff curve, while simple smoothing yields even further improvement. Finally, forecasting with a damped trend shifts the tradeoff curve to the position where management is able to achieve the minimal investment to meet any customer service target.

3. Inventory Decision Rules

To understand how the aggregate inventory investment totals in Figure 1 are determined, it is necessary to review the decision rules for order quantities and safety stocks. The decision rules generally follow the classic text by Hadley and Whitin (1960, Chapter 4), although there are some differences worth discussion here. The objective in inventory decisions is to minimize the sum of total variable costs:

$$TVC = A(4D/Q) + IC(Q/2 + R - LD + B) + SE(B/U) \quad \text{where} \quad (1)$$

A = administrative cost of placing an order on procurement plus the manufacturer's production set-up cost.

D = quarterly demand forecast.

Q = order quantity.

I = inventory holding cost rate including storage, obsolescence, and opportunity costs.

C = unit purchase cost of the item.

R = reorder point, composed of leadtime demand plus safety stock.

L = procurement leadtime expressed in number of quarters.

B = expected number of units of stock backordered at any random point in time.

S = shortage cost per customer requisition backordered.

E = essentiality code for the inventory item.

U = number of units of stock per customer requisition.

The first term in (1) is the expected ordering cost for one year. While Hadley and Whitin assume that all parameters of the distribution of demand are known, in this system simple exponential smoothing is used to produce D , the quarterly forecast. Since simple smoothing yields the same forecast for every quarter in the future, $4D$ is taken to be annual demand and $4D/Q$ is an estimate of the expected number of procurement orders per year.

The second term in (1) approximates the carrying costs for the expected number of units of stock on hand at any random point in time. Expected units are defined as average cycle stock investment (approximated by $Q/2$) plus the reorder point minus leadtime demand plus backorders. The sum of the second term in (1) for all inventory items is the aggregate inventory investment (the X axis in Figure 1).

The third term in (1), for shortage costs, contains two differences from Hadley and Whitin. First, shortage costs are weighted by an essentiality index ranging from one to five. These indices reflect the importance of the inventory item to the missions of customers. Another difference is that Hadley and Whitin compute shortage costs as a function of the number of units of stock on backorder while in (1) shortage costs depend on the number of customer requisitions on backorder. A requisition is defined as a demand from a single customer for any number of units of stock, all required for immediate delivery. Thus B/U estimates the number of requisitions on backorder by assuming that the number of units per requisition is constant.

Differentiating the cost expression and solving the first-order conditions leads to the following decision rule for order quantities:

$$Q = [8DA/(IC(1 - P))]^{1/2}. \quad (2)$$

This is the same as the well-known EOQ formula except that it is adjusted by $1 - P$, where P is the probability of being out of stock at any random point in time. As this probability increases, the order quantity increases. P can be written as:

$$P = 1/Q \int_R^{R+Q} [1 - F(x)] dx \quad (3)$$

where $F(x)$ is the cumulative probability distribution of leadtime demand. $F(x)$ is assumed to be normal and simple exponential smoothing is used to estimate the mean and variance of the distribution. As one of the referees for this paper pointed out, a normality assumption for inventory demands is usually suspect. However, there seems to be no reasonable alternative to use of the normal distribution in this application. Extensive study showed that no other standard distribution gave better results. Iterative methods are used to solve these equations simultaneously for Q , P , and R . Details of the solution routine are available in Hadley and Whitin.

One problem with implementing these decision rules is that the inventory costs are unknown. The problem of cost measurement is bypassed by using the costs as policy variables. Different combinations of costs determine a tradeoff curve (see Figure 1) showing how customer service varies as a function of total inventory investment.

For a variety of reasons, forecasting performance should be an important determinant of the effectiveness of these inventory rules. In equation (3) the probability of being out of stock depends directly on forecast accuracy. For a given reorder point quantity, as forecast accuracy improves, the probability of being out of stock declines and thus delay time improves. Another way to view the effects of forecast accuracy is to consider that the reorder point R is composed of leadtime demand and safety stock. As forecast accuracy improves, management has the option of reducing investment because less safety stock is needed to support any given customer service target. In equation (2), it is also important to understand that the validity of the order quantity Q depends on the quarterly demand forecast D . Normal procurement leadtime (not considering expedited procurements) ranges from four to eight quarters. However, the current forecasting model in this system assumes that mean demand will be constant for every quarter in the future. Thus annual demand is taken to be $4D$ and leadtime demand is taken to be LD . If trends or shifts in level of the time series occur, the forecasts will be biased, which will in turn distort order quantities and reorder points. A positive trend in demand will cause shortages of stock while a negative trend will cause excess stocks to accumulate.

4. Time-Series Data Analysis

A data base was made available for this research including a complete daily history of receipt and issue transactions for the previous nine years for all 50,000 items stocked. Fortunately, it was not necessary to analyze all of this data. Management concentrates attention on a sample of 5,661 "Class A" items that account for more than 80% of customer requisitions. This sample was used for data analysis and the forecasting tests discussed below. Data for the Class A items were aggregated into quarterly time series to correspond to established time frames for updating forecasts and inventory decision rules. These series contained an average of 24 observations. About half of the series contained a full nine years (36 quarters) of demand information. The other series were shorter because they represented items that had been added to or deleted from the inventory during the nine-year period.

To gain an understanding of the structure of the series, autocorrelation coefficients were computed. Only 12% of the series had a significant autocorrelation coefficient at lag 1. The percentage of significant coefficients declined steadily to about 3% at lag 5. Thus ARIMA modeling was ruled out. Since the autocorrelations were so small at seasonal lags, there was no need to bother with seasonality in forecasting.

Complex modeling efforts seem even more unlikely when the coefficients of variation of the series are considered. The mean coefficient of variation is about 1.5 with the distribution skewed left. Extreme variability is not unusual in inventory time series (Brown 1982 and Dancer and Gray 1977). In these series, variability is due to at least three factors. First, demand is lumpy because the series contain many zero observations. About 78% of the series have at least one quarter with a zero observation and about half have at least four consecutive quarters with zero demand. Second, graphical analysis revealed that many series contain sudden jump shifts in the level of demand. Third, the series are contaminated by outliers. A random sample of 100 series was examined for outliers by constructing 95% confidence limits around the ex post forecast errors (one-step-ahead) from simple exponential smoothing. In 92 series, at least one outlier was found. Similar results were obtained by constructing confidence limits around the errors from a variety of other smoothing methods. Many of the outliers are huge, on the order of five to six times the level of the time series before the outlier occurred.

Graphical analysis also revealed numerous trends in these series. The trends are erratic so it was not obvious that trend-adjusted forecasting models should be used. Trends may appear to be unlikely from the autocorrelations, but the coefficients are distorted by the zero observations.

To sum up data analysis, the time series are nonseasonal and contain zero observations, jump shifts in level, and outliers. The series are also very short so any conclusions from data analysis are likely to be subject to large sampling error. Under these conditions, some form of exponential smoothing is the only reasonable forecasting method if indeed forecasting should be attempted at all. Prudence suggests the use of a naive (random walk) model as a benchmark for any other forecasting model applied to this data.

It is impractical to make all of these time series available to other researchers. The series are imbedded in a large, complex data base system maintained in a proprietary format on a mainframe computer. However, a random sample of 100 time series was extracted from the data base and placed on a diskette. Any researcher interested in obtaining a copy of the diskette should write to the author.

5. Forecasting Alternatives

The standard fixed-parameter exponential smoothing models, simple smoothing and a linear trend, were selected for testing. The adaptive smoothing approach of Trigg and Leach (1967) was tested since it was designed to deal with jump shifts in the level of demand. In brief, the Trigg and Leach scheme sets the smoothing parameter in the simple model equal to the absolute value of the ratio of smoothed error to smoothed mean absolute deviation. When the time series is well-behaved, the ratio is near zero; when a sudden change occurs, the ratio moves toward one to eliminate bias in the forecasts.

The damped-trend model by Gardner and McKenzie (1985) was also tested since it performed well in the noisy data used in the *M*-competition (Makridakis et al. 1982). Linear-trend smoothing models tend to overshoot the data at long forecast horizons. The damped trend is more conservative and operates by damping the trend according to the level of noise in the time series. The greater the noise, the greater the rate of damping applied.

The models were programmed using the following formulation:

$$S_t = S_{t-1} + \phi T_{t-1} + h_1 e_t, \quad (4)$$

$$T_t = \phi T_{t-1} + h_2 e_t, \quad (5)$$

$$\hat{X}_t(m) = S_t + \sum_{i=1}^m \phi^i T_t. \quad (6)$$

The one-step-ahead forecast error is e_t . S_t and T_t are the level and trend components of the series, with h_1 and h_2 as the smoothing parameters. The type of model is controlled by the damping parameter ϕ . If ϕ is zero, the model is simple smoothing. If ϕ is 1.0, the trend is linear and the model is usually known as Holt's linear-trend model. Finally, when $0 < \phi < 1$, the trend is damped.

To minimize computational problems, parameters in the linear and damped trend models were constrained using discounted-least-squares (DLS). With β as the discount factor, the smoothing parameters become:

$$h_1 = 1 - (\beta/\phi)^2; \quad h_2 = (1 - \beta/\phi)(1 - \beta/\phi^2). \quad (7)$$

DLS simplifies matters considerably since the linear-trend model effectively has only one parameter, β , and the damped trend has two parameters, β and ϕ . The linear-trend model constrained by DLS is also known as double exponential smoothing (Brown 1963).

These models are by no means the only possibilities for the time series. For example, many other adaptive smoothing systems have been proposed in the literature. However, it was necessary to restrict the number of models tested due to the computational problems discussed in the next section.

6. Experimental Design

The traditional research design for testing a forecasting model on a collection of time series is as follows (Makridakis et al. 1982). Each series is divided into two samples. The first sample is used to fit the model and a set of forecasts is made covering all observations in the second or holdout sample. Forecast errors in the holdout sample are averaged across all series. The mean absolute percentage error (MAPE) by forecast horizon in the holdout sample is often used to evaluate competing models.

In this inventory system, it is unlikely that such a research design would yield information convincing management to implement a new forecasting model. Because of the zero observations, the MAPE is undefined for many time series. There is also no way to use the MAPE averaged across all time series to estimate the impact on management's primary concerns, inventory investment and customer service. The mean-squared-error (MSE) is an alternative, but the time series vary significantly in magnitude, making it misleading to average squared errors across the series. The geometric MSE is another alternative but this measure is difficult to interpret and again there is no way to estimate the effects on investment and customer service.

Because of these concerns a large-scale simulation model of the inventory system was developed. The aim of the simulation was to develop tradeoff curves between investment and customer service for each forecasting model. The hypothesis was that the set of models would produce substantially different tradeoff curves. Nine years of daily operations were simulated. On the first day of the simulation, actual values from the past were used to initialize stock on hand, backorders, and outstanding purchase orders. During the simulation, actual demands from customers were processed on the same days that they occurred in the past. As each demand was processed, stock on hand was reduced and backorders were recorded as necessary. The decision rules in equations (2) and (3) were used to determine whether to reorder after each demand transaction. When reorders

were placed, leadtimes were assigned using an empirical distribution of actual leadtimes experienced. New items were brought into stock and obsolete items were deleted on the same days that the actions occurred in the past. At the end of each quarter, forecasts and decision rules were updated. Aggregate inventory investment and customer service were also recorded on a quarterly basis. To determine the variance of the forecast errors needed to solve equation (3) for the probability of shortage, the smoothed MSE was used and forecast errors were assumed to be normal as in the current forecasting system.

The forecasting models were initialized as follows. The level in simple smoothing and adaptive smoothing was set equal to the mean of the first four quarters of demand. For the linear and damped trends, the initial level and trend were computed by regression on the first four quarters.

To deal with outliers and jump shifts in the time series, the current forecasting system uses demand filters (control limits) set at plus and minus three standard deviations of demand. If a new observation falls outside the filters, the forecast is not updated. A check is then made in the subsequent quarter. If two successive observations are both higher or both lower than the filters, a new forecast is made by re-initializing the model using the last four quarters of demand. Tests showed that this procedure made marginal improvements in performance so it was retained in the simulations. The procedure is similar to Whybark's (1973) adaptive smoothing system except that Whybark changes the smoothing parameter when control limits are broken. In this application, parameters are retained and the model is simply restarted to catch up with the data.

The fundamental problem in research design was in choosing smoothing parameters for the forecasting models. The series are too short to rely on fitted error measures as a means of choosing parameters. Even if the series were longer, the forecasting system is so large that it is infeasible to tailor parameters for individual series. Computation times dictate that the same model with the same parameters be used for all series.

These considerations led to the development of a cross-sectional model-fitting procedure, using a sample of 1,000 times series: For each forecasting model, multiple simulation runs for the nine-year period were made using trial parameters in the range 0.0 to 1.0. For each parameter, the three costs in the decision rules for order quantities and safety stocks were varied to produce aggregate investment values over a target range of \$375-\$425 million (the typical operating range for the inventory). Parameters were chosen to yield minimum delay time for values of investment in the target range. The resulting parameters were:

Simple smoothing	$h_1 = 0.2$
Adaptive smoothing	$h_1 = 0.2$
Linear trend	$\beta = 0.7$
Damped trend	$\beta = 0.7, \phi = 0.6$

The h_1 value for adaptive smoothing refers to the parameter used to smooth the forecast errors in the Trigg and Leach scheme.

This model-fitting procedure was tedious, requiring thousands of simulation runs, so no attempt was made to refine the parameters in increments smaller than 0.1. The search for investment values in the target range was also constrained. Experimentation with costs was halted when six investment values within the target range were obtained. It is certainly possible that better parameters could be found in an unconstrained search.

The sample parameters were used to estimate investment/service tradeoffs for the remaining 4,661 series. Finally, before presenting the results to management, a search was made for parameters using the entire set of 5,661 series. The reason is that parameters based on all series were needed to implement the results of the study. Optimal parameters

using all series were the same as sample parameters so the results in the next section are based on all series.

The simulation was validated by comparing simulated investment and service to actual values in the past, using simple exponential smoothing as the forecasting model. The correlation between simulated and actual values was statistically significant. The simulation reached steady-state after four years of operations for the new forecasting models. The results in the next section are based on average investment and service for the last five years of the simulation. The results were stable during this period and it makes little difference how the results are presented.

7. Investment/Service Tradeoffs

Figure 1 summarizes the final results of the study. The damped trend proved superior to the other models at all investment levels. For example, a typical operating value for investment in the Class A items is \$420 million. At this investment, simple smoothing, the current forecasting model, produces delay time of 32 days. The damped trend produces delay time of 26 days, a reduction of 19%. Rather than reduce delay time, management also has the option of using the damped-trend model to reduce investment. For example, if target delay time is set at 32 days, this can be achieved with the damped trend for an investment of about \$390 million compared to \$420 million for simple smoothing. This is a modest percentage reduction (7%) but still worthwhile in absolute terms.

One reason for the improvement of the damped trend over simple smoothing is straightforward. Detailed examination of the results showed that trending series were common for high-value items having a significant impact on inventory investment. Simple smoothing produces a constant forecast for all periods in the procurement leadtime. Thus simple smoothing forecasts are biased in the face of trends whereas damped-trend forecasts are not. Another reason for the better performance of the damped trend is more subtle and also more important. Simple smoothing lags behind jump shifts in the level of demand while the trend component in the damped-trend model structure makes the forecasts more responsive to jump shifts in level. A response lag still exists but it is not as serious.

A separate curve for adaptive smoothing is not shown in Figure 1 since the results were very close to simple smoothing with a fixed parameter. This is not surprising because the same result has occurred in numerous prior studies (for a review see Gardner 1985).

The linear-trend model performed worse than simple smoothing because it consistently overestimated demand. One reason for this performance is that the version of the linear-trend model tested had only one parameter, the discount factor. When a jump shift in demand was encountered, this parameter increased the slope as well as the estimated level of demand. Subsequently the increased trend caused the forecasts to overshoot demand. A two-parameter model such as the Holt linear-trend model would likely perform much better in this data, at least for the time series with jump shifts.

Finally, the naive model produced very poor forecasts. At \$420 million in investment, the naive model's delay time is 41 days, about 58% greater than the delay time for the damped trend. This is reassuring since it indicates there is much to be gained through better forecasting.

The graph in Figure 1 covers a relatively narrow range of values for both delay time and investment. What about values near the origin of the graph? The answer to this question is that there are mathematical limits on both investment and customer service in any inventory. Smaller investment and delay time values than those displayed are infeasible in this inventory. Investment values less than about \$375 million are infeasible because at that level almost all investment is taken up by order quantity stocks. Smaller investment values will not support routine customer demands and will cause infinite

delay times. There also appears to be no way to obtain delay time values less than 26 days. This is the minimum point on the best tradeoff curve, for the damped trend model.

8. Conclusions

Management accepted the results of this study and established plans to implement the damped-trend model during the next few years as part of a general modernization program for computer systems. Gaining management acceptance was not difficult since the forecasting results were presented in terms of management's primary concerns, inventory investment and customer service.

Prior to this research, it was not apparent that the choice of forecasting technique would make any important difference in determining inventory investment and customer service. For example, Armstrong (1986) argues that there are few differences in accuracy among time series models and that there is little to be gained by further research in time series forecasting. Armstrong may be correct when forecast accuracy is judged only by measures such as the MAPE. However, such measures are of little interest to managers, who are concerned instead with whether forecasting will improve decision-making. The major conclusion of this research is that differences in forecast accuracy can be substantial when measured in terms meaningful to managers.

There is no guarantee that the forecasting models tested here will give similar results in other inventory systems. However, the general methodology in this research should be useful in analyzing other systems. The impact of forecasting should be presented to management in the form of a tradeoff curve between inventory investment and customer service. Alternative forecasting models should be evaluated by comparing the position of their tradeoff curves. The aim in improving forecast accuracy should be to shift the optimal tradeoff curve down and to the left, that is to improve customer service and reduce inventory investment.

Tradeoff curves based on alternative forecasting models should have applications in other operational decision problems. In staffing problems, tradeoff curves could be developed to show managers how forecasting affects measures such as the number of personnel required, overtime/undertime, and customer waiting time. In production scheduling, tradeoff curves should be helpful in finding a good balance between production setups and the number of late jobs. Such tradeoff information should be more useful to managers than the mean accuracy measures typically found in the literature of forecasting.

At present there seems to be no alternative to a complex simulation study as a means of developing tradeoff curves based on forecasting models. How can simulation work be justified? A common-sense approach is simply to compute the percentage reduction in inventory investment or operating costs that must be achieved to pay for the simulation study. Often this percentage is trivial. To illustrate, the simulation work reported here incurred costs of no more than \$150,000, mostly in manpower. This figure is a worst-case estimate and includes some fixed costs that would have been incurred regardless of whether the study was conducted. These worst-case costs are less than 0.04% of the typical inventory investment of \$420 million. The simulation study was approved because it was difficult to believe that percentage savings in investment could be less than 0.04%. Actual savings will depend on the tradeoff options selected by management. If target delay time is set at 32 days, investment savings will be about \$30 million. Thus the cost of the study will be about 0.50% of savings.

One important qualification to these comments is that the present value of the savings will be less than \$30 million because it will take some time for inventory investment to be reduced. For most items in the inventory, better forecasting means that less safety stock investment is needed. Unfortunately, management will have to wait for routine

customer demand to reduce stock levels. Another complication is that, for some items with trending demand, stock levels may actually go up. Previous forecasts were biased low due to the use of simple exponential smoothing; in severe cases, more order-quantity stock will have to be ordered immediately to catch up with trends. That is, the additional order-quantity stock required by the revised forecasts will exceed any projected reduction in safety stock. Thus there may be a temporary increase in aggregate inventory investment before the system reaches a steady state at a reduced investment level. The transitional behavior of any inventory is very difficult to predict. However, it seems reasonable to conclude that the costs of this study will be recouped many times over on a present value basis.

References

- ARMSTRONG, J. S., "Research on Forecasting: A Quarter-Century Review, 1960-1984," *Interfaces*, 16 (January-February 1986), 52-66.
- BROWN, R. G., *Advanced Service Parts Inventory Control* (second ed.), Materials Management Systems, Inc., Norwich, Vt, 1982.
- CROSTON, J. D., "Forecasting and Stock Control for Intermittent Demands," *Oper. Res. Quart.*, 23 (1972), 289-303.
- DANCER, R. E. AND C. E. GRAY, "An Empirical Evaluation of Constant and Adaptive Computer Forecasting Models for Inventory Control," *Decision Sci.*, 8 (January 1977), 228-238.
- EPPEN, G. D. AND R. K. MARTIN, "Determining Safety Stock in the Presence of Stochastic Lead Time and Demand," *Management Sci.*, 34 (November 1988), 1380-1390.
- GARDNER, E. S., "Exponential Smoothing: The State of the Art," *J. Forecasting*, 4 (January-March 1985), 1-28.
- AND E. MCKENZIE, "Forecasting Trends in Time Series," *Management Sci.*, 31 (October 1985), 1237-1246.
- HADLEY, G. AND T. M. WHITIN, *Analysis of Inventory Systems*, Prentice-Hall, Englewood Cliffs, NJ, 1963.
- LEE, T. S. AND E. E. ADAM, JR., "Forecasting Error Evaluation in Material Requirements Planning (MRP) Production-Inventory Systems," *Management Sci.*, 32 (September 1986), 1186-1205.
- MAKRIDAKIS, S. ET AL., "The Accuracy of Extrapolation (Time Series) Methods: Results of a Forecasting Competition," *J. Forecasting*, 1 (January-March 1982), 111-153.
- SILVER, E. A. AND R. PETERSON, *Decision Systems for Inventory Management and Production Planning*, John Wiley and Sons, New York, 1987.
- WATSON, R. B., "The Effects of Demand-Forecast Fluctuations on Customer Service and Inventory Cost When Demand Is Lumpy," *J. Oper. Res. Soc.*, 38 (January 1987), 75-82.
- WHYBARK, D. C., "A Comparison of Adaptive Forecasting Techniques," *Log. Transportation Rev.*, 8 (1973), 13-36.