models can be superior to causal models. In general, complexity in forecasting models does not increase predictive accuracy.[1] Therefore, the following discusses three basic time series forecasting methodologies: exponential smoothing, classic time series decomposition, and multiple regression analysis.

Exponential Smoothing

Probably the most useful technique for short-term forecasting is exponential smoothing. It is simple, requires a minimum amount of data to be retained for continued application, has been observed to be the most accurate among competing models in its class, and is self-adapting to fundamental changes in the forecasted data. It is a type of moving average, where the past observations are not given equal weight. Rather, observations that are more recent are weighted more heavily than older ones.

Such a geometric weighting scheme can be reduced to a simple expression involving only the forecast from the most recent period and the actual demand for the current period. Thus, the demand forecast for the next period is given by

$$\text{New forecast} = \alpha(\text{actual demand}) + (1 - \alpha)(\text{previous forecast}) \qquad \textbf{(8-1)}$$

where α is a weighting factor, commonly called the exponential smoothing constant, with values between 0 and 1. Note that the effect of all of history is included in the previous forecast so that only one number needs to be retained at any time to represent demand history.

Example

Suppose that a demand level of 1,000 units was forecasted for the current month. Actual demand for the current month is 950 units. The value of the smoothing constant is = 0.3. The expected value for demand next month, according to Equation (8-1), would be

$$\begin{aligned}\text{New forecast} &= 0.3(950) + 0.7(1{,}000)\\ &= 985 \text{ units}\end{aligned}$$

This forecast becomes the previous forecast when the procedure is repeated one month from now. And so it goes.

For convenience, we can write this "level only" model as

$$F_{t+1} = \alpha A_t + (1 - \alpha)F_t \qquad \textbf{(8-2)}$$

[1]For a summary of these results, see Robin M. Hogarth and Spyros Makridakis, "Forecasting and Planning: An Evaluation," *Management Science*, Vol. 27, No. 2 (February 1981), pp. 115–138.

where

t = current time period
α = exponential smoothing constant
A_t = demand at period t
F_t = forecast for period t
F_{t+1} = forecast for period following t, or the next period

It is identical to Equation (8-1).

Example

The following quarterly data represent a demand time series for a product:

	QUARTER			
	1	2	3	4
Last year	1,200	700	900	1,100
This year	1,400	1,000	$F_3 = ?$	

We wish to forecast the demand for the third quarter of this year. We will assume that $\alpha = 0.2$ and the previous forecast is constructed from the average for the four quarters of last year. Hence, $F_0 = (1{,}200 + 700 + 900 + 1{,}100)/4 = 975$. We begin forecasting the first quarter of this year and carry the computations forward until we reach the third quarter.

The forecast for the first quarter of this year is

$$F_1 = 0.2A_0 + (1 - 0.2)\,F_0 = 0.2(1{,}100) + 0.8(975) = 1{,}000$$

The forecast for the second quarter of this year is

$$F_2 = 0.2A_1 + (1 - 0.2)\,F_1 = 0.2(1{,}400) + 0.8(1{,}000) = 1{,}080$$

The forecast for the third quarter of this year is

$$F_3 = 0.2A_2 + (1 - 0.2)\,F_2 = 0.2(1{,}000) + 0.8(1{,}080) = 1{,}064$$

Summarizing,

	QUARTER			
	1	2	3	4
Last year	1,200	700	900	1,100
This year	1,400	1,000		
Forecast	1,000	1,080	**1,064**	

Choosing the proper value for the exponential smoothing constant requires a degree of judgment. The higher the value of α, the greater is the weight placed on the more recent demand levels. This allows the model to respond more quickly to changes in the time series. However, too high an α value may make the forecast "nervous" and track random variations in the time series rather than the fundamental changes. The lower the α value, the greater is the weight given to demand history in forecasting future demand and the longer is the time lag in responding to fundamental changes in the demand level. Low values provide very "stable" forecasts that are not likely to be heavily influenced by randomness in the time series.

Compromise values for α typically range from 0.01 to 0.3, although higher values may be used for short time periods when anticipated changes will occur, such as a recession, an aggressive but temporary promotional campaign, the discontinuing of some products in the line, or the starting of the forecasting procedure when little or no historical sales results are available. A good rule to follow when searching for an α value is to choose one that will allow the forecast model to track major changes occurring in the time series and average the random fluctuations. This is an α to minimize forecast error.

Correcting for Trend

The basic exponential smoothing model gives good performance when applied to a time series pattern, as shown in Figure 8-1(a), or where the changes in trend and seasonal components are not great. However, when there is a substantial trend or a significant seasonal pattern in the data, the inherent forecast lag in this type of model may give unacceptable forecast error. Fortunately, the model can be expanded to provide better tracking when trend and seasonal elements are significant from the randomness in the data, as shown in Figure 8-1(b) and (c).

Correcting the basic model for forecast time lag due to trend is a simple embellishment to the "level only" model in Equation (8-2). The trend-corrected version of the model is a set of equations that can be stated as

$$S_{t+1} = \alpha A_t + (1 - \alpha)(S_t + T_t) \quad (8\text{-}3)$$

$$T_{t+1} = \beta(S_{t+1} - S_t) + (1 - \beta)T_t \quad (8\text{-}4)$$

$$F_{t+1} = S_{t+1} + T_{t+1} \quad (8\text{-}5)$$

where the additional symbols not previously defined are

F_{t+1} = trend-corrected forecast for period $t+1$
S_t = initial forecast for period t
T_t = trend for period t
β = trend smoothing constant

Example

Recall the previous example having the following data:

	QUARTER			
	1	2	3	4
Last year	1,200	700	900	1,100
This year	1,400	1,000	F_3=?	

We still want to make a forecast for the third period of this year, but with a correction for trend. We will use an arbitrary starting value of $S_t = 975$ (average of last year's demand) and $T_t = 0$ (no trend). The smoothing constant β is assumed 0.3 and α remains the previous value of 0.2. Now begin the forecasting procedure.

The forecast for the first quarter of this year is

$$S_1 = .2(1{,}100) + .8(975 + 0) = 1{,}000$$
$$T_1 = .3(1{,}000 - 975) + .7(0) = 7.5$$
$$F_1 = 1{,}000 + 7.5 = 1{,}007.5$$

Using the results from the first quarter, the forecast for the second quarter of this year is

$$S_2 = .2(1{,}400) + .8(1{,}000 + 7.5) = 1{,}086$$
$$T_2 = .3(1{,}086 - 1{,}000) + .7(7.5) = 31.05$$
$$F_2 = 1{,}086 + 31.05 = 1{,}117.05$$

Using the results from the second quarter, the forecast for the third quarter of this year is

$$S_3 = .2(1{,}000) + .8(1{,}086 + 31.05) = 1{,}093.64$$
$$T_3 = .3(1{,}093.64 - 1{,}086) + .7(31.05) = 24.03$$
$$F_3 = 1{,}093.64 + 24.03 = 1{,}117.67, \text{ or } \mathbf{1{,}118}$$

Summarizing,

	QUARTER			
	1	2	3	4
Last year	1,200	700	900	1,100
This year	1,400	1,000		
Forecast	1,008	1,117	**1,118**	

Correcting for Trend and Seasonality

In addition to trend, the effects of seasonal fluctuations in the time series may also be taken into account. Before applying this type of model, two conditions should be met.

1. There must be a known reason for the periodic peaks and valleys in the demand pattern, and these peaks and valleys should occur at the same time every year.
2. The seasonal variation should be greater than the random variations, or "noise."

If seasonal demand is not stable, significant, and discernible from random variations, then it becomes extremely difficult to develop a model that will accurately predict the direction of the next period's demand. If this is the case, a basic form of the exponential smoothing model, with a high value for the smoothing constant to reduce the effects of lag, may give a lower forecast error than the more complicated model. Caution is needed in model choice.

The level-trend-seasonal model is built around the concept of forecasting the index of actual demand to the trend, and then deseasonalizing it to produce the forecast. The equations for this model are

$$S_{t+1} = \alpha (A_t / I_{t-L}) + (1 - \alpha)(S_t + T_t) \quad \textbf{(8-6)}$$

$$T_{t+1} = \beta(S_{t+1} - S_t) + (1 - \beta)T_t \quad \textbf{(8-7)}$$

$$I_t = \gamma(A_t / S_t) + (1 - \gamma)I_{t-L} \quad \textbf{(8-8)}$$

$$F_{t+1} = (S_{t+1} + T_{t+1})I_{t-L+1} \quad \textbf{(8-9)}$$

where symbols not previously defined are

F_{t+1} = trend and seasonally corrected forecast for period $t + 1$
γ = smoothing constant on the seasonal index
I_t = seasonal index for period t
L = the time period for one full season

Solving this model involves the same procedures as the previous versions. The number of calculations makes it rather impractical to compute forecasts manually. Computer packages, such as the forecast module in LOGWARE,[2] have been written not only to make the forecast, but also to assist the user in setting the initial values to start the forecasting process and determine the best smoothing constants.

Forecast Error Defined

To the extent that the future is not perfectly mirrored by the past, the forecast of future demand will generally be in error to some degree. Since the exponential smoothing forecast is a prediction of the average demand, we seek to project a range within which the actual demand will fall. This requires a statistical forecast.

The error in the forecast refers to how close the forecast comes to the actual demand level. It is properly expressed statistically as a standard deviation, variance, or mean absolute deviation. Historically, the mean absolute deviation (MAD) has been used as the measure of forecast error with reference to exponential smoothing. Early proponents of exponential smoothing may have preferred standard deviation as the proper measure, but accepted the simpler MAD computation because of the limited memory of early computers. Since computers now have adequate memory for the forecasting task, the standard deviation is developed as the measure of forecast error.

[2]Software available with this text.

The forecast error is defined as

$$\text{Forecast error} = \text{actual demand} - \text{forecasted demand} \tag{8-10}$$

Since the forecasted demand is an arithmetic mean value, the sum of the forecast errors over a number of periods should be zero. However, the magnitude of the forecast error can be found by squaring the errors, thus eliminating the canceling of positive and negative errors. The common form of the standard deviation is developed, and it is corrected for the one degree of freedom that is lost in producing the forecast; that is, the α in the "level only" forecast model. The expression for this standard deviation is[3]

$$S_F = \sqrt{\frac{\sum_t (A_t - F_t)^2}{N-1}} \tag{8-11}$$

where

S_F = standard error of the forecast
A_t = actual demand in period t
F_t = forecast for period t
N = number of forecast periods t

The form of the frequency distribution of forecast errors becomes important when making probability statements about the forecast. Two typical generalized forms of the forecast error distribution are shown in Figure 8-3. Assuming that the forecast model is tracking the average of actual demand levels quite well and the variation of actual demand about the forecast is small relative to the forecast level, the normal frequency distribution, or approximations to it, is a likely form to be found in practice. This will especially be the case for the distribution of average forecast errors. The central limit theorem[4] applies, and the normal frequency distribution is the proper distribution form. Where the forecast interval is short, a skewed distribution may result like that shown in Figure 8-3(b).

A way to determine the frequency distribution that applies in any particular situation is through using the chi-square goodness of fit test.[5] Alternately, the following test can be used to select between the normal (symmetrical) distribution and the exponential distribution form as a simple representation of a skewed distribution.

> In a normal distribution, about 2% of the observations exceed a level two standard deviations above the mean. In an exponential distribution, the probability of exceeding the mean by more than 2.75 standard deviations is about

[3]Alternately, the mean squared error and the root mean squared error are popular formulations. They differ as to whether the square root is taken of the sum of the squared errors and whether a correction is made for the degrees of freedom lost. The degrees of freedom lost depends on the number of smoothing constants estimated in the model equations.
[4]For a definition, see any good book on applied statistics, or John Neter, William Wasserman, and G. A. Gilmore, *Applied Statistics* (Boston: Allyn and Bacon, 1988), pp. 262–263.
[5]Ibid.

Figure 8-3
Typical Forecast Error Distributions

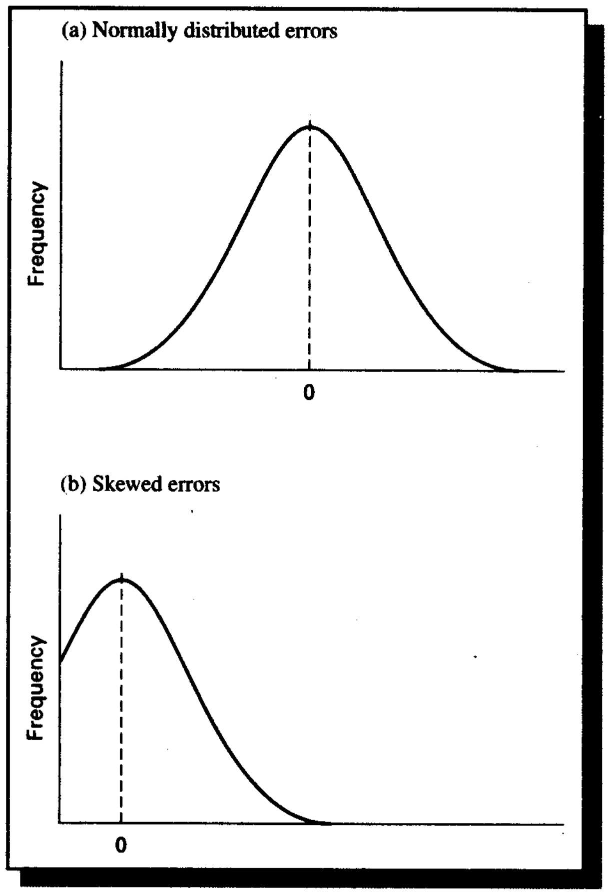

2%. Therefore, if the number of standard deviations it takes to account for all but about 2% of the observations is near 2, a normal distribution should be used. If it is above 2.7, the exponential distribution should be used.[6]

Example

Recall the "level only" forecast that had the following data and results:

	QUARTER			
	1	2	3	4
Last year	1,200	700	900	1,100
This year	1,400	1,000		
Forecast	1,000	1,080	**1,064**	

[6]Robert G. Brown, *Materials Management Systems* (New York: John Wiley & Sons, 1977), p. 146.

Now, let's estimate the standard error of the forecast (S_F) for the two periods (N = 2) for which the forecast has been made and actual demand values are available. Assuming that demand is normally distributed about the forecast, we can develop a 95 percent confidence band around the third quarter forecast. Based on Equation (8-11), we estimate S_F.

$$S_F = \sqrt{\frac{(1,400 - 1,000)^2 + (1,000 - 1,080)^2}{2 - 1}}$$
$$= 407.92$$

The best estimate for the actual demand level (Y) for the third quarter with $z_{@95\%}$ = 1.96 from a normal distribution table (see Appendix A) is

$$Y = F_3 \pm z(S_F)$$
$$= 1,064 \pm 1.96(407.92)$$
$$= 1,064 \pm 800$$

Hence, the 95 percent confidence range for the forecast of actual demand (Y) is

$$264 < Y < 1,864$$

Monitoring Forecast Error

One of the notable advantages of using exponential smoothing for short-term forecasting is its ability to adapt to changing patterns in the time series. How well the model maintains its accuracy is directly related to the smoothing constant value at any point in time. Therefore, sophisticated forecasting procedures involve monitoring the forecast error and making adjustments in the smoothing constant values. If the time series is stable, relatively low values would be selected. During periods of rapid change, high values would be used. By not being limited to single values, the forecast error can be reduced, especially when demand patterns are dynamic.

A popular method for monitoring the forecast error is by means of a tracking signal. The tracking signal is a comparison, usually a ratio, of the current forecast error to an average of past forecast errors. This ratio may be continuously or periodically evaluated. As a result of this computation, the exponential smoothing constants may be recomputed or specified again if the ratio exceeds a specified control limit.

In general, the best smoothing constant values are the ones that minimize the forecast error over time for a stable time series. Adjusting the values as the characteristics of the time series change offers further opportunity to reduce forecast error. Adaptive models that allow the smoothing constants to be revised continuously perform well when the demand time series is changing rapidly, but they do not seem to perform as well during stable periods. Conversely, smoothing constants revised to specified limits offer good performance during stable demand periods and can give remarkably good performance during periods of sudden or rapid change in the series.[7] Flowers has suggested optimum values for these specified smoothing

[7]From computer simulation experiments as reported in D. Clay Whybark, "A Comparison of Adaptive Forecasting Techniques," *Logistics and Transportation Review*, Vol. 8, No. 3 (1972), pp. 13–25.

Figure 8-4
Example of the Performance of a Well-Specified Exponential Smoothing Model

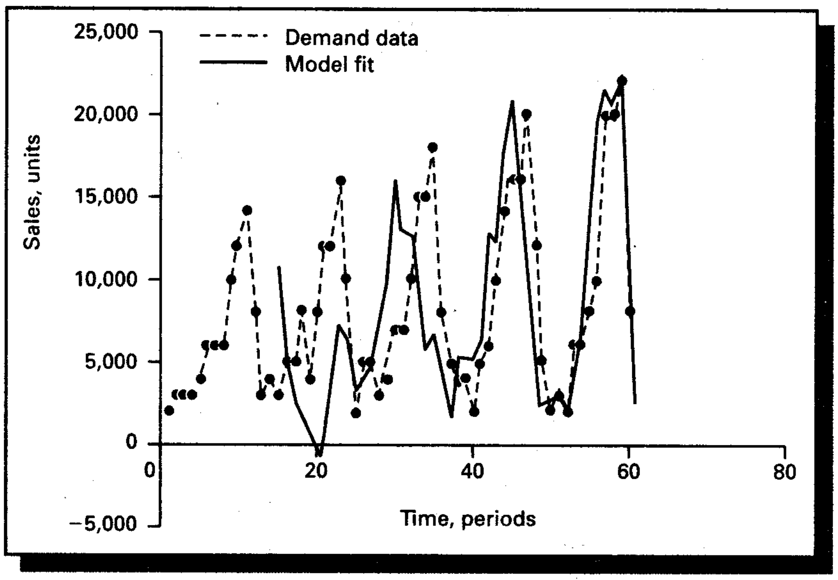

constants.[8] Performance of a well-specified exponential smoothing model should look like the latter periods of the time series shown in Figure 8-4.

Classic Time Series Decomposition

A class of forecasting models that has been useful over the years is that of time series decomposition. These methods include spectral analysis, classic time series analysis, and Fourier series analysis. Classic time series decomposition analysis is discussed here mainly because of its mathematical simplicity and its popularity, and because more elegant methods have not offered increased accuracy.

Classic time series decomposition forecasting is built on the philosophy that a historical sales pattern can be decomposed into four categories: trend, seasonal variation, cyclical variation, and residual, or random, variation. Trend represents the long-term movement in sales caused by factors such as changes in population, changes in marketing performance of the firm, and fundamental changes in market acceptance of the firm's products and services. Seasonal variation refers to the regular hills and valleys in the time series that usually repeat every 12 months. The forces causing this regular variation include climatic changes, buying patterns pegged to calendar dates, and the availability of goods. Cyclical variation is the long-term (more than one year) undulations in the demand pattern. Residual, or random, variation is that portion of total sales that is unaccounted for by trend, seasonal, or

[8] A. Dale Flowers, "A Simulation Study of Smoothing Constant Limits for an Adaptive Forecasting System," *Journal of Operations Management*, Vol. 1, No. 2 (November 1980), pp. 85–94.

cyclical components. If the time series is well described by the other three components, the residual variation should be random.

Classic time series analysis combines each type of sales variation in the following way:

$$F = T \times S \times C \times R \tag{8-12}$$

where

F = demand forecast (units or $)
T = trend level (units or $)
S = seasonal index
C = cyclical index
R = residual index

In practice, the model is often reduced to only trend and seasonal components. This is done because a well-specified model has a residual index value (R) of 1.0 and thus does not affect the forecast, and because it is difficult in many cases to decompose cyclical variation from random variation. Treating the cyclical index (C) as equal to 1.0 is not as serious as it first seems because the model is usually updated when new data become available. The effect of cyclical variation tends to be compensated for in the updating process.[9]

The trend value (T) in the model may be determined by several methods, such as fitting a line "by eye," using some form of the moving average, or using the method of least squares.

The popular least squares method is a mathematical technique that minimizes the sum of the squared differences between the actual data and the proposed trend line. A least squares line can be found for any trend-line form, whether linear or nonlinear. The mathematical expression for a linear trend line is $T = a + bt$, where t is time, T is the average demand level, or trend, and a and b are coefficients to be determined for the particular time series. These coefficients are found by

$$b = \frac{\sum D_t(t) - N(\bar{D})(\bar{t})}{\sum t^2 - N\bar{t}^2} \tag{8-13}$$

and

$$a = \bar{D} - b\bar{t} \tag{8-14}$$

where

N = the number of observations used in the development of the trend line
D_t = the actual demand in time period t
$\bar{D}$ = average demand for N time periods
$\bar{t}$ = average of t over N time periods

[9]The model is sometimes expressed in an additive form of $F = T + S + C + R$.

Nonlinear trend lines are more complex mathematically and are not discussed here.[10]

The seasonality component of the model is represented by an index value that changes for each period being forecasted. This index is a ratio of the actual demand in a given time period to the average demand. The average demand may be represented by a single average of the actual demand over a specified period, usually one year; a moving average; or the trend line. Inasmuch as the trend line was previously discussed, it will be used as the seasonal index base. Therefore,

$$S_t = D_t / T_t \tag{8-15}$$

where

S_t = seasonal index in time period t
T_t = trend value determined from $T = a + bt$

Finally, the forecast is made for time period t in the future as follows:

$$F_t = (T_t)(S_{t-L}) \tag{8-16}$$

where

F_t = the forecasted demand in time period t
L = number of periods in the seasonal cycle

These ideas are best illustrated by an example.

Example

A manufacturer of young women's clothing had to make purchase quantity decisions and set production and logistics schedules based on forecasts of market bookings (sales). Five seasons of the year were specified for planning and promotional purposes—summer, trans-season, fall, holiday, and spring. Sales data for approximately two and one-half years were obtained (see Table 8-2). A forecast was needed for two seasons ahead of the current accounting period to ensure adequate purchasing and production lead time. In this case, the forecast period was the holiday season, even though the sales for the intervening fall period were not yet known.

The first task was to find the trend line using Equations (8-13) and (8-14). Assuming a straight-line trend, the b coefficient was

$$b = \frac{1,218,217 - (12)(14,726.92)(6.5)}{650 - (12)(6.5)^2} = 486.13$$

and the a coefficient follows as

$$\begin{aligned} a &= 14,726.92 - 486.13(6.5) \\ &= 11,567.08 \end{aligned}$$

[10]A discussion of nonlinear trend lines can be found in John Neter, William Wasserman, and Michael H. Kutner, *Applied Linear Regression Models* (Homewood, IL: Irwin, 1983), Chapter 14.

Sales Period	(1) Time Period (t)	(2) Sales (D_t) ($000s)	(3) $D_t \times t$	(4) t^2	(5) Trend Value (T_t)	(6) = (2)/(5) Seasonal Index	Forecast ($000s)
Summer	1	$ 9,458	9,458	1	$12,053	0.78	
Trans-season	2	11,542	23,084	4	12,539	0.92	
Fall	3	14,489	43,467	9	13,025	1.11	
Holiday	4	15,754	63,016	16	13,512	1.17	
Spring	5	17,269	86,345	25	13,998	1.23	
Summer	6	11,514	69,084	36	14,484	0.79	
Trans-season	7	12,623	88,361	49	14,970	0.84	
Fall	8	16,086	128,688	64	15,456	1.04	
Holiday	9	18,098	162,882	81	15,942	1.14	
Spring	10	21,030	210,300	100	16,428	1.28	
Summer	11	12,788	140,668	121	16,915	0.76	
Trans-season	12	16,072	192,864	144	17,401	0.92	
Fall	13	?			17,887[a]		$18,602[b]
Holiday	14	?			18,373		20,945
Total	78	$176,723	1,218,217	650			

$N = 12$

$\sum D_t x t = 1,218,217$

$\sum t^2 = 650$

$\bar{D} = (176,723 / 12) = \$14,726.92$

$\bar{t} = (78 / 12) = 6.5$

[a] Forecasted values. For example, $T_t = 11,567.08 + 486.13(13) = 17,887$.

[b] $F_{13} = T_{13} \times S_{13\text{-}5}$ or $18,602 = 17,887 \times 1.04$.

Table 8-2 Time Series Forecast from a Clothing Manufacturer's Sales Data

Therefore, the trend equation was

$$T_t = 11,567.08 + 486.13t$$

From this trend-line equation, the values were projected by substituting into the previous equation each value of t; see column 5 in Table 8-2.

The seasonal indices were computed according to Equation (8-15) and are displayed in column 6 of Table 8-2. For forecasting purposes, the most recently available season was used, mainly because the indices did not vary greatly from year to year. If this were not the case, the indices for several years might be averaged.

The forecast for the holiday season (period 14) was

$$Y_{14} = [11,567.08 + 486.13(14)] \times 1.14 = \$20,945 \text{ (in \$000s)}$$

The forecast for the fall period (period 13) was made in a similar manner.

Multiple Regression Analysis

In the forecasting models discussed thus far, time is the only variable that has been considered. To the extent that other variables show a relationship to demand, they may also be included in a model to forecast sales. Multiple regression analysis is a statistical technique that helps to determine the degree of association between a number of selected variables and demand. From this analysis, a model is developed that may use more than one variable to predict future demand. Information about the predictor (independent) variables is then converted by the regression equation to give a demand forecast.

Example

Reconsider the clothing manufacturer's problem discussed previously. An alternative approach to forecasting over the two-season interval was to use a regression model, preferably where the independent variables "lead" the demand variable in time. This permitted data to be obtained about independent variables in advance of the forecast period. One such forecasting equation was developed for the summer selling season:

$$F = -3{,}016 + 1{,}211X_1 + 5.75X_2 + 109X_3 \tag{8-17}$$

where

F = estimated average summer season sales (in thousands of dollars)
X_1 = time in years (1991 = 1)
X_2 = number of accounts of purchasing during the season (from advanced bookings)
X_3 = monthly net change in consumer installment debt (percent)

The model explained 99 percent ($R^2 = 0.99$) of the total variation in demand and was statistically significant at the 5 percent level. The equation was deemed an accurate predictor of demand. For example, the actual sales for the summer season of 1996 were \$20,750,000. The model inputs for 1996 were $X_1 = 6$, $X_2 = 2{,}732$, and $X_3 = 8.63$, and, when substituted in Equation (8-17), gave a sales forecast of \$20.9 million, or \$20,900,000.

Although a reasonable knowledge of statistical methodology is required to construct such a model, computer software, such as SPSS[11] and BMDP,[12] for performing a regression analysis is readily available for both microcomputer and mainframe computer installations. These programs perform the necessary computations for fitting an ordinary least squares line to the data and providing statistical information to evaluate the fit. However, care should be exercised in the use of these statistical packages since they alone cannot guarantee a valid model, that is, one that is free of specification and statistical problems.[13]

---

[11]A product of SPSS, Inc., 444 N. Michigan Ave., Chicago, IL.
[12]A product of BMDP Statistical Software, 1964 Westwood Blvd., Los Angeles, CA.
[13]For a discussion of these problems, see Marija J. Norusis, SPSS/PC+ (Chicago: SPSS, Inc., 1986), Chapter 17; and Neter, Wasserman, and Kutner, *Applied Linear Regression Models*, op. cit.

SPECIAL PREDICTION PROBLEMS FOR LOGISTICIANS

Special problems are sometimes encountered when attempting to predict requirements. These problem areas are start-up, lumpy demand, regional forecasting, and forecast error. Although not all of these problems are necessarily unique to logistics, they are of great concern to the logistician in accurately determining requirements.

Start-Up

The logistician is often faced with the problem of predicting requirement levels for products or services for which there is not enough history to start the forecasting process. New product or service introductions and the need to provide logistical support for them create the common start-up conditions. Several approaches have been used during this early forecasting period.

First, put the initial estimation in the hands of the marketing personnel until a sales history begins to develop. They will know best the level of promotional effort, early customer response, and expected customer acceptance. Once a reasonable demand history has been generated, say, in six months, the established forecasting methods can be used with some confidence.

Second, an estimate may be made from the demand pattern of similar products in the line. Although many companies turn their product line on the average of once every five years, few products are radically new. They often represent changes in size, style, or revision of existing products. Therefore, demand patterns previously experienced may provide insight and a basis for estimating initial demand for new products.

Third, if the exponential smoothing model is used for forecasting, the exponential smoothing constant may be set at a high level (0.5 or higher) during the initial forecasting period. It will be reduced to a normal level once an adequate demand history has been generated.

Lumpy Demand

The problem of lumpy, or irregular, demand has been described previously and illustrated in Figure 8-2. It represents the condition where there is so much random variation in the demand pattern that trend and seasonal patterns can be obscured. The lumpy demand condition occurs when two or three times the standard deviation of the historical data exceed the forecast of the best model that can be fit to the time series. The lumpy demand pattern occurs frequently for a variety of reasons: The demand pattern is dominated by large, infrequent customer orders; demand may be derived from the demand of other products or services; seasonal peaking may not have been taken into account; and the demand pattern may be a result of exceptional data, outliers, or unusual conditions.

Lumpy demand patterns are, by nature, difficult to predict accurately by mathematical methods due to the wide variability in the time series; however, some suggestions on how to treat them can be offered. First, look for obvious reasons for the lumpiness and use them to produce the forecast. Separate the forecasting of lumpy demand products from those showing a regular pattern and use forecasting methods tailored to each.

Example

A chemical manufacturer had a product in its line used for cleaning apples at harvest time. Depending on the size of the apple crop, sales of the product could vary considerably from year to year. Exponential smoothing was used to forecast this as well as other products in the line. Inventory levels at warehouses that were set based on this forecast were typically either short or greatly in excess of reasonable needs. Grouping this lumpy demand product in with those having regular patterns did not permit the company to take advantage of the basic reasons why the demand level was changing throughout the year.

Second, do not react quickly to changes in the demand pattern for such products or services if no assignable causes can be found for the demand shifts. Rather, use a simple, stable forecasting method that does not react rapidly to change, such as a basic exponential smoothing model with a low smoothing constant value or a regression model that is refitted no more frequently than on an annual basis.

Third, because lumpy demand frequently occurs in low demand items, forecast accuracy may not be an overriding issue. If the forecast is used to establish inventory levels, carrying a little more inventory to compensate for forecast inaccuracy may be more economical than attempting to manage the forecast carefully.

Regional Forecasting

Although most of the discussion in this chapter has been directed toward forecasting time-related demand, geographic aggregation or disaggregation of the forecast is also of concern. That is, the logistician must decide whether to take a forecast of total demand and apportion it by regions, such as by plant or warehouse territories, or to forecast each region separately. Achieving the greatest accuracy in the forecast at the regional level is the concern. Forecasting all demands simultaneously very often will be more accurate than the sum of individual regional forecasts. If this is so, apportioning the aggregated forecast to the individual regions may preserve enough accuracy to give better results than individual forecasting. Research on the subject has not provided a definitive answer about which approach is better. Hence, the logistician should be aware of both possibilities and compare the methods in his or her particular situation.

Forecast Error

The final concern is to make the most of the available forecasting techniques. The discussion so far has centered on the use of individual models and methods. In practice, no single forecasting model may be best at all times. Rather, combining results of several models may give more stable and accurate forecasts.[14]

[14]M. J. Lawrence, R. H. Edmundson, and M. J. O'Connor, "The Accuracy of Combining Judgemental and Statistical Forecasts," *Management Science*, Vol. 32, No. 12 (December 1986), pp. 1521–1532, Essam Mahmoud, "Accuracy in Forecasting: A Survey," *Journal of Forecasting* (April-June 1984), p. 139; Spyros Makridakis and Robert L. Winkler, Average of Forecasts: Some Empirical Results," *Management Science*, (September 1983), p. 987; and Victor Zarnowitz, "The Accuracy of Individual and Group Forecasts from Business Outlook Surveys," *Journal of Forecasting* (January–March 1984), p. 10.

The following example shows combining multiple forecast methods according to their forecast error. This generally works well for long-term forecasts. For short-term forecasts, equally weighted forecasts have been shown to be particularly robust, and they give greater forecast accuracy than unequally weighted ones.[15]

Example

Reconsider the clothing manufacturer[15] forecasting problem. Because there were five selling seasons, there was no guarantee that one method of forecasting would consistently be superior throughout all seasons. In fact, four methods were used. There was a regression model (R) that predicted sales based on the two variables: (1) number of accounts; and (2) the change in consumer debt. Two versions (ES_1, ES_2) of an exponential smoothing model were used. The fourth model was the company's internal forecast based on managerial judgment and experience (MJ). The average forecast error realized by using each method during the different selling seasons is shown in Figure 8-5.

One way to combine the information from each forecast model is to weight the results according to the average historical error that they produced. In this way, no model results would be eliminated or would there be total reliance on the model result that happened to appear best historically.

To illustrate the weighting scheme, consider the fall selling season results shown in Figure 8-5. The average error for each model was $MJ = 9.0$ percent, $R = 0.7$ percent, $ES_1 = 1.2$ percent, and $ES_2 = 8.4$ percent. The weights should be inversely proportional to the forecast error and in the same ratio to their respective percentages. Table 8-3 shows the computation of the weighting factors.

Finally, given the forecast results of each model and the weighting factors, a weighted average forecast can be calculated, as shown in Table 8-4. The final forecast value for the fall selling season is \$20,208,000 and represents inputs from a number of forecasting sources.

An expert system approach to combining forecast methods has shown encouraging results, especially for short-term forecast periods of less than one year. Known as rule-based forecasting, several time series methods are combined and rules derived from forecasting experts are applied to the input data and to the model application. The rules are various IF-THEN statements that guide actions to improve the forecast. Collopy and Armstrong develop 99 such rules, a few examples of which are

- IF an observation is an outlier, THEN set the observation equal to two standard deviations from the mean.
- For an exponential smoothing model, IF alpha (smoothing constant) is calculated to be greater than 0.7, THEN use 0.7. IF alpha is less than 0.2, THEN use 0.2.

[15]Fred Collopy and J. Scott Armstrong, "Rule-Based Forecasting: Development and Validation of an Expert Systems Approach to Combining Time Series Extrapolations," *Management Science*, Vol. 38, No. 10 (1992), pp. 1394–1414.

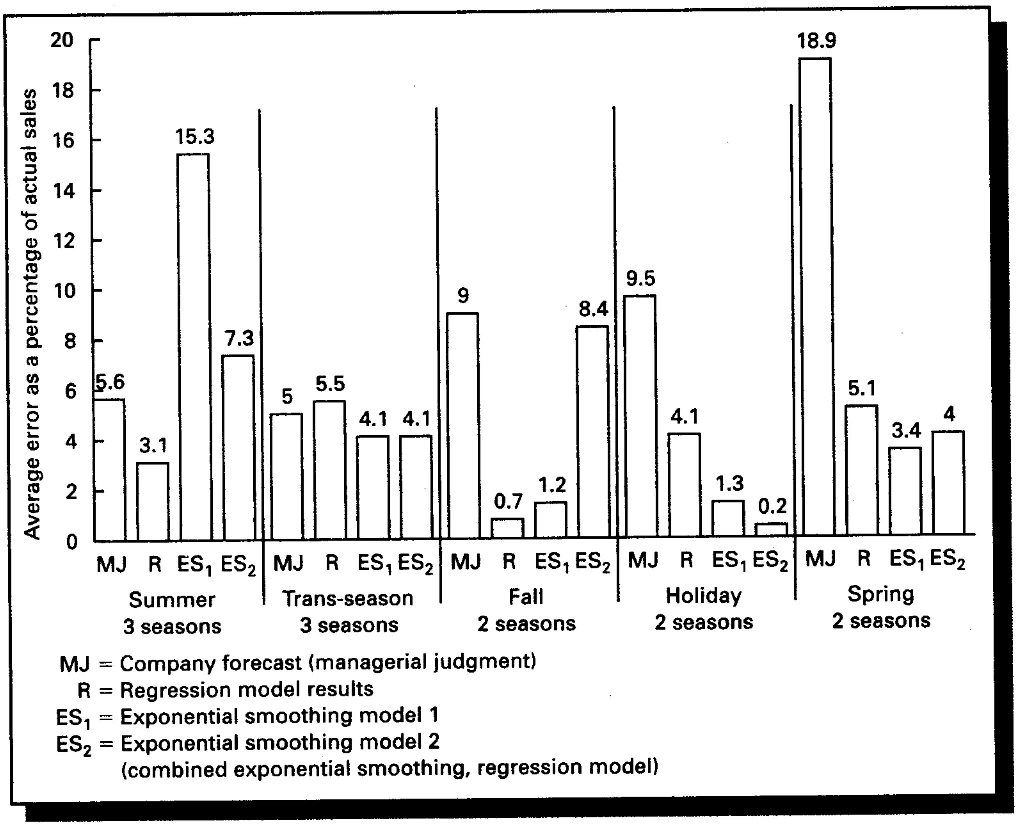

Figure 8-5 Forecast Error for Four Forecasting Techniques Applied to the Sales of a Clothing Manufacturer

- IF early data are irrelevant, THEN delete these data.
- IF observations are judged irregular based on knowledge of the application, THEN adjust the observations prior to analysis to remove their short-term effects.[16]

Table 8-3 Calculation of Model Weights

Model Type	(1) Forecast Error	(2) = (1)/19.3 Proportion of Total Error	(3) = 1/(2) Inverse of Error Proportion	(4) = (3)/48.09 Model Weights
MJ	9.0	0.466	2.15	0.04
R	0.7	0.036	27.77	0.58
ES_1	1.2	0.063	15.87	0.33
ES_2	8.4	0.435	2.30	0.05
	19.3	1.000	48.09	1.00

[16]Ibid.

Forecast Type	Model Forecast	Weighting Factor[a]	Weighted Proportion[b]
Regression model *(R)*	$20,367,000	0.58	$11,813,000
Exponential smoothing (ES_1)	20,400,000	0.33	6,732,000
Combined exponential smoothing–regression model (ES_2)	17,660,000	0.05	883,000
Managerial judgment *(MJ)*	19,500,000	0.04	780,000
Weighted average forecast			$20,208,000

[a]From Table 8-3
[b]Model forecast multiplied by the weighting factor

Table 8-4 Weighted Average Fall Selling Season Forecast Using Several Forecasting Techniques

Applying these and additional rules to multiple forecasting methods such as random walk (using the most recent observation), exponential smoothing, and regression, significant reduction in forecast error can be achieved. This is especially true when the time series shows significant trends, low uncertainty, and stability, and the forecaster has good knowledge of the application. Each of the methods is weighted equally.

Collaborative Forecasting

The forecasting methods that have been illustrated to this point in the chapter work best when demand does not show great variability. However, lumpy, highly uncertain, and dynamic demand brought about by such factors as promotions, few buyers purchasing in large qualities, seasonal/cyclical buying, and demand created by "acts of God" creates a special problem. Although some guidelines have already been offered to treat the lumpy demand case, collaborative forecasting is being suggested as an improved approach, especially for channel planning of the business processes. It is based on the premise that "two heads are better than one." That is, multiple parties have the likelihood of producing more accurate forecasts than a single party.

Collaborative forecasting refers to the development of forecasts using the inputs from multiple participants, whether they are from various functional areas within a single firm (marketing, operations, logistics, finance, purchasing, etc.) or from the various members in a supply channel such as vendors, carriers, and buyers. The goal is to reduce forecast error. This can best be accomplished when each party brings unique perspective to the forecasting process. Buyers or marketing personnel may be close to the final consumer and have the best "feel" of end demand. On the other hand, vendors or purchasing personnel may be attuned to supply shortages and capacity limitations that may place a cap on demand or that may ultimately affect product price that in turn affects price and product demand levels. Transportation personnel or carriers may be able to predict delivery times that affect customer service and sales.

Forecasting by means of collaboration requires administering a team of diverse parties with all the complications inherent in such a process. However, the key administrative steps can be identified, many of which are listed here:

- Someone should champion the process and provide the necessary communication and group meeting schedules.
- The kinds of information needed in the forecast and the processes for collecting them should be identified, including the timing, quantities, and person(s) responsible.
- Methods for processing the information from multiple sources, types, and formats should be established as well as the weights to be used for combining and reconciling forecasts from the multiple parties.
- Methods are needed for translating the final forecast into the form needed by each party, such as sales, shipments, and SKUs in total and by customer account, service territory, and so on.
- A process should be available for revising and updating the forecast on a real-time basis.
- Metrics should be established for appraising the forecast and determining whether collaborative forecasting is an improvement over traditional approaches.
- The benefits to each party of collaborative forecasting should be obvious and real.

Collaborative forecasting is a complex process that is inherently unstable, that is, forecasting will have the tendency to fall back to the individual members making their own predictions. Successful collaboration requiring sharing, coordination, compromise, consideration, commitment, and understanding is not easy to realize. However, the benefits of more accurate forecasting in the most difficult forecasting situations as well as the benefits of improved communication interfunctionally and interorganizationally may justify the extra effort required to operate in a coalition.

Applications

Although the supporting software for collaborative forecasting is new and under constant revision, several notable companies have reported early successes as they experiment with the collaborative forecasting approach. Heineken USA (the brewer) has close to 100 of its independent beer distributors submit forecasts electronically to Heineken USA's White Plains, New York, office using a third-party software product. Involving about 40 percent of Heineken's total volume, this setup has cut order-cycle times from 12 weeks to just four or five.[17]

Ace Hardware, a $2.8 billion hardware retailer, experimented with joint forecasting for stock replenishment with Manco, a supplier of tape, glues, and adhesives. Using Web-based software and the Internet, Manco can gain access to Ace's database. Ace presents Manco with its forecast for items through a Web browser screen, but Manco has the opportunity to change the forecast before it brings that forecast into its production planning system. Ace and Manco look at the same screens in real

[17]John Verity, "Collaborative Forecasting: Vision Quest" *Computerworld Commerce*, Vol. 31, No. 45 (November 1997), pp. 12–14.

time and exchange messages before coming to a forecast consensus. Forecast accuracy is reviewed on a monthly basis. In the past, forecast accuracy was 20 percent over or under the actual demand. Now, it is less than 10 percent.[18]

Bajaj Auto is one of the top manufacturers of two-and three-wheeled vehicles in India. It has introduced SAP R/3 Enterprise Resource Planning (ERP) software to improve forecasting and sales planning, product planning, reducing inventory, and scheduling. Its wide area network (WAN) covers three plants, 14 regional offices, 12 warehouses, and two liaison offices in Mumbai and Delhi. Its SAP Enterprise Portal links its dealers and suppliers. The portal speeds up processing of incoming goods and reduces waiting time at the plant's gate. The dealers and sales employees get relevant information of their respective territory from the SAPR/3 ERP database. This is helping the company to beat stagnation and competition.[19] Major FMCG companies in India have programs in place to link their distributors to the firm via Internet- or Intranet-based systems. For example, HLL uses RSNet, Marico MINet, and Dabur has its Daburnet. As a result, forecasting accuracies are increasing.[20]

Shoppers' Stop has been an early adopter of information technology (IT) and has made significant IT investments in supply chain, demand chain, store operations, existing and new store projects, and customer relationship management. The retail chain's IT solutions have been sourced from global companies such as JD Armstrong (JDA), Oracle, IBM, Microsoft, AutoCAD, Dell, and Cisco. The company's investments in E3, an advanced replenishment application from JDA, for its HyperCITY business unit has helped the company manage high inventory turnovers and provide better forecasts for the inventory requirements.[21] A team in a retail chain specifically does range planning too, where it also decides what to buy, how many options to buy and how much to buy of each option, and when to put it on sale. This is derived from the financial plans formulated for each season for each division and department. The team with the help of software fine tunes the range for each store.[22]

Flexibility and Quick Response—An Alternative to Forecasting

The sales of some products or services are so unpredictable that using the types of forecasting methods already described results in such a high potential forecasting error that it makes them impractical. Lumpy demand patterns are an example, and so an alternative is needed. Recognizing that there can be no better forecast than to wait until customer demand materializes is a basis for responding accurately to demand. If the

[18]James A. Cooke, "Why ACE Is Becoming THE PLACE," *Logistics Management & Distribution Report*, Vol. 41, No. 3 (March 2002), pp. 32–36.
[19]*INDUSTRY 2.0* (January 31, 2006), p. 89.
[20]Samir K. Srivastava, "Logistics and Supply Chain Practices in India," *Vision: The Journal of Business Perspective*, Vol. 10, No. 3 (2006), pp. 69–79.
[21]*Benefit* (November 30, 2005), p. 53.
[22]Samir K. Srivastava, "RFID Technology in Retail: Indian Scenario," *International Journal of Manufacturing Technology and Management* (special issue on Connective Technologies and Their Impact on Manufacturing and Logistics), Vol. 10, No. 1 (2007), pp. 71–91.

processes of the supply chain can be made to be flexible and to respond quickly to demand requirements, there is little need for forecasting. After all, statistical forecasting assumes the usual properties that the observations in the time series are random, independent, and each observation is a small portion of the total. When there is a time lag in matching supply to demand, forecasting serves to set the levels of production, purchasing, and inventories so that supply is available when demand occurs. Changing the nature of the supply chain so that the processes can respond flexibly and efficiently to the specific requirements of each customer request, and to do this almost instantaneously, makes forecasting unnecessary. Where demand is very unpredictable, this alternative approach should be explored. However, in many cases where demand is "regular," supplying to a forecasted demand remains the preferred choice.

Example

National Bicycle found that sports bikes—ten-speed and mountain bikes—had become fashion items, sold in part because of their bright, intricate color patterns that changed every year. National's inability to predict which color patterns would be hot each year was causing it to overproduce some colors and underproduce others, generating huge losses. To circumvent this forecasting problem, the company created a custom-ordering system by which customers were measured for their ideal frame dimensions and invited to choose their favorite color pattern from a wide selection. Their ideal bike was then created in the company's remarkably flexible plant in Kashiwara and delivered to their door two weeks later.[23]

Pantaloon Retail (India) has taken up radio frequency identification (RFID) pilot project at one of its warehouses in Tarapur using 1,000 RFID tags. The company selected shirts and trousers of John Miller formals and casuals for its RFID pilot project. The RFID application has been developed by Wipro and is integrated with Oracle database 10g and middleware along with implementation of the RFID hardware. RFID integrates with the current IT infrastructure of the company. Through the use of RFID, recording of data was smoother at the inward and outward terminals, which enabled the company to save time and gain accuracy. The company was able to save 80 percent time in inward warehouse processing and 12 percent in outward processing. It plans to extend the application to production routing and scheduling, product recall and returns, and real time data for category managers for effective forecasting.[24]

CONCLUDING COMMENTS

The logistician frequently finds it necessary to provide his or her own forecasts of demand, lead times, prices, and costs for use in strategic and operational planning and control. Many times, the long-term forecasts needed are provided from outside the logistics function or are only partially the responsibility of the logistician. This is

[23]Marshall L. Fisher, Janice H. Hammond, Walter R. Obermeyer, and Ananth Raman, "Making Supply Meet Demand in an Uncertain World," *Harvard Business Review*, Vol. 72 (May–June 1994), pp. 83–89+.
[24]*Financial Express* (August 22, 2005).

particularly true for strategic planning. Therefore, this chapter has focused on short to medium-term forecasting methods that the logistician is most likely to use. Within this period, those techniques that have proven to have the greatest utility are discussed—exponential smoothing, classic time series decomposition, and multiple regression.

Several special problems of producing a forecast are briefly discussed. These include starting the forecast with little or no previous information about the time series; dealing with lumpy or irregular time series patterns; forecasting demand within geographic segments; and using forecasting models in combination to reduce forecast error.

The logistics manager should also be aware of an alternative to forecasting that may be needed when demand is so unpredictable that forecasted results are unsatisfactory. By designing the supply chain for flexibility and quick response, supply can meet demand as it occurs and forecasting may not be needed at all.

QUESTIONS

A number of the problems in this chapter can be solved, or partially solved, with the aid of computer software. The software package that is most important in this chapter is FORECAST (*F*) in LOGWARE. The CD icon will appear where the FORECAST software is appropriate. A database has been prepared for the World Oil[25] case study. In general, the problems may be solved manually.

1. Why, and to what extent, is the logistician interested in demand forecasting? How do you suppose the interest might be different if the logistician were associated with
   a. a food manufacturer?
   b. an aircraft producer?
   c. a large retail chain?
   d. a hospital?
2. Give illustrations of
   a. spatial versus temporal demand
   b. lumpy versus regular demand
   c. derived versus independent demand
3. Contrast qualitative, historical projection, and causal models for forecasting. What strengths do you see in each type? How might the logistician use each? Categorize the models in Table 8-1 into these three basic types.
4. The Ace Trucking Company must determine the number of drivers and trucks to have available on a weekly basis. The standard schedule is to send drivers over the pickup and delivery route on Monday and return them to the originating point on Friday. The trucking requirements can be determined from the total volume to be moved for the week; however, they must be known a week in advance for planning purposes. The volume for the last ten weeks is given here:

[25]A database for this case study has been prepared in LOGWARE.

Week	Volume	Week	Volume
10 weeks ago	2,056,000	5 weeks ago	2,268,000
9	2,349,000	4	2,653,000
8	1,895,000	3	2,039,000
7	1,514,000	2	2,399,000
6	1,194,000	1 (this week)	2,508,000

a. Using the simplest (level only) exponential smoothing model, predict the expected volume for the next week. [*Note*: You will need to estimate an exponential smoothing constant (α) that will minimize the forecast error. Use the four oldest weeks of data to start the forecast process, that is, find F_0, and search for α in increments of 0.1.]
b. Estimate the forecast error (S_F). Use the last six weekly periods.
c. Find the range over which the actual volume is likely to vary. (*Hint*: You must compute a statistical confidence band here. Assume a 95-percent confidence band and a normal distribution of requirements.)

5. Suppose the data in problem 4 were given as follows:

Week	Volume	Week	Volume
10 weeks ago	1,567,000	5 weeks ago	2,056,000
9	1,709,000	4	2,088,000
8	1,651,000	3	1,970,000
7	1,778,000	2	1,925,000
6	1,897,000	1 (this week)	2,003,000

a. Using the trend-corrected version of the exponential smoothing model, with $\alpha = \beta = 0.2$, forecast next week's volume.
b. Estimate the error in the above forecast (S_F). Use the last six weekly periods.
c. Construct a 95-percent confidence band on the forecast assuming a normal distribution of requirements.

6. The High-Volt Electric Company has a difficult time predicting the quarterly sales for its room air conditioner line due to the substantial seasonality in product sales. Quarterly sales data for the last three years are shown as follows:

LAST YEAR		TWO YEARS AGO		THREE YEARS AGO	
Quarter	Units	Quarter	Units	Quarter	Units
1	34,000	1	30,000	1	27,000
2	82,000	2	73,000	2	70,000
3	51,000	3	48,000	3	41,000
4	16,000	4	15,000	4	13,000

a. Determine the best straight-line trend using simple regression analysis.
b. Determine the seasonal indices for each quarter using the trend line values in your seasonal index computations.
c. By means of classic time series decomposition, forecast the sales for the next four quarters.

7. The materials manager at Metropolitan Hospitals must plan for inventories at three hospital locations within the region. His plan is to allocate stock to these locations. It is necessary to predict sales in order to have a basis for stock allocation. The manager wonders whether it would be more accurate to generate a forecast for each hospital or to generate one forecast from the aggregated data and apportion it to each region. (The more accurate the forecast for each region, the lower will be the inventories.)

   To test the idea, the manager assembled the following monthly usage data for a particular syringe over the last year:

	Region 1	Region 2	Region 3	Combined
Jan.	236	421	319	976
Feb.	216	407	295	918
Mar.	197	394	305	896
Apr.	247	389	287	923
May	256	403	300	959
June	221	410	295	926
July	204	427	290	921
Aug.	200	386	285	871
Sept.	185	375	280	840
Oct.	199	389	293	881
Nov.	214	401	305	920
Dec.	257	446	337	1,040
Total	2,632	4,848	3,591	11,071

   If the manager were to use simple (level only) exponential smoothing with $\alpha = 0.2$, which approach should he use? Why? [*Hint*: Find the initial forecast (F_0) by averaging the first four values in each series. Also, recall the law of variances where $S_T^2 = S_{E_1}^2 + S_{E_2}^2 + S_{E_3}^2$ and compare total errors of the forecast.]

8. A Texas manufacturer of concrete pipe and other precast concrete products for highway, farm, and commercial construction wished to project sales for improved planning of production and logistics operations. A number of variables were thought to affect sales—time, population, housing starts, construction employment, number of residential units, highway budget projections, number of farms, commercial structure construction permits, and number of competing firms in the state. A multiple regression analysis showed three variables to be key to projecting sales: population, construction employment, and construction permits of the previous year.
   a. Do you believe that there is a cause and effect relationship between these variables and company sales?
   b. Are there other important variables that should be considered that were not in the original list?

9. The purchasing agent for a hospital has collected data over the last five years on average monthly unit prices for a commonly used surgical item.

	Last Year	2 Years Ago	3 Years Ago	4 Years Ago	5 Years Ago
Jan.	210	215	211	187	201
Feb.	223	225	210	196	205
Mar.	204	230	214	195	235
Apr.	244	214	208	246	243
May	274	276	276	266	250
June	246	261	269	228	234
July	237	250	265	257	256
Aug.	267	248	253	233	231
Sept.	212	229	244	227	229
Oct.	211	221	202	188	185
Nov.	188	209	221	195	187
Dec.	188	214	210	191	189
Total	2,704	2,792	2,783	2,609	2,645

She believes that an accurate forecast of price would help improve the timing of her buying. Using the oldest four years of data as the base data and saving the most recent data (last year) for checking the forecast accuracy, do the following:

a. Plot the data on a graph. What important observations can you make about the data that would be useful to forecasting?
b. Construct a classic time series decomposition forecasting model based on two full years of data (years 2 and 3) and compute the error of the forecast (S_F) for the last full year. *Hint*: Use

$$S_F = \sqrt{\frac{\Sigma(A_t - F_t)^2}{N - 2}}$$

c. Construct an exponential smoothing model ($\alpha = 0.14$, $\beta = 0.01$, and $\gamma = 0.7$) with level, trend, and seasonal components, and compute S_F for the last year.
d. Create a weighted average model that combines both model types.

10. Hudson Paper Company is a small family-owned business that purchases paper in rolls from large mills. It then cuts and prints the paper into a variety of products, such as bags and wrapping paper, to customer order. The seasonality of sales has made forecasting a particularly difficult problem, especially because of the exact timing of seasonal surges. Management would like to develop an exponential smoothing forecasting model that will aid in forecasting sales. The model should be one that minimizes forecast error.
    a. Based on the following aggregate product sales data that have been collected over the last five years, what forecasting model type and smoothing constant values would you suggest?

	SALES, ROLLS				
Month	2003	2002	2000	1999	1998
Jan.	7,000	8,000	7,000	10,000	10,000
Feb.	8,000	9,000	10,000	9,000	7,000
Mar.	8,000	8,000	10,000	9,000	8,000
Apr.	8,000	10,000	8,000	7,000	7,000

May	9,000	10,000	9,000	10,000	11,000
June	11,000	13,000	12,000	11,000	11,000
July	11,000	9,000	12,000	15,000	13,000
Aug.	11,000	13,000	15,000	19,000	15,000
Sept.	15,000	17,000	20,000	21,000	25,000
Oct.	17,000	17,000	20,000	21,000	25,000
Nov.	19,000	21,000	23,000	25,000	27,000
Dec.	13,000	15,000	13,000	17,000	13,000
Total	137,000	150,000	159,000	174,000	172,000

b. What is your forecast for January 2004?
c Construct a 95-percent confidence band on the forecast in part (b).

11. A steel distributor cuts sheets of steel from coils supplied by major mills. Accurate forecasting of coil usage can be very beneficial in controlling raw material inventories. Eighty percent of the sales price is in the cost of purchased materials. Although determining purchase quantities involves many considerations, a three-month moving average is used to project the usage rate to the next month. Actual coil usage rates in lb for two products are given in the following table:

	COIL A569 CQ P&O			COIL A366 CQ CR		
	Two Years Ago	Last Year	This Year	Two Years Ago	Last Year	This Year
Jan.	206,807	304,580	341,786	794,004	735,663	633,160
Feb.	131,075	293,434	521,878	703,091	590,202	542,897
Mar.	124,357	273,725	179,878	757,610	601,401	692,376
Apr.	149,454	210,626	226,130	499,022	529,784	703,151
May	169,799	150,587	177,400	445,703	672,040	917,967
June	216,843	289,621	182,109	483,058	450,735	532,171
July	288,965	168,590	123,957	446,770	567,928	654,445
Aug.	219,018	171,470	54,074	806,204	549,355	546,480
Sept.	65,885	209,351	136,795	646,300	481,355	472,664
Oct.	179,739	203,466		470,551	419,846	
Nov.	251,969	145,866		682,611	612,346	
Dec.	205,806	203,742		606,968	447,021	

a. Will an exponential smoothing model provide improved forecasts compared with the three-month moving average? If so, what type of model and smoothing constants would you suggest?
b. What is your forecast of the usage for October of this year?
c. If the actual October usage for A569 is 369,828 lb and for A366 is 677,649 lb, how do you rationalize the difference from your forecasts of these items?

CASE STUDY

World Oil

World Oil is a worldwide refiner and distributor of fuel products for automobiles, aircraft, trucks, and marine operations, service stations, and bulk facilities as outlets. Keeping more than 1,000 such outlets supplied is a significant operating problem for the company. Maintaining adequate fuel levels at the auto service stations is its major concern, because fuel generates the most revenue for the firm and has the greatest demand for customer service (product availability). Being able to forecast usage rates by product at these service stations is one of the key elements of good distribution operations. In particular, the tanker truck dispatchers need an accurate forecast of fuel usage in order to schedule fuel deliveries at service stations to avoid stockouts.

SERVICE STATION OPERATION

Service stations may carry three or four different grades of fuel including 87-, 89-, and 92-octane gasoline and diesel fuels. These are stored in underground tanks. Due to the variations in the usage rates among the stations and the limited capacities of these tanks, the frequency of replenishment may range from two or three times per day to only several times per week. Each tank is dedicated to one type of fuel. Fuel levels are measured periodically by placing a calibrated stick into a storage tank, although some of the more modern stations have electronic metering devices on their tanks. Tanker trucks, typically having four fuel compartments, are used for replenishment.

A FORECASTING SITUATION

Each service station's fuel grade represents a specific forecasting situation. A case in point is one of the lower-volume stations selling 87-octane fuel. With replenishment occurring only a few times per week, forecast of usage rates on a daily basis is adequate. Because usage does depend on the day of the week, forecasting for a particular day of the week may be quite different from any other day of the week. In Table 1, a history of Monday 87-octane fuel usage rates for the last 2-plus years is given for one low-volume station. A plot of this time series is shown in Figure 1. ■

QUESTIONS

1. Develop a forecasting procedure for this service station. Why did you select your method?
2. How should promotions, holidays, or other such periods where fuel usage rates deviate from normal patterns be handled in the forecast?
3. Forecast next Monday's usage and indicate the probable accuracy of the forecast.

Table 1 Historical Daily (Monday) Usage Rates of 87-Octane Fuel for a Low-Volume Auto Service Station

Two Years Ago		Last Year		This Year	
Week	Usage, Gal.	Week	Usage, Gal.	Week	Usage, Gal.
1 (Jan.)	530	1 (Jan.)	660	1 (Jan.)	790
2	570	2	640	2	860
3	560	3	810[b]	3	890
4	530	4	790[b]	4	780
5	510	5	820[b]	5	810
6	560	6	650	6	?
7	610	7	710		
8	560	8	700		
9	580	9	670		
10	610	10	690		
11	650	11	730		
12	700	12	730		
13	670	13	760		
14	700	14	790		
15	760	15	810		
16	730	16	870		
17	760	17	890		
18	820	18	870		
19	780	19	890		
20	900	20	880		
21	840	21	930		
22	770	22	980		
23	820	23	900		
24	800	24	860		
25	760	25	890		
26	760	26	880		
27	770	27	870		
28	790	28	840		
29	760	29	860		
30	740	30	910		
31	720	31	870		
32	670	32	860		
33	690	33	840		
34	470[a]	34	540[a]		
35	670	35	780		
36	690	36	750		
37	620	37	780		
38	650	38	760		
39	610	39	710		

Two Years Ago		Last Year		This Year	
Week	Usage, Gal.	Week	Usage, Gal.	Week	Usage, Gal.
40	620	40	730		
41	640	41	750		
42	590	42	750		
43	610	43	710		
44	600	44	750		
45	630	45	720		
46	600	46	770		
47	630	47	740		
48	640	48	750		
49	610	49	760		
50	590	50	780		
51	610	51	800		
52 (Dec.)	630	52 (Dec.)	850		
Totals	34,690		41,030		

[a]Holiday
[b]Promotional period

Figure 1 Fuel Usage on Mondays at a Low-Volume Service Station Over Approximately the Last Two Years

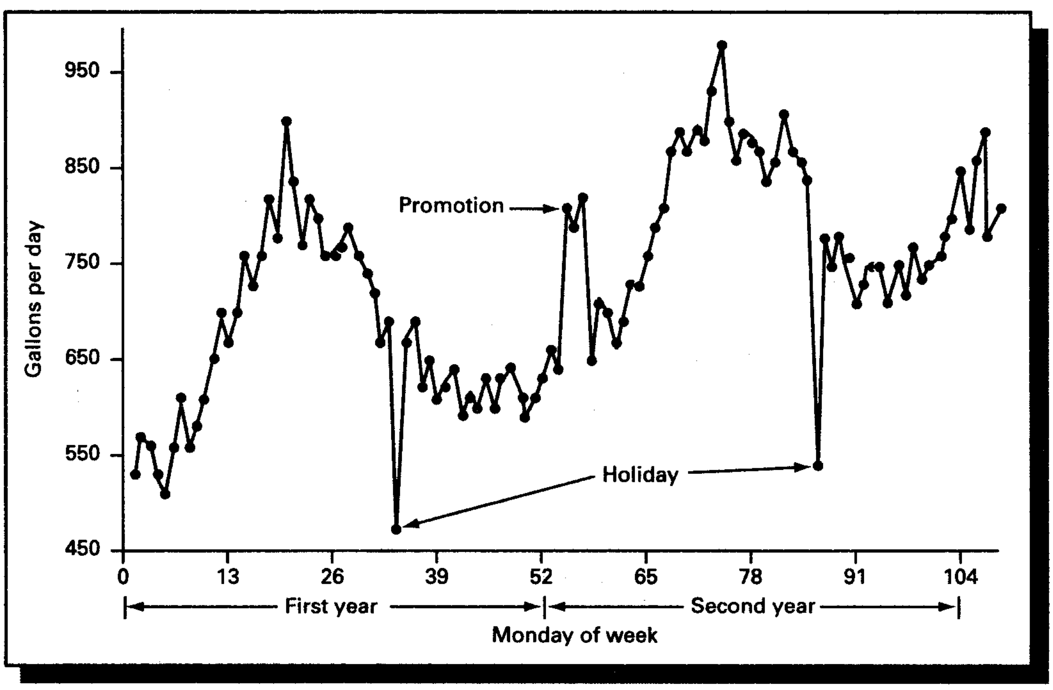

store shelves. Having these inventories on hand can cost between 20 and 40 percent of their value per year. Therefore, carefully managing inventory levels makes good economic sense. Even though many strides have been taken to reduce inventories through just-in-time, time compression, quick response, and collaborative practices applied throughout the supply channel, the annual investment in inventories by manufacturers, retailers, and merchant wholesalers, whose sales represent about 99 percent of GNP, is about 12 percent of the U.S. gross domestic product.[1] This chapter is directed toward managing the inventories that remain in the supply channel.

There is much to learn about inventory management and this chapter is rather lengthy because of it; however, the subject can be viewed in three major parts. First, inventories are most frequently managed as individual items located at single stocking points. This inventory control type has been researched extensively with methods for many specific applications. Second, inventory control will be viewed as management of inventory in the aggregate. Top managers are particularly interested in this perspective because of their need to control the overall inventory investment rather than individual stock-keeping units. Finally, managing inventories among multiple locations and multiple echelons within the supply channel will be examined.

APPRAISAL OF INVENTORIES

There are numerous reasons why inventories are present in a supply channel, yet in recent years inventory holding has been roundly criticized as unnecessary and wasteful. Consider why a firm might want inventories at some level in its operations, and why that firm would want to keep them to a minimum.

Arguments for Inventories

Reasons for holding inventories relate to customer service or to cost economies indirectly derived from them. Briefly consider some of these.

Improve Customer Service

Operating systems may not be designed to respond to customer requests for product or services in an instantaneous manner. Inventories provide a level of product or service availability, which, when located in the proximity of the customer, can meet high customer expectations for product availability. The availability of these inventories to customers can not only maintain sales, but they can also increase them.

[1]U.S. Bureau of the Census, *Statistical Abstract of the United States: 2001*, 121th ed. (Washington, DC: 2001), pp. 623, 644, and 657.

Application

Auto repair shops are faced with maintaining thousands of repair parts for a variety of automobiles from different model years. An automobile can contain 15,000 parts. To provide the fastest turnaround, repair shops carry a limited inventory of the more popular parts such as spark plugs, fan belts, and batteries. The auto manufacturer maintains a second inventory tier in regional warehouses from which parts can be transported using airfreight. The repair shops can, in some cases, receive these parts the same day they are requested. A high level of parts availability can be achieved with a minimum of on-site inventory.

Samsonite South Asia, headquartered in India, was facing long inventory cycles in Middle East and Asia. Production of soft luggage for Samsonite is predominantly in Southeast Asia, but the luggage used to be shipped to Europe and then distributed to the Middle East and Asia. The company took charge of a 20,000-square foot logistics center in Dubai to handle the logistics of soft luggage movement in South Asia and the Middle East. Earlier, this region was handled from Belgium in Europe. The new logistics center in Dubai now handles all this cutting short inventory cycles.[2]

Reduce Costs

Although holding inventories has an associated cost, their use can indirectly reduce operating costs in other supply channel activities that may more than offset the inventory carrying cost. First, holding inventories may encourage economies of production by allowing larger, longer, and more level production runs. Production output can be decoupled from the variation in demand requirements when inventories exist to act as buffers between the two.

Second, holding inventories fosters economies in purchasing and transportation. A purchasing department may buy in quantities beyond the firm's immediate needs in order to realize price-quantity discounts. The cost of holding the excess quantities until they are needed is balanced with the price reduction that can be achieved. In a similar manner, transportation costs can often be reduced through shipping in larger quantities that require less handling per unit. However, increasing the shipment size causes increased inventory levels that need to be maintained at *both* ends of the transportation channel. The reduction in transportation costs justifies the carrying of an inventory.

Third, forward buying involves purchasing additional product quantities at lower current prices rather than at higher anticipated future prices. Buying in quantities greater than immediate needs causes a larger inventory than does purchasing in quantities that more closely match immediate requirements. However, if prices are expected to rise in the future, some inventory resulting from forward buying can be justified.

Fourth, variability in the time that it takes to produce and transport goods throughout the supply channel can cause uncertainties that impact on operating costs as well as customer service levels. Inventories are frequently used at many points in the channel to buffer the effects of this variability and, thereby, help to smooth operations.

[2]Available at http://www.etintelligence.com.

Fifth, unplanned and unanticipated shocks can befall the logistics system. Labor strikes, natural disasters, surges in demand, and delays in supplies are the types of contingencies against which inventories can afford some protection. Having some inventory at key points throughout the supply channel allows the system to operate for a period while the effect of the shock can be diminished.

Application

Papermaking requires expensive Fourdrinier machines and other pieces of equipment that have large capacities. The high fixed cost of this equipment dictates that it constantly be kept busy. Demand for industrial paper products (e.g., kraft wrapping papers, multiwall bags, and bulk products) is anything but stable and known for sure. Although large orders can be scheduled directly to the process, production of small orders would be too costly, considering that changeovers can take 30 minutes on machines costing $3,500 per hour to operate. Producing to an inventory and servicing the small-order demand for the more standardized products from that inventory reduces setup time, which more than compensates for the inventory-carrying cost.

Arguments Against Inventories

It has been claimed that management's job is much easier having the security of inventories. Being overstocked is much more defensible from criticism than being short of supplies. The major portion of inventory-carrying costs is of an opportunity cost nature and, therefore, goes unidentified in normal accounting reports. To the extent that inventory levels have been too high for the reasonable support of operations, the criticism is perhaps deserved.

Critics have challenged the holding of inventories along several lines. First, inventories are considered wasteful. They absorb capital that might otherwise be put to better use, such as to improve productivity or competitiveness. In addition, they do not contribute any direct value to the firm's products, although they do store value.

Second, they can mask quality problems. When quality problems occur, reducing existing inventories to protect the capital investment is often a first consideration. Correcting quality problems can be slow.

Finally, using inventories promotes an insular attitude about the management of the supply channel as a whole. With inventories, it is often possible to isolate one stage of the channel from another. The opportunities arising from integrated decision making that considers the entire channel are not encouraged. Without inventories, it is difficult to avoid planning and coordinating across several echelons of the channel at one time.

Types of Inventories

Inventories can be categorized in five distinct forms. First, inventories may be in the *pipeline*. These are inventories in transit between echelons of the supply channel. Where movement is slow and/or over long distances, or movement must take place between many echelons, the amount of inventory in the pipeline may well exceed that held at the stocking points. Similarly, work-in-process inventories between manufacturing operations can be considered as inventories in the pipeline.

Second, some stocks may be held for *speculation,* but they are still part of the total inventory base that must be managed. Raw materials such as copper, gold, and silver are purchased as much for price speculation as they are to meet operating requirements. Where price speculation takes place for periods beyond the foreseeable needs of operations, such resulting inventories are probably more the concern of financial management than logistics management. However, when inventories are built up in anticipation of seasonal selling or occur due to forward buying activities, these inventories are likely to be the responsibility of logistics.

Third, stocks may be *regular* or *cyclical* in nature. These are the inventories necessary to meet the average demand during the time between successive replenishments. The amount of cycle stock is highly dependent on production lot sizes, economical shipment quantities, storage space limitations, replenishment lead times, price-quantity discount schedules, and inventory carrying costs.

Fourth, inventory may be created as a hedge against the variability in demand for the inventory and in replenishment lead time. This extra measure of inventory, or *safety* stock, is in addition to the regular stock that is needed to meet average demand and average lead-time conditions. Safety stock is determined from statistical procedures that deal with the random nature of the variability involved. The amount of safety stock maintained depends on the extent of the variability involved and the level of stock availability that is provided. Accurate forecasting is essential to minimizing safety stock levels. In fact, if lead time and demand could be predicted with 100 percent accuracy, no safety stock would be needed.

Finally, some of the inventory deteriorates, becomes out of date, or is lost or stolen when held for a time. Such inventory is referred to as *obsolete, dead,* or *shrinkage* stock. Where the products are of high value, perishable, or easily stolen, special precautions must be taken to minimize the amount of such stock.

Classifying Inventory Management Problems

Managing inventories involves a variety of problem types. Since managing inventories cannot be handled using a single solution method, we need to categorize the methods into several major groups. Inventory management using just-in-time methods will not be included in this grouping, since the technique is discussed in Chapter 10. With the remaining inventory management methods, we assume that the conditions of demand level and its variability, lead time and its variability, and inventory-related costs are known, and that we must do the best job of inventory control, given these conditions. In contrast, the just-in-time philosophy (supply directly to demand as it occurs) is to eliminate inventories by reducing the variability in demand and replenishment cycle time, reducing lot sizes, and forging strong relationships with a limited number of suppliers to ensure quality products and accurate order filling.

Nature of Demand

The nature of demand over time plays a significant role in determining how we treat the control of inventory levels. Several common types of demand patterns are shown in Figure 9-2. Perhaps the most common demand characteristic is for it to continue into the indefinite future. The demand pattern is referred to as *perpetual*. Although demand for most products rises and falls through their life cycles, many products have a selling life that is sufficiently long to be considered infinite for planning purposes. Even though brands turn over at the rate of 20 percent per year, a life cycle of three to five years can be long enough to justify treating them as having a perpetual demand pattern.

On the other hand, some products are highly seasonal or have a one-time, or *spike*, demand pattern. Inventories that are held to meet such a demand pattern usually cannot be sold off without deep price discounting. A single inventory replenishment order must be placed with little or no opportunity to reorder or return goods if demand has been inaccurately projected. Fashion clothing, Christmas trees, and political campaign buttons are examples of this type of demand pattern.

Similarly, demand may display a lumpy, or erratic, pattern. The demand may be perpetual, but there are periods of little or no demand followed by periods of high demand. The timing of lumpy demand is not as predictable as for seasonal demand, which usually occurs at the same time every year. Items in inventory are typically a mixture of lumpy and perpetual demand items. A reasonable test to separate these is to recognize that lumpy items have a high variance around their mean demand level. If the standard deviation of the demand distribution, or the forecast error, is greater than

Figure 9-2
Examples of Common Product Demand Patterns

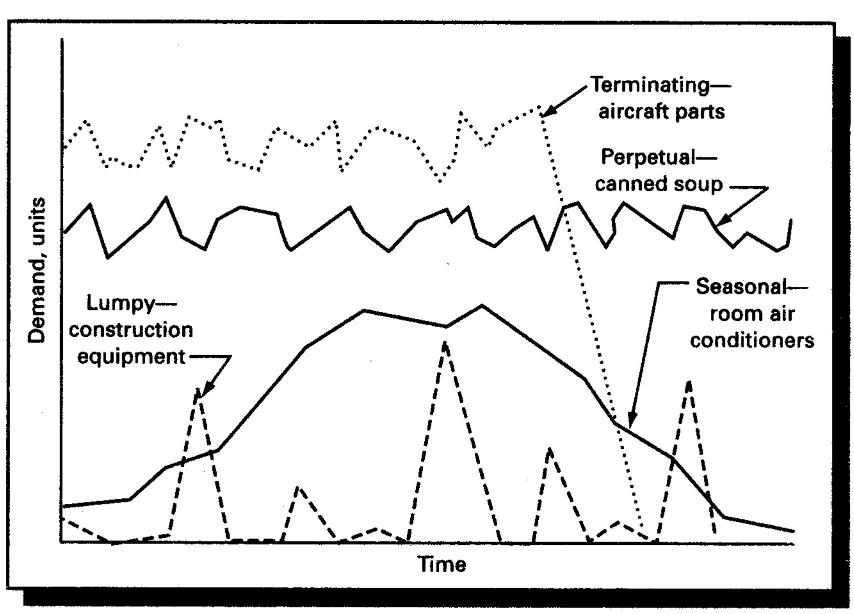

the average demand, or forecast, the item is probably lumpy. Inventory control of such items is best handled by intuitive procedures, or by a modification of the mathematical procedures discussed in this chapter, or through collaborative forecasting.

There are products whose demand terminates at some predictable time in the future, which is usually longer than one year. Inventory planning here involves maintaining inventories to just meet demand requirements, but some reordering within the limited time horizon is allowed. Textbooks with planned revisions, spare parts for military aircraft, and pharmaceuticals with a limited shelf life are examples of products with a defined life. Since the distinction between these products and those with a perpetual life is often blurred, they will not be treated differently from perpetual-life products for the purposes of developing a methodology to control them.

Finally, the demand pattern for an item may be derived from demand for some other item. The demand for packaging materials is derived from the demand for the primary product. The inventory control of such dependent demand items is best handled with some form of just-in-time planning such as MRP or DRP, which are discussed in Chapter 10.

Management Philosophy

Inventory management is developed around two basic philosophies. First, there is the *pull* approach. This philosophy views each stocking point, for example, a warehouse, as independent of all others in the channel. Forecasting demand and determining replenishment quantities are accomplished by taking into account only local conditions, as illustrated in Figure 9-3. No direct consideration is given to the effect

Figure 9-3
Pull Versus Push Inventory Management Philosophies

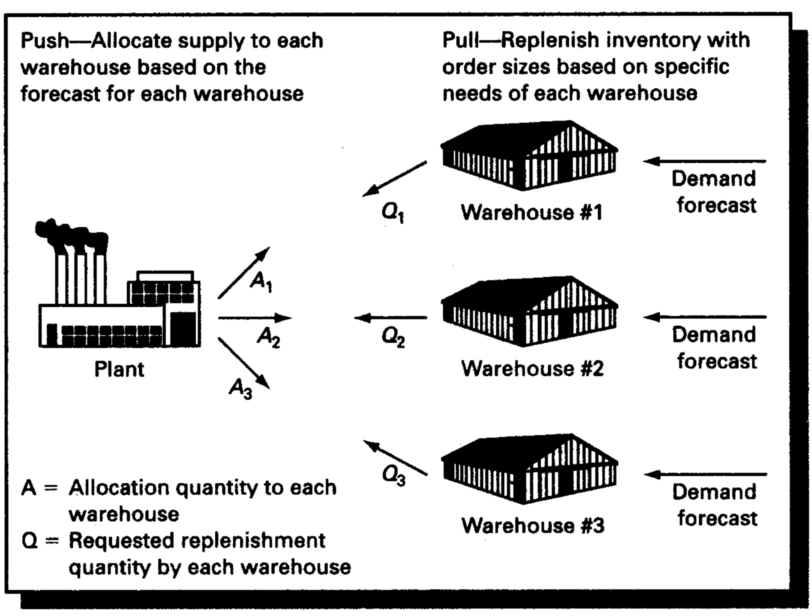

that the replenishment quantities, each with their different levels and timing, will have on the economics of the sourcing plant. However, this approach does give precise control over inventory levels at each location. Pull methods are particularly popular at the retail level in the supply channel where over 60 percent of the hard goods and almost 40 percent of the soft goods are under replenishment programs.[3]

Alternatively, there is the *push* approach to inventory management (see Figure 9-3). When decisions about each inventory are made independently, the timing and replenishment order sizes are not necessarily well coordinated with production lot sizes, economical purchase quantities, or order size minimums. Therefore, many firms choose to allocate replenishment quantities to inventories based on projected needs for inventories at each location, available space, or some other criteria. Inventory levels are set collectively across the entire warehousing system. Typically, the push method is used when purchasing or production economies of scale outweigh the benefits of minimum collective inventory levels as achieved by the pull method. In addition, inventories can be managed centrally for better overall control, production and purchase economies can be used to dictate inventory levels for lower costs, and forecasting can be made on aggregate demand and then apportioned to each stocking point for improved accuracy.

Collaborative replenishment can be used as a hybrid of the pull and push methods. In this case, the channel members representing the source point and the stocking point jointly agree on the replenishment quantities and their timing. The result can be order replenishment that is more economical for the supply channel than if either party alone were to make the replenishment decision.

Degree of Product Aggregation

Much of inventory control is directed at controlling each item in inventory. Precise control of each item can lead to precise control of the sum of all item inventory levels. This is a bottoms-up approach to inventory management.

Management of product groups rather than individual items is an alternate, or top-down, approach—a common perspective of top management. Although daily operation of inventories may require item-level control, strategic planning of inventory levels can be accomplished by substantially aggregating products into broad groups. This is a satisfactory approach when managing the inventory investment of all items collectively is the issue, and the effort associated with an item-by-item analysis for thousands of items at many locations is not warranted. Methods of control tend to be less precise for aggregate inventory management than for item management.

Multi-Echelon Inventories

As supply chain management has encouraged managers to think about including increasingly more of the supply channel in their planning processes, inventories that

[3]Tom Andel, "Manage Inventory, Own Information," *Transportation & Distribution* (May 1996), p. 54ff.

span more than one channel echelon become a focus. Rather than planning inventories at each location separately, planning their levels in concert can lead to lower overall inventory quantities. Multi-echelon inventory planning has been a particularly difficult problem to solve, but some progress is being made in methods useful to managers.

Virtual Inventories

Historically, customers have been served from inventories to which they were assigned. If product was out of stock, either a sale was lost or the product was placed on back order. Improved information systems changed that. It became possible for firms to know product inventory levels at every stocking point in the logistics network, creating a virtual inventory of products. Because of this, out-of-stock items could be replaced by cross filling them from other locations. Satisfying demand when cross filling is an option can result in lower overall inventory levels and higher product fill rates.

INVENTORY OBJECTIVES

Inventory management involves balancing product availability, or customer service, on the one hand with the costs of providing a given level of product availability on the other. Since there may be more than one way of meeting the customer service target, we seek to minimize inventory-related costs for each level of customer service (see Figure 9-4). Let us begin the development of the methodology to control inventories with a way to define product availability and an identification of the costs relevant to managing inventory levels.

Figure 9-4
Design Curves for Inventory Planning

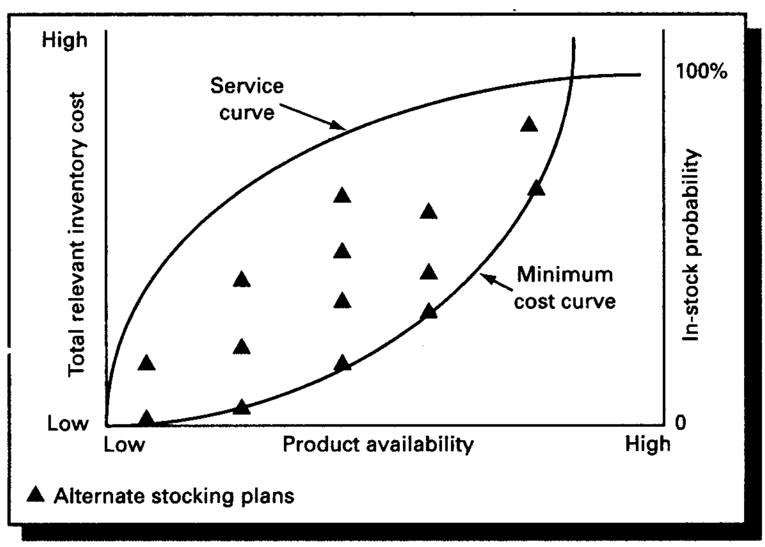

Product Availability

A primary objective of inventory management is to ensure that product is available at the time and in the quantities desired. This is commonly judged based on the probability of fulfillment capability from current stock. This probability, or item fill rate, is referred to as the service level, and, for a single item, can be defined as

$$\textit{Service level} = 1 - \frac{\textit{Expected number of units out of stock annually}}{\textit{Total annual demand}} \tag{9-1}$$

Service level is expressed as a value between 0 and 1. Since a target service level is typically specified, our task will be to control the expected number of stock out units.

We will see that controlling the service level for single items is computationally convenient. However, customers frequently request more than one item at a time. Therefore, the probability of filling the customer order completely can be of greater concern than single-item service levels. For example, suppose that five items are requested on an order where each item has a fill rate of 0.95, that is, only a 5 percent chance of not being in stock. Filling the entire order without any item being out of stock would be

$$0.95 \times 0.95 \times 0.95 \times 0.95 \times 0.95 = 0.77$$

The probability of filling the order completely is somewhat less than the individual item probabilities.

A number of orders from many customers will show that a mixture of items can appear on any one order. The service level is then more properly expressed as a *weighted average fill rate* (*WAFR*). The WAFR is found by multiplying the frequency with which each combination of items appears on the order by the probability of filling the order completely, given the number of items on the order. If a target WAFR is specified, then the fill rates for each item must be adjusted to achieve this desired WAFR.

Example

A specialty chemical company receives orders for one of its paint products. The paint product line contains three separate items that customers order in various combinations. From a sampling of orders over time, the items appear on orders in seven different combinations with frequencies as noted in Table 9-1. Also from the company's historical records, the probability of having each item in stock is $SL_A = 0.95$; $SL_B = 0.90$; and $SL_C = 0.80$. As the calculations in Table 9-1 show, the WAFR is 0.801. There will be about one order in five where the company cannot supply all items at the time of the customer request.

Recall that additional measures for customer service were discussed in Chapter 4. Some of these measures encompass more than inventory and are not appropriate for

Item Combination on Order	(1) Frequency of Order	(2) Probability of Filling Order Complete	(3) = (1) × (2) Marginal Value
A	0.1	(.95) = 0.950	0.095
B	0.1	(.90) = 0.900	0.090
C	0.2	(.80) = 0.800	0.160
A, B	0.2	(.95)(.90) = 0.855	0.171
A, C	0.1	(.95)(.80) = 0.760	0.076
B, C	0.1	(.90)(.80) = 0.720	0.072
A, B, C	0.2	(.95)(.90)(.80) = 0.684	0.137
	1.0	WAFR =	0.801

Table 9-1 Computation of the Weighted Average Fill Rate

the discussion here. However, additional inventory performance measures might include percent of items on back order, percent of orders filled complete, percent of orders filled complete to a given percentage, and percent of items cross filled from secondary locations. These are not discussed further.

Relevant Costs

Three general classes of costs are important to determining inventory policy: procurement costs, carrying costs, and stockout costs. These costs are in conflict, or in trade-off, with each other. For determining the order quantity to replenish an item in inventory, these relevant costs trade-off are shown in Figure 9-5.

Figure 9-5
Trade-Off of the Relevant Inventory Costs with the Order Quantity

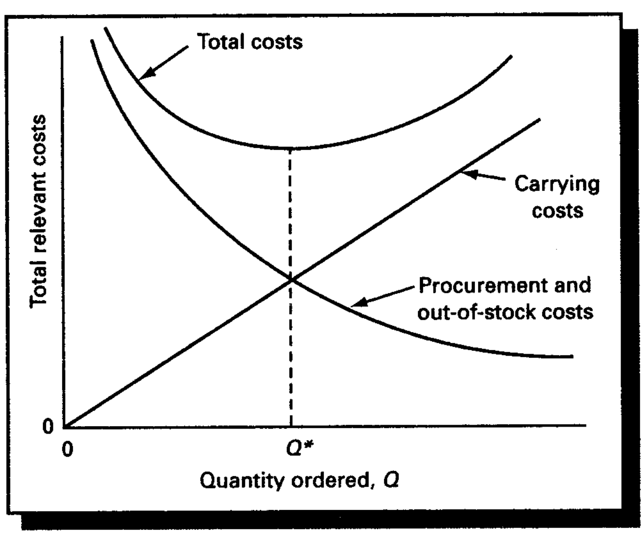

Procurement Costs

Costs associated with the acquisition of goods for the inventory replenishment are often a significant economic force that determines the reorder quantities. When a stock replenishment order is placed, a number of costs are incurred that are related to the processing, setup, transmitting, handling, and purchase of the order. More specifically, procurement costs may include the price, or manufacturing cost, of the product for various order sizes; the cost for setting up the production process; the cost of processing an order through the accounting and purchasing departments; the cost of transmitting the order to the supply point, usually using mail or electronic means; the cost of transporting the order when transportation charges are not included in the price of the purchased goods; and the cost of any materials handling or processing of the goods at the receiving point. When the firm is self-supplied, as in the case of a factory replenishing its own finished goods inventories, procurement costs are altered to reflect production setup costs. Transportation costs may not be relevant if a delivered pricing policy is in effect.

Some of these procurement costs are fixed per order and do not vary with the order size. Others, such as transportation, manufacturing, and materials-handling costs, vary to a degree with order size. Each requires slightly different analytical treatment.

Carrying Costs

Inventory carrying costs result from storing, or holding, goods for a period and are roughly proportional to the average quantity of goods on hand. These costs can be collected into four classes: space costs, capital costs, inventory service costs, and inventory risk costs.

Space Costs. Space costs are charges made for the use of the volume inside the storage building. When the space is rented, storage rates are typically charged by weight for a period of time, for example, $/cwt./month. If the space is privately owned or contracted, space costs are determined by allocating space-related operating costs, such as heat and light, as well as fixed costs, such as building and storage equipment costs, on a volume-stored basis. Space costs are irrelevant when calculating carrying costs for in-transit inventories.

Capital Costs. Capital costs refer to the cost of the money tied up in inventory. This cost may represent over 80 percent of total inventory cost (see Table 9-2), yet it is the most intangible and subjective of all the carrying cost elements. There are two reasons for this. First, inventory represents a mixture of short-term and long-term assets, as some stocks may serve seasonal needs and others are held to meet longer-term demand patterns. Second, the cost of capital may vary from the prime rate of interest to the opportunity cost of capital.

The exact cost of capital for inventory purposes has been debated for some time. Many firms use their average cost of capital, whereas others use the average rate of return required of company investments. The hurdle rate has been suggested as most accurately reflecting the true capital cost.[4] The hurdle rate is the rate of return on the most lucrative investments that the firm does not accept.

---

[4]Douglas M. Lambert and Bernard J. LaLonde, "Inventory Carrying Costs," *Management Accounting* (August 1976), pp. 31–35.

Table 9-2
Relative Percentages of Cost Elements in Inventory Carrying Costs

Interest and opportunity costs	82.00%
Obsolescence and physical depreciation	14.00
Storage and handling	3.25
Property taxes	0.50
Insurance	0.25
Total	100.00%

Source: Adapted from Robert Landeros and David M. Lyth, "Economic-Lot-Size Models for Cooperative Inter-Organizational Relationships," *Journal of Business Logistics*, Vol. 10, No. 2 (1989), p. 149.

Inventory Service Costs. Insurance and taxes are also a part of inventory carrying costs because their level roughly depends on the amount of inventory on hand. Insurance coverage is carried as a protection against losses from fire, storm, or theft. Inventory taxes are levied on the inventory levels found on the day of assessment. Although the inventory at the point in time of the tax assessment only crudely reflects the average inventory level experienced throughout the year, taxes typically represent only a small portion of total carrying cost. Tax rates are readily available from accounting or public records.

Inventory Risk Costs. Costs associated with deterioration, shrinkage (theft), damage, or obsolescence make up the final category of carrying costs. In the course of maintaining inventories, a certain portion of the stock will become contaminated, damaged, spoiled, pilfered, or otherwise unfit or unavailable for sale. The costs associated with such stock may be estimated as the direct loss of product value, as the cost of reworking the product, or as the cost of supplying it from a secondary location.

Out-of-Stock Costs

Out-of-stock costs are incurred when an order is placed but cannot be filled from the inventory to which the order is normally assigned. There are two kinds of out-of-stock costs: lost sales costs and back order costs. Each presupposes certain actions on the part of the customer, and, because of their intangible nature, they are difficult to measure accurately.

A *lost sales cost* occurs when the customer, faced with an out-of-stock situation, chooses to withdraw his or her request for the product. The cost is the profit that would have been made on this particular sale and may include an additional cost for the negative effect that the stockout may have on future sales. Products for which the customer is very willing to substitute competing brands, such as bread, gasoline, or soft drinks, are those that are most likely to incur lost sales.

A *back order cost* occurs when a customer will wait for his or her order to be filled so that the sale is not lost, only delayed. Back orders can create additional clerical and sales costs for order processing, and additional transportation and handling costs when such orders are not filled through the normal distribution channel. These costs are tangible, so measurement of them is not too difficult. There also may be the intangible cost of lost future sales. This cost is very difficult to measure. Products (automobiles and major appliances) that can be differentiated in the consumer's mind are more likely to be back ordered than substituted.

Push Inventory Control

Let's begin to develop methods for controlling inventory levels with the *push* philosophy. Recall that this method is appropriate where production or purchase quantities exceed the short-term requirements of the inventories into which the quantities are to be shipped. If these quantities cannot be stored at the production site for lack of space or other reasons, then they must be allocated to the stocking points, hopefully in some way that makes good economic sense. Push is also a reasonable approach to inventory control where production or purchasing is the dominant force in determining the replenishment quantities in the channel. In either case, the following questions need to be addressed. How much inventory should be maintained at each stocking point? How much of a purchase order or production run should be allocated to each stocking point? How should the excess supply over requirements be apportioned among the stocking points?

A method for pushing quantities into stocking points involves the following steps:

1. Determine through forecasting or other means the requirements for the period between now and the next expected production run or vendor purchase.
2. Find the current on-hand quantities at each stocking point.
3. Establish the stock availability level at each stocking point.
4. Calculate total requirements from the forecast plus additional quantities needed to cover uncertainty in the demand forecast.
5. Determine net requirements as the difference between total requirements and the quantities on hand.
6. Apportion the excess over total net requirements to the stocking points on the basis of the average demand rate, that is, the forecasted demand.
7. Sum the net requirements and prorate the excess quantities to find the amount to be allocated to each stocking point.

Example

When the tuna boats are sent to the fishing grounds, a packer of tuna products must process all the tuna caught since storage is limited and, for competitive reasons, the company does not want to sell the excess of this valued product to other packers. Therefore, this packer processes all fish brought in by the fleet and then allocates the production to its three field warehouses on a monthly basis. There is only enough storage at the plant for one month's demand. The current production run is 125,000 lb.

For the upcoming month, the needs of each warehouse were forecasted, the current stock levels checked, and desired stock availability level noted for each warehouse. The findings are tabulated in Table 9-3.

Warehouse	Current Stock Level	Forecasted Demand	Forecast Error[a] (Std. Dev.)	Stock Availability Level[b]
1	5,000 lb	10,000 lb	2,000 lb	90%
2	15,000	50,000	1,500	95%
3	30,000	70,000	20,000	90%
		130,000		

[a]Assumed to be normally distributed.
[b]Stock availability level is defined as the probability of stock being available during the forecast period.

Table 9-3 Basic Inventory Planning Data for a Tuna Packer

Now we need to compute the total requirements for each warehouse. Total requirements for warehouse 1 will be the forecast quantity and the added amount needed to ensure a 90 percent stock availability level. This is found from

$$\text{Total requirements} = \text{Forecast} + (z \times \text{Forecast error})$$

where z is the number of standard deviations on the normal distribution curve beyond the forecast (the distribution mean) to the point where 90 percent of the area under the curve is represented (see Figure 9-6). From the normal distribution curve in Appendix A, $z = 1.28$. Hence, the total requirement for warehouse 1 is $10{,}000 + (1.28 \times 2{,}000) = 12{,}560$. Other warehouse total requirements are computed similarly. The information is recorded in Table 9-4.

Net requirements are found as the difference between total requirements and the quantity on hand in the warehouse. Summing the net requirements (110,635) shows that $125{,}000 - 110{,}635 = 14{,}365$, which is the excess production that needs to be prorated to the warehouses.

Figure 9-6
Area Under the Forecast Distribution for Warehouse 1

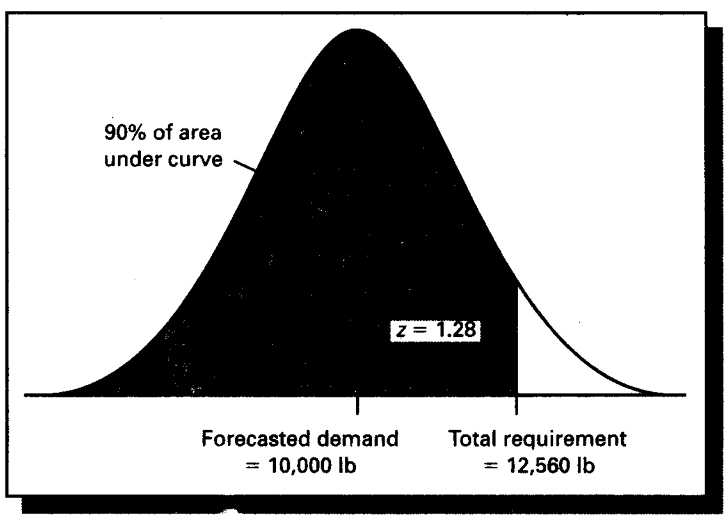

Warehouse	(1) Total Requirements	(2) On Hand	(3) = (1) − (2) Net Requirements	(4) Prorated Excess	(5) = (3) + (4) Allocation
1	12,560 lb	5,000	7,560 lb	1,105 lb	8,665 lb
2	52,475	15,000	37,475	5,525	43,000
3	95,600	30,000	65,600	7,735	73,335
	160,635		110,635	14,365	125,000

Table 9-4 Allocation of Tuna Production to Three Warehouses

Prorating the excess production of 14,365 lb is made in proportion to the average demand rate for each warehouse. Average demand for warehouse 1 is 10,000 lb against a total demand rate for all warehouses of 130,000 lb. The proportion of the excess allocated to warehouse 1 should be (10,000 ÷ 130,000)(14,365) = 1,105. Prorate the excess for the remaining warehouses in a similar manner. The total allocation to a warehouse is the sum of its net requirement plus its portion of the production excess. The results are tabulated in Table 9-4.

Basic Pull Inventory Control

Recall that pull inventory control gives low inventory levels at stocking points because of its response to the demand and cost conditions particular to each stocking point. Although many specific methods have been developed to handle a variety of situations, the discussion here will attempt to highlight the fundamental ideas. Specifically, a contrast will be made between (1) demand that is one-time, highly seasonal, or perpetual; (2) ordering that is triggered from a particular inventory level or from a process of inventory level review; and (3) the degree of uncertainty in demand and replenishment lead time.

Single-Order Quantity

Many practical inventory problems exist where the products involved are perishable or the demand for them is a one-time event. Products such as fresh fruits and vegetables, cut flowers, newspapers, and some pharmaceuticals have a short and defined shelf life, and they are not available for subsequent selling periods. Others, such as toys and fashion clothes for the immediate selling season, hotdog buns for a baseball game, and posters for a political campaign, have a one-time demand level that usually cannot be estimated with certainty. Only one order can be placed for these products to meet such demand. We wish to determine how large the single order should be.

To find the most economic order size (Q^*), we can appeal to marginal economic analysis. That is, Q^* is found at the point where the marginal profit on the next unit sold equals the marginal loss of not selling the next unit. The marginal profit per unit obtained by selling a unit is

$$\text{Profit} = \text{Price per unit} - \text{Cost per unit} \tag{9-2}$$

The per-unit loss incurred by not selling a unit is

$$\text{Loss} = \text{Cost per unit} - \text{Salvage value per unit} \tag{9-3}$$

Considering the probability of a given number of units being sold, the expected profits and losses are balanced at this point. That is,

$$CP_n\,(\text{Loss}) = (1 - CP_n)(\text{Profit}) \tag{9-4}$$

where CP_n represents the cumulative frequency of selling at least n units of the product. Solving the above expression for CP_n, we have

$$CP_n = \frac{\text{Profit}}{\text{Profit} + \text{Loss}} \tag{9-5}$$

This says that we should continue to increase the order quantity until the cumulative probability of selling additional units just equals the ratio of Profit ÷ (Profit + Loss).

Example

A grocery store estimates that it will sell 100 pounds of its specially prepared potato salad in the next week. The demand distribution is normally distributed with a standard deviation of 20 pounds. The supermarket can sell the salad for $5.99 per pound. It pays $2.50 per pound for the ingredients. Since no preservatives are used, any unsold salad is given to charity at no cost.

Finding the quantity to prepare that will maximize profit requires that we first compute CP_n. That is,

$$CP_n = \frac{\text{Profit}}{\text{Profit} + \text{Loss}} = \frac{(5.99 - 2.50)}{(5.99 - 2.50) + 2.50} = 0.583$$

From the normal distribution curve (Appendix A), the optimum Q^* is at the point of 58.3 percent of the area under the curve (see Figure 9-7). This is a point where $z = 0.21$. The salad preparation quantity should be

$$Q^* = 100 \text{ lb} + 0.21(20 \text{ lb}) = 104.2 \text{ lb}$$

When demand is discrete, the order quantity may be between whole values. In such cases, we will round up Q to the next higher unit to ensure that at least CP_n is met.

Example

An equipment repair firm wishes to order enough spare parts to keep a machine tool running throughout a trade show. The repairperson prices the parts at $95 each if

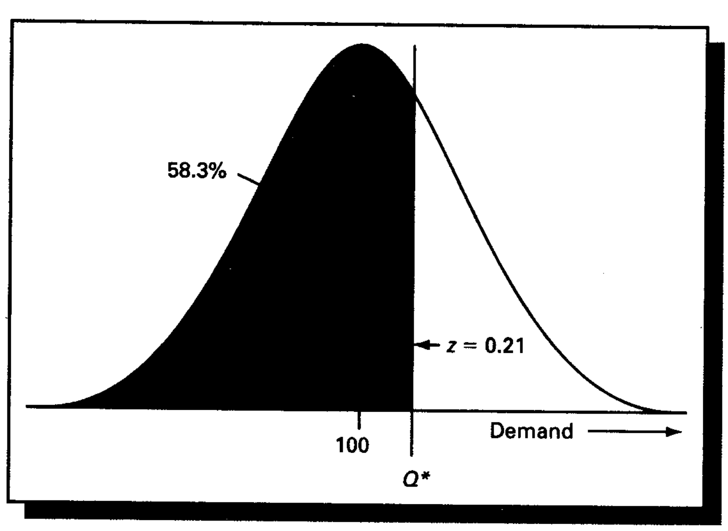

Figure 9-7
Normally Distributed Demand for Potato Salad Problem

needed for a repair. He pays $70 for each part. If all the parts are not needed, they may be returned to the supplier for a credit of $50 each. The demand for the parts is estimated according to the following distribution:

Number of Parts	Frequency of Need	Cumulative Frequency
0	0.10	0.10
1	0.15	0.25
2	0.20	0.45
3	0.30	0.75 ⇐ Q^*
4	0.20	0.95
5	0.05	1.00
	1.00	

We should set the order quantity at

$$CP_n = \frac{\text{Profit}}{\text{Profit} + \text{Loss}} = \frac{(95 - 70)}{(95 - 70) + (70 - 50)} = 0.555$$

The CP_n value is between 2 and 3 units on the cumulative frequency column. Rounding up, we choose $Q^* = 3$.

Repetitive Order Quantities

In contrast to demand that occurs only periodically or possibly only once, demand may be perpetual. Inventory replenishment orders repeat over time and may be supplied instantaneously in their entirety, or the items in the orders may be supplied over time. Both cases are illustrated.

Instantaneous Resupply

When demand is continuous and the rate is essentially constant, controlling inventory levels is accomplished by specifying (1) the quantity that will be used to replenish the inventory on a periodic basis and (2) the inventory replenishment frequency. This is a problem of balancing conflicting cost patterns. In the simplest case, it requires balancing procurement costs against carrying costs, as was shown in Figure 9-5. Ford Harris recognized this problem as early as 1913 in his work at Westinghouse. The model that he developed for finding the optimum order quantity has become known as the basic *e*conomic *o*rder *q*uantity (*EOQ*) formula,[5] and it serves as the basis for many of the pull inventory policies currently used in practice.

The basic *EOQ* formula is developed from a total cost equation involving procurement cost and inventory carrying cost. It is expressed as

$$\text{Total cost} = \text{Procurement cost} + \text{Carrying cost}$$

$$TC = \frac{D}{Q}S + \frac{ICQ}{2} \qquad \textbf{(9-6)}$$

where

TC = total annual relevant inventory cost, dollars
Q = order size to replenish inventory, units
D = Item annual demand occurring at a certain and constant rate over time, units/year
S = Procurement cost, dollars/order
C = Item value carried in inventory, dollar/unit
I = carrying cost as a percent of item value, percent/year

The term D/Q represents the number of times per year a replenishment order is placed on its supply source. The term $Q/2$ is the average amount of inventory on hand.

As Q varies in size, one cost goes up as the other goes down. It can be shown mathematically that an optimal order quantity (Q^*) exists where the two costs are in balance and the minimal total cost results. The formula for this *EOQ* is

$$Q^* = \sqrt{\frac{2DS}{IC}} \qquad \textbf{(9-7)}$$

The optimal time between orders is therefore

$$T^* = \frac{Q^*}{D} \qquad \textbf{(9-8)}$$

and the optimal number of times per year to place an order is

$$N = \frac{D}{Q^*} \qquad \textbf{(9-9)}$$

[5]F. W. Harris, "How Many Parts to Make at Once," *Factory, The Magazine of Management*, Vol. 10, No. 2 (February 1913), pp. 135–136, 152.

Example

An industrial machine tools manufacturer supplies replacement parts from its inventory. For a particular part, the annual demand is expected to be 750 units. Machine setup costs are $50, carrying costs are 25 percent per year, and the part is valued in inventory at $35 each. The economic order quantity placed on production is

$$Q^* = \sqrt{\frac{2DS}{IC}} = \sqrt{\frac{2(750)(50)}{(0.25)(35)}} = 92.58 \text{ or } 93 \text{ units}$$

This order size is expected to be placed in production every $T^* = Q^*/D = 92.58/750 =$ 0.12 years, or 0.12 (years) $\times$ 52 (weeks per year) = 6.4 weeks. For practical reasons, we may wish to round this to 6 or 7 weeks with some slight increase in total costs.

A Lead Time for Resupply

Using this formula as part of a basic inventory control procedure, we see that a sawtooth pattern of inventory depletion and replenishment occurs, as illustrated in Figure 9-8. We can now introduce the idea of a *reorder point,* which is the quantity to which inventory is allowed to drop before a replacement order is placed. Since there is generally a time lapse between when the order is placed and when the items are

Figure 9-8 A Basic Pull Inventory Control Model for a Replenishment Part

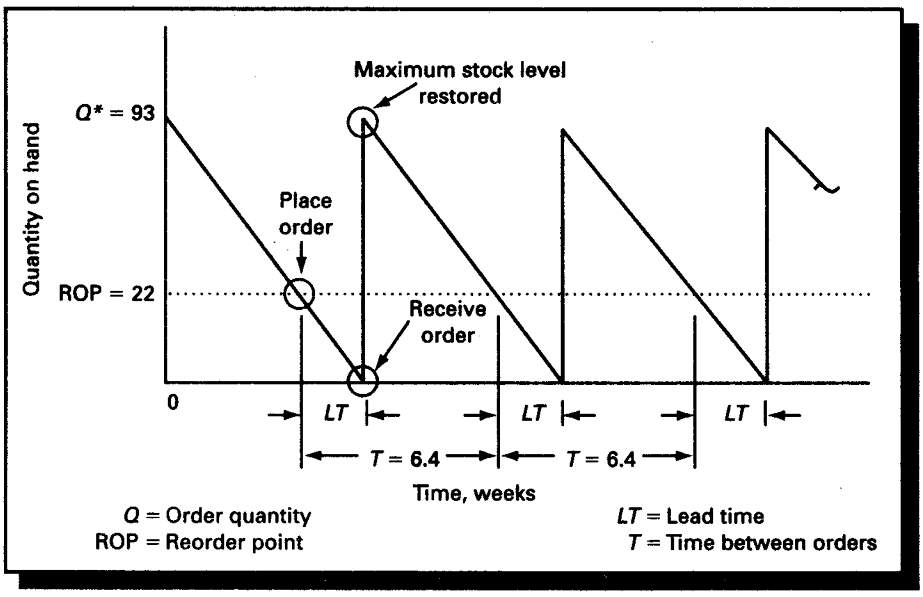

available in inventory, the demand that occurs over this lead time must be anticipated. The reorder point (*ROP*) is

$$ROP = d \times LT \tag{9-10}$$

where

ROP = reorder point quantity, units
d = demand rate, in time units
LT = average lead time, in time units

The demand rate (*d*) and the average lead time (*LT*) must be expressed in the same time dimension.

Example

Continuing the previous machine replacement part example, suppose that it takes 1.5 weeks to set up production and make the parts. The demand rate is d = 750 (units per year)/52 (weeks per year) = 14.42 units per week. Therefore, $ROP = 14.42 \times 1.5 = 21.6$, or 22 units. We can now state the inventory policy: When the inventory level drops to 22 units, place a replenishment order for 93 units.

Sensitivity to Data Inaccuracies

Demand and costs cannot always be known for sure. However, our computation of the economic order quantity is not very sensitive to incorrect data estimations. For example, if demand is in fact 10 percent higher than anticipated, Q^* should only be increased by $\sqrt{1.10}$ = 4.88 percent. If the carrying cost is 20 percent lower than assumed, Q^* should be increased by only $\sqrt{1/(1-0.20)}$ = 11.8 percent. These percentage changes are inserted into the *EOQ* formula without changing the remaining cost and/or demand factors since they remain constant. Notice the stability of the Q^* values. If the incorrect order quantity were used in these two cases, total costs would have been in error by only 0.11 percent and 0.62 percent, respectively.

Noninstantaneous Resupply

A built-in assumption of Ford Harris's original *EOQ* formula was that resupply would be made instantaneously in a single batch of size Q^*. In some manufacturing and resupply processes, output is continuous for a time, and it may take place simultaneously with demand. The basic sawtooth pattern of on-hand inventory is modified, as shown in Figure 9-9. The order quantity now becomes the production run, or production lot size, quantity (*POQ*) labeled Q_p^*. To find Q_p^*, the basic order quantity formula is modified as follows:

$$Q_p^* = \sqrt{\frac{2DS}{IC}}\sqrt{\frac{p}{p-d}} \tag{9-11}$$

where p is the output rate. Computing Q_p^* only makes sense when the output rate p exceeds the demand rate d.

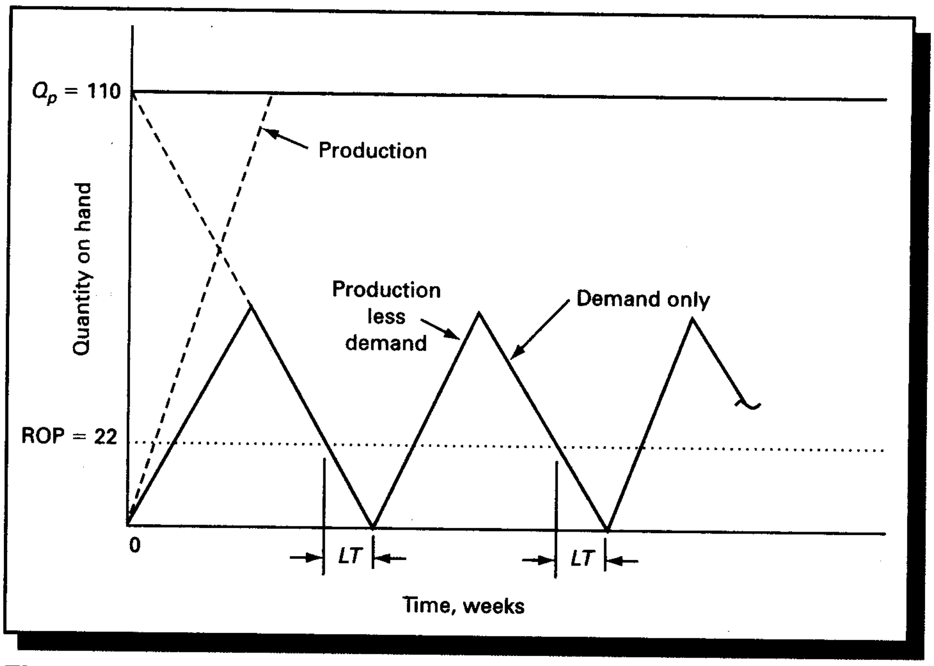

Figure 9-9 Noninstantaneous Resupply for a Parts Replacement Problem

Example

Again, for the previous parts replacement problem, suppose that the production rate for these parts is 50 units per week. The production run quantity is

$$Q_p^* = \sqrt{\frac{2(70)(50)}{(0.25)(35)}}\sqrt{\frac{50}{50-14.42}}$$
$$= 92.5 \times 1.185 = 109.74, \text{ or } 110 \text{ units}$$

The *ROP* quantity remains unchanged.

Advanced Pull Inventory Control

Advanced pull control of inventories means that we recognize that demand and lead time cannot be known for sure. Therefore, we must plan for the situation where not enough stock may be on hand to fill customer requests. In addition to the regular stock that is maintained for meeting average demand and average lead time, an incremental quantity is added to inventory. The amount of this safety, or buffer, stock sets the level of stock availability provided to customers by controlling the probability of a stockout occurring.

Two inventory control methods form the foundation for most pull-type management philosophies with perpetual demand patterns. These are (1) the reorder point method and (2) the period review method. Practical control systems may be based on either of these methods or on a combination of them.

A Reorder Point Model with Uncertain Demand

Finding Q and ROP*

Reorder point inventory control assumes that demand is perpetual and continually acts on inventory to reduce its level. When inventory is depleted to the point where its level is equal to or less than a specified quantity called the reorder point, an economic order quantity of Q^* is placed on the supplying source to replenish the inventory. The effective inventory level at a particular point in time is the quantity on hand plus the quantity on order, less any commitments against the inventory, such as customer back orders or allocations to production or customers. The entire quantity Q^* arrives at a point in time offset by the lead time. Between the time when the replenishment order is placed at the reorder point and when it arrives in stock, there is a risk that demand will exceed the remaining amount of inventory. The probability of this occurring is controlled through raising or lowering the reorder point and by adjusting Q^*.

In Figure 9-10, the operation of the reorder point system is illustrated for a single item where the demand during the lead time is known only to the extent of a normal probability distribution. This demand during lead time (DDLT) distribution has a mean of X' and a standard deviation of s_d'. The values for X' and s_d' are usually not

Figure 9-10 Reorder Point Inventory Control Under Uncertainty for Item

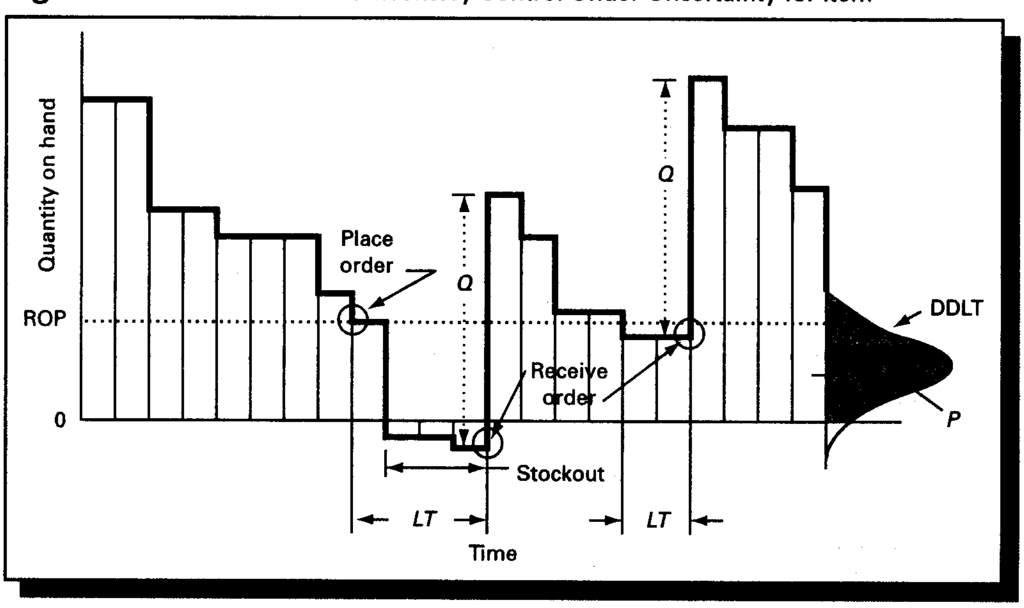

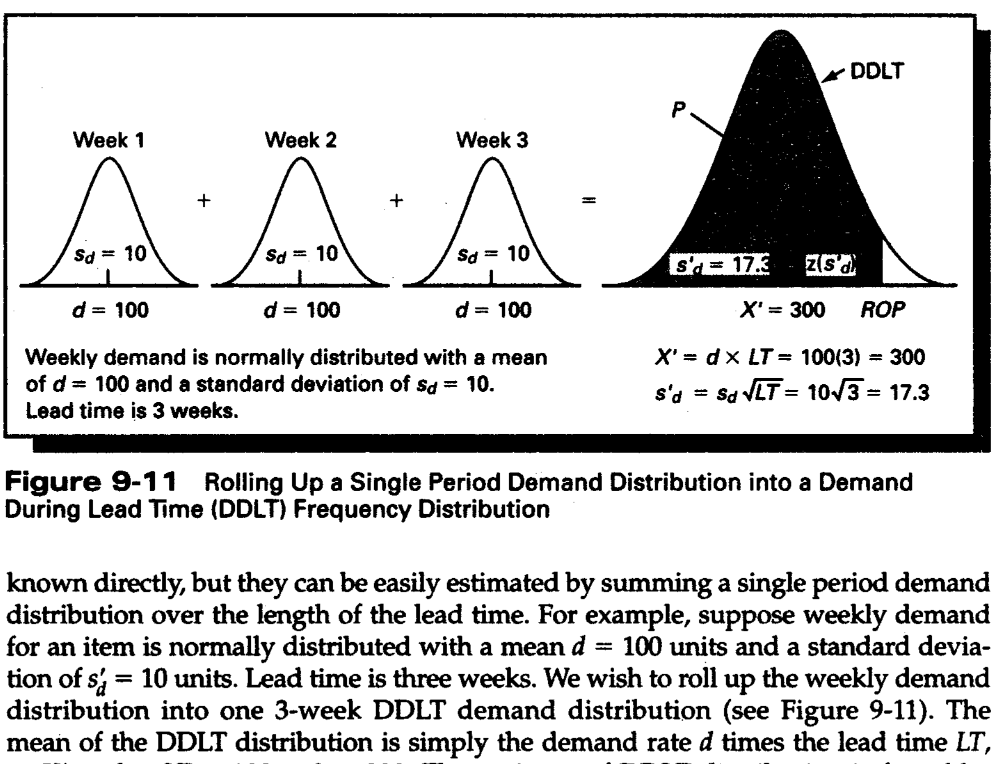

Figure 9-11 Rolling Up a Single Period Demand Distribution into a Demand During Lead Time (DDLT) Frequency Distribution

known directly, but they can be easily estimated by summing a single period demand distribution over the length of the lead time. For example, suppose weekly demand for an item is normally distributed with a mean $d = 100$ units and a standard deviation of $s_d' = 10$ units. Lead time is three weeks. We wish to roll up the weekly demand distribution into one 3-week DDLT demand distribution (see Figure 9-11). The mean of the DDLT distribution is simply the demand rate d times the lead time LT, or $X' = d \times LT = 100 \times 3 = 300$. The variance of DDLT distribution is found by adding the variances of the weekly demand distributions (see Figure 9-11). That is, $s_d'^2 = LT(s_d^2)$. The standard deviation is the square root of $s_d'^2$, which is $s_d'^2 = s_d\sqrt{LT} = 10\sqrt{3} = 17.3$.

Finding Q^* and the ROP is rather mathematically complex; however, a satisfactory approximation can be found if we first determine Q^* according to the basic EOQ formula (Equation 9-7).[6] Then, find

$$ROP = d \times LT + z(s_d') \quad \textbf{(9-12)}$$

The term z is the number of standard deviations from the mean of the DDLT distribution to give us the desired probability of being in stock during the lead time period (P). The value for z is found in a normal distribution table (Appendix A) for the area under the curve P.

Example

Buyers Products Company distributes an item known as a tie bar, which is a U-bolt used on truck equipment. The following data have been collected for this item held in inventory:

[6]Sven Axsäter, "Using the Deterministic EOQ Formula in Stochastic Inventory Control," *Management Science*, Vol. 42, No. 6 (June 1996), p. 830.

Monthly demand forecast, d	11,107 units
Std. error of forecast, s_d	3,099 units
Replenishment lead time, LT	1.5 months
Item value, C	\$0.11/unit
Cost for processing vendor order, S	\$10/order
Carrying cost, I	20%/year
In-stock probability during lead time, P	75%

The reorder quantity is

$$Q^* = \sqrt{\frac{2DS}{IC}} = \sqrt{\frac{2(11,107)(10)}{(0.20/12)(0.11)}} = 11,008 \text{ units}$$

The reorder point is

$$ROP = d \times LT + z(s'_d)$$

where $s'_d = s_d\sqrt{LT} = 3,099\sqrt{1.5} = 3,795$ units. The value for z is 0.67 from Appendix A, where the fraction of the area under the normal distribution curve is 0.75. Thus,

$$ROP = (11,107 \times 1.5) + (0.67 \times 3,795) = 19,203 \text{ units}$$

So, when the effective inventory level drops to 19,203 units, place a replenishment order for 11,008 units.

It is common for the reorder point quantity to exceed the order quantity, as was the case in the previous example. This frequently happens when lead times are long or demand rates are high. For the reorder point control system to work properly, make sure that the timing of a replenishment order is based on the effective inventory level. Recall that the effective inventory level requires that all stock on order be added to the current quantity on hand when making a comparison to the reorder point. When $ROP > Q^*$, the result of this procedure is that a second order will be placed before the first arrives in stock.

Average Inventory Level

The average inventory level for this item is the total of the regular stock plus safety stock. That is,

$$\begin{aligned} \text{Average inventory} &= \text{Regular stock} + \text{Safety stock} \\ AIL &= Q/2 + z(s'_d) \end{aligned} \tag{9-13}$$

Example

For the previous tie bar problem, the average inventory would be $AIL = (11,008/2) + (0.67 \times 3,795) = 8,047$ units

---

Total Relevant Cost

The total relevant cost is useful for comparing alternative inventory policies or determining the impact of deviations from optimum policies. We add two new terms to the total cost formula stated in Equation (9-6), which account for uncertainty. These are safety stock and out-of-stock terms. Total cost can now be expressed as

Total cost = Order cost + Carrying cost, regular stock
+ Carrying cost, safety stock + Stockout cost

$$TC = \frac{D}{Q}S + IC\frac{Q}{2} + ICzs'_d + \frac{D}{Q}ks'_dE_{(z)} \qquad \textbf{(9-14)}$$

where k is the stockout cost per unit. The stockout cost term requires some explanation. First, the combined term of $s'_d E_{(z)}$ represents the expected number of units out of stock during an order cycle. $E_{(z)}$ is called the unit normal loss integral whose values are tabled as a function of the normal deviate z (see Appendix B) Second, the term D/Q is the number of order cycles per period of time, usually a year. Hence, the number of order cycles times the expected number of units out of stock during each order cycle gives the total expected number of units out of stock for the entire period. Then, multiplying by the out-of-stock cost yields the total period cost.

Example

Continuing the tie bar example, suppose the stockout cost is estimated at \$0.01 per unit. The total annual cost for the item would be

$$TC = \frac{11,107(12)10)}{11,008} + 0.20(0.11)\left(\frac{11,008}{2}\right)$$
$$+0.20(0.11)(0.67)(3,795) + \frac{11,107(12)}{11,008}(0.01)(3,795)(0.150)$$
$$= 121.08 + 121.09 + 55.94 + 68.92 = \$367.03 \text{ per year}$$

Note: The value of 0.150 for $E_{(z)} = E_{(0.67)}$ is from the body of the table in Appendix B for $z = 0.67$.

Service Level

The customer service level, or item fill rate, achieved by a particular inventory policy was previously defined in Equation (9-1). Restating it in the symbols now being used, we have

$$SL = 1 - \frac{(D/Q)(s'_d \times E_{(z)})}{D} = 1 - \frac{s'_d(E_{(z)})}{Q} \qquad \textbf{(9-15)}$$

Example

The service level achieved for the tie bar problem is

$$SL = 1 - \frac{3,795(0.150)}{11,008} = 0.948$$

That is, the demand for tie bars can be met 94.8 percent of the time. Note that this is somewhat higher than the probability of a stockout during the lead time of $P = 0.75$.

Application

A manufacturer of quick-connect hose couplings uses an easy method for implementing a reorder point method of inventory control. A finished goods inventory is maintained at the factory from which customer orders are filled. The stock is divided into two sections. An amount of an item equal to the reorder point quantity is placed in a covered tray in the reserve section of the stocking area. A second tray contains the remainder of the stock. All orders are filled from the second tray first. When all stock is depleted from the second tray, the first tray is brought from reserve storage and inserted into its position. This action is the trigger to place a replenishment order on production. Little or no paperwork is needed to make a rather sophisticated inventory control system operate effectively.

The Reorder Point Method with Known Stockout Costs

When the stockout costs are known, it is not necessary to assign a customer service level. The optimum balance between service and cost can be calculated. An iterative computational procedure is outlined as follows:

1. Approximate the order quantity from the basic *EOQ* formula [Equation (9-7)] that is,

$$Q = \sqrt{\frac{2DS}{IC}}$$

2. Compute the probability of being in stock during the lead-time if back ordering is allowed

$$P = 1 - \frac{QIC}{Dk} \qquad \textbf{(9-16)}$$

or if during a stockout the sales are lost

$$P = 1 - \frac{QIC}{Dk + QIC} \qquad \textbf{(9-17)}$$

Find s'_d. Find the z value that corresponds to P in the normal distribution table (Appendix A). Find $E_{(z)}$ from the unit normal loss integral table (Appendix B).

3. Determine a revised Q from a modified EOQ formula, which is

$$Q = \sqrt{\frac{2D[S + ks'_d E_{(z)}]}{IC}} \tag{9-18}$$

4. Repeat steps 2 and 3 until there is no change in P or Q. Continue.
5. Compute ROP and other statistics as desired.

Example

Repeating the tie bar example, with the known stockout cost of $0.01 per unit and back orders are allowed.

Estimate Q

$$Q = \sqrt{\frac{2DS}{IC}} = \sqrt{\frac{2(11,107)(12)(10)}{0.20(0.11)}} = 11,008 \text{ units}$$

Estimate P

$$P = 1 - \frac{QIC}{Dk} = 1 - \frac{11,008(0.20)(0.11)}{11,107(12)(0.01)} = 0.82$$

From Appendix A, $z_{@0.82} = 0.92$. From Appendix B, $E_{(0.92)} = 0.0968$.

Revise Q The standard deviation of DDLT was calculated previously to be $s'_d = 3{,}795$ units. Now,

$$Q = \sqrt{\frac{2D[S + ks'_d E_{(z)}]}{IC}} = \sqrt{\frac{2(11,107)(12)[(10 + 0.01(3,795)(0.068)]}{0.20(0.11)}} = 12,872 \text{ units}$$

Revise P

$$P = 1 - \frac{12,872(0.20)/(0.11)}{11,107(12)(0.01)} = 0.79$$

Now, $z_{@0.79} = 0.81$ and $E_{(0.81)} = 0.1181$

Revise Q

$$Q = \sqrt{\frac{2(11,107)(12)[10 + 0.01(3,795)(0.1181)]}{0.20(0.11)}} = 13,246 \text{ units}$$

We continue this revision process until the changes in P and Q are so small that further calculation is impractical. The results are $P = 0.78$, $Q^* = 13{,}395$ units, and $ROP = 19{,}583$ units, with a total relevant cost of $TC = \$15{,}019$ and an actual service level (item fill rate) of $SL = 96$ percent.

The Reorder Point Method with Demand and Lead Time Uncertainty

Accounting for uncertainty in the lead time can extend the realism of the reorder point model. What we wish to do is find the standard deviation (s_d') of the DDLT distribution based on uncertainty in both demand and lead time. Adding the demand variance to the lead time variance gives a revised formula for s_d', which is

$$s_d' = \sqrt{LTs_d^2 + d^2 s_{LT}^2} \tag{9-19}$$

where s_{LT} is the lead time standard deviation.[7]

Example

In the tie bar problem, s_{LT} is 0.5 months. The value for s_d' would now be

$$s_d' = \sqrt{1.5(3{,}099)^2 + 11{,}107^2(0.5)^2} = 6{,}727 \text{ units}$$

Combining demand and lead time variability in this way can greatly increase s_d, and the resulting safety stock. Brown warns that demand and lead time distributions may be dependent on each other.[8] Rather, when a replenishment order is placed, a fair idea is known as to the lead time for that order. Therefore, application of Equation (9-19) may lead to an overstatement of s_d' and the resulting amount of safety stock. If lead times do vary unpredictably, Brown suggests the following precise procedure for determining the standard deviation of demand during lead time:

> Forecast demand per lead time. A lead time starts when you trigger a replenishment order. Record the demand year-to-date at that time. Later, whenever material is received is, by definition, the end of the lead-time. Examine the demand year-to-date. The difference between the current demand year-to-date and the value when the order was released is precisely, by definition, the demand during the lead time. The values of this variable can be forecasted (usually with very simple forecast models) and the mean square error is the variance of demand during lead time, precisely the value being sought.[9]

Alternately and less precisely, the longest lead time may be used as the average lead time with s_{LT} set at zero (0). The standard deviation is then computed as $s_d' = s_d\sqrt{LT}$.

[7] Note that if demand is known for sure ($s_d = 0$) and lead time is uncertain, then $s_d' = ds_{LT}$.
[8] Robert G. Brown, *Materials Management Systems* (New York: John Wiley & Sons, 1977), pp. 150–151.
[9] Ibid.

Figure 9-12 Multiple Time Elements Throughout a Supply Channel

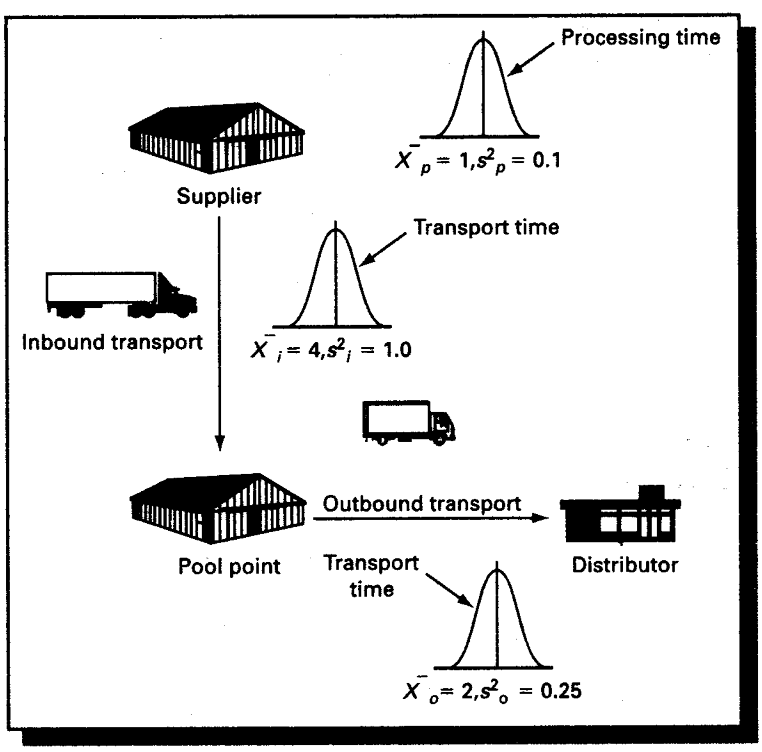

Example

Suppose inventory is to be maintained on a distributor's shelf for an item whose demand is forecasted to be $d = 100$ units per day and $s_d = 10$ units per day. A reorder point is the method of inventory control. There are multiple points throughout the supply channel where time is incurred in the product flow between source point and customer. The distributions of these times that form the order replenishment lead time are shown in Figure 9-12. No significant amount of inventory is maintained at the pool point or in the trucks.

We also know that

$$I = 10\%/\text{year}$$
$$S = \$10/\text{order}$$
$$C = \$5/\text{unit}$$
$$P = 0.99$$

Determine the average inventory to be held at the distributor.

Solution The reorder point inventory control method applies. However, determining the statistics of the demand-during-lead-time distribution requires taking the lead time for the *entire* channel into account.

Recall:

$$s_d' = \sqrt{LTs_d^{\,2} + d^2 s_{LT}^{\,2}}$$

where from Figure 9-12

$$s_{LT}^2 = s_p^2 + s_i^2 + s_o^2 = 0.1 + 1.0 + 0.25 = 1.35 \text{ days}$$

and

$$LT = \overline{X}_p + \overline{X}_i + \overline{X}_o = 1 + 4 + 2 = 7 \text{ days}$$

Now,

$$s_d' = \sqrt{7 \times 10^2 + 100^2 \times 1.35} = \sqrt{14,200} = 119.16 \text{ units}$$

and

$$AIL = \frac{Q^*}{2} + zs_d'$$

where

$$Q^* = \sqrt{\frac{2(100)(10)}{0.1(5)}} = 63 \text{ units}$$

Finally, the average inventory level is

$$AIL = \frac{63}{2} + 2.33(119.16) = 309 \text{ units}$$

A Periodic Review Model with Uncertain Demand

An alternative to the reorder point method of control is the periodic review method. Although the reorder point method offers precise control over each item in inventory and, therefore, the lowest total relevant cost, it has some economic disadvantages. For example, each item is possibly ordered at a different time, thus missing joint production, transportation, or buying economies. Administratively, reorder point control requires constant monitoring of the inventory levels. Alternatively, under periodic review control, inventory levels for multiple items can be reviewed at the same time so that they may be ordered together, thus realizing production, transportation, or purchasing economies. Periodic review control results in slightly more inventory, but the added carrying costs may be more than offset by reduced administrative costs, lower prices, or lower procurement costs. Reasons for preferring a periodic review method can be summarized as follows:

1. A manual bookkeeping inventory system is used, and it is convenient to review inventory stocks on a definite schedule. This might be done on a *cycle count* basis, in which a portion of stock is reviewed each day or week, perhaps on an *ABC* basis (reordering *A* items more often than *B* items, etc). This also allows balancing of clerical workload.

2. A large number of items are to be jointly ordered from the same vendor sources.
3. Items ordered have a significant effect on the supplying plant's production output, and order predictability is desirable.
4. Significant transportation savings can sometimes result when several items are ordered at the same time.[10]

Single Item Control

The periodic review model is very similar to the reorder point model under uncertain demand conditions. However, one important difference in the periodic review model is that demand fluctuations during the order interval and the lead time must be protected against, whereas only demand fluctuations during the lead time are important in calculating safety stock using the reorder point method. This makes the periodic review model more complex to formulate precisely than the reorder point model, but an approximate solution will provide reasonable answers. Approximate solutions in inventory control are reasonable since the total cost curve usually has a flat bottom such that slight deviations from optimum values for the policy variables result in only small changes to the total cost.

Periodic review control operates as shown in Figure 9-13. That is, the inventory level for an item is audited at predetermined intervals (T). The quantity to be placed on order is the difference between a maximum quantity (M) and the amount on hand at the review time. Thus, inventory is controlled through the setting of T^* and M^*.

A good approximation for the optimum review interval begins with the basic inventory control model. That is,

$$Q^* = \sqrt{\frac{2DS}{IC}}$$

and the review interval is

$$T^* = \frac{\text{Order quantity}}{\text{Annual demand}} = \frac{Q^*}{D}$$

The order interval may also be assigned a particular value that best conforms to the practices of the firm. Of course, this does not necessarily assure an optimum policy.

Next, construct the distribution for demand over the order interval plus the lead time [$DD(T^* + LT)$], as shown in Figure 9-14. The point where the probability of a stockout during the protection period ($1 - P$) equals the area under the normal distribution curve is the point of the maximum level (M^*). This point may be calculated as

$$M^* = d(T^* + LT) + z(s_d') \tag{9-20}$$

where $d(T^* + LT)$ is the mean of the DD($T^* + LT$) distribution, d is the average daily demand rate, and s_d' is the standard deviation of the $DD(T^* + LT)$ distribution. This standard deviation is now calculated as

[10]Lynn E. Gill, George Isoma, and Joel L. Sutherland, "Inventory and Physical Distribution Management," in James F. Robeson and Robert G. House (eds.), *The Distribution Handbook* (New York: The Free Press, 1985), p. 673.

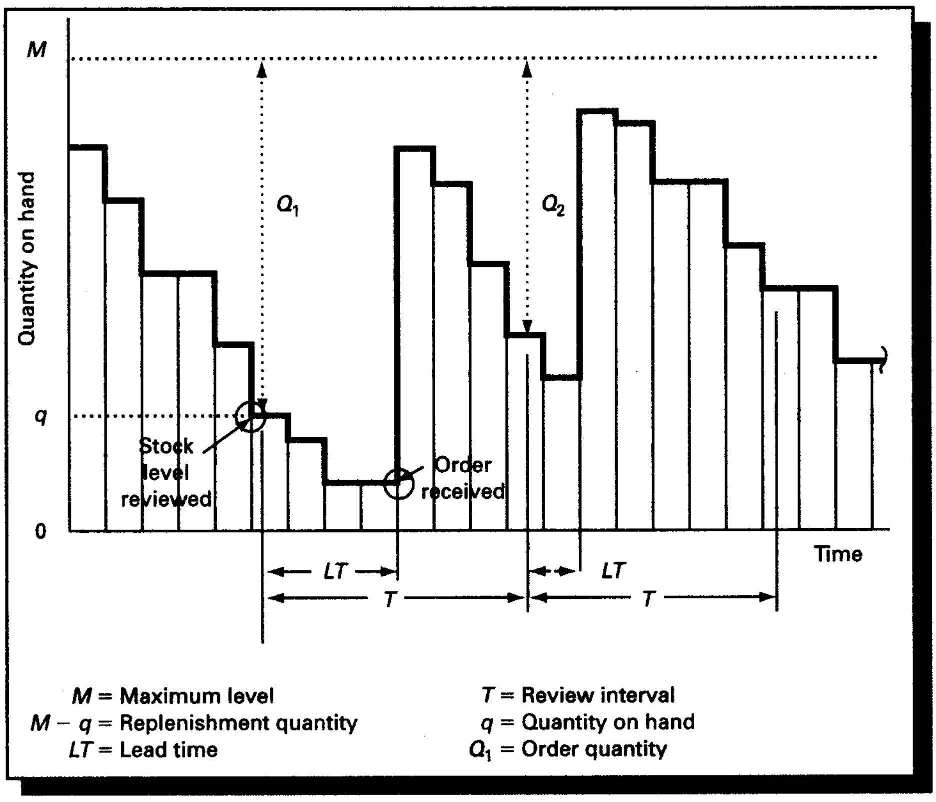

Figure 9-13 Periodic Review Inventory Control with Uncertainty for an Item

Figure 9-14
A Distribution of Demand Over the Order Interval Plus Lead Time for the Periodic Review Inventory Control Method

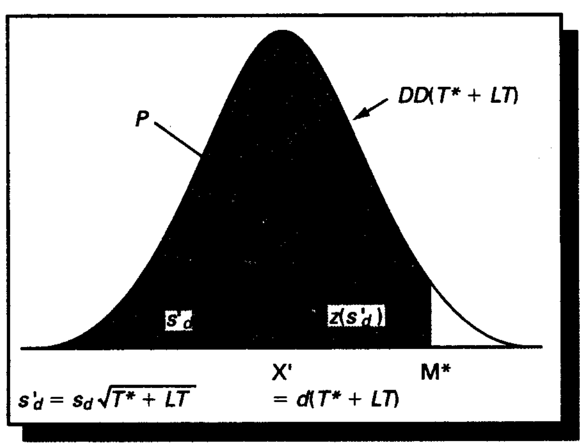

$$s'_d = s_d\sqrt{T^* + LT} \quad (9\text{-}21)$$

where lead time is known for sure.

The average inventory level is found from

$$AIL = \frac{dT^*}{2} + z(s'_d) \quad (9\text{-}22)$$

and the total relevant cost is computed with the same formula as under the reorder point method, that is, Equation (9-14).

Example

Let's use the tie bar problem data, but now develop a periodic review policy for it.

Find T^* and M^* The optimal order quantity is the same as under the reorder point policy, or 11,008 units. The order interval is

$$T^* = \frac{Q^*}{d} = \frac{11,008}{11,107} = 0.991, \text{ or 1 month}$$

Then, the demand standard deviation during the review period plus lead time is

$$s'_d = s_d\sqrt{T^* + LT} = 3,099\sqrt{0.991 + 1.5} = 4,891 \text{ units}$$

The maximum level for a $P = 0.75$ is

$$\begin{aligned} M^* &= d(T^* + LT) + z(s'_d) \\ &= 11,107(0.991 + 1.5) + 0.67(4,891) \\ &= 30,945 \text{ units} \end{aligned}$$

The inventory policy is to review the inventory level every month and to place a replenishment order for the difference between the quantity on hand and 30,945 units.

Average Inventory Level This inventory policy can be expected to produce an average inventory level of

$$AIL = \frac{dT^*}{2} + z(s'_d) = \frac{11,107(0.99)}{2} + 0.67(4,891) = 8,780 \text{ units}$$

Total Cost The total relevant cost according to Equation (9-14) is

$$\begin{aligned} TC &= 121.08 + 121.09 + 0.20(0.11)(0.67)(4,891) + \frac{11,107(12)}{11,008}(0.01)(4,891)(0.150) \\ &= 121.08 + 121.09 + 72.09 + 88.83 = \$403.09 \end{aligned}$$

Note the slightly higher annual cost ($367.03 vs. $403.09) for the periodic review policy compared with the reorder point policy.

Service Level The service level (item fill rate) achieved according to Equation (9-15) is

$$SL = 1 - \frac{4{,}891(0.150)}{11{,}008} = 0.933$$

Note: When using this method of determining service level (fill rate) in periodic inventory systems, researchers warn that accurate estimates are achieved when the fill rate is above 90 percent and demand variability is low.[11]

Joint Ordering

Both the reorder point and periodic review models discussed so far have been for single items. This assumes that each item in inventory is controlled independently of the others. In many cases, this is not the best practice since multiple items may be purchased from the same supplier or produced at the same time and location. Ordering multiple items at the same time and on the same order can result in economic benefits such as qualifying for price-quantity discounts or meeting vendor, carrier, or production minimum quantities, so inventory policy should reflect joint ordering. An inventory joint ordering policy involves determining a common inventory review time for all jointly ordered items, and then finding each item's maximum level (M^*) as dictated from its particular costs and service level.

The common review time for jointly ordered items is

$$T^* = \sqrt{\frac{2(O + \sum_i S_i)}{I \sum_i C_i D_i}} \qquad \textbf{(9-23)}$$

where O is the common cost for procuring an order and the subscript i refers to a particular item. The maximum level for each item is

$$M_i^* = d_i(T^* + LT) + z_i(s_d')_i \qquad \textbf{(9-24)}$$

The total relevant cost is

Total cost = Order cost + Regular stock carrying cost + Safety stock carrying cost + Stockout stock

$$TC = \frac{O + \sum_i S_i}{T} + \frac{TI \sum_i C_i D_i}{2} + I \sum_i C_i z_i (s_d')_i + \frac{1}{T} \sum_i k_i (s_d')_i (E_{(z)})_i \qquad \textbf{(9-25)}$$

An example with only two jointly ordered items will be used. Using more items increases the computations needlessly.

[11]M. Eric Johnson, Hau L. Lee, Tom Davis, and Robert Hall, "Expressions for Item Fill Rates in Periodic Inventory Systems," *Naval Research Logistics*, Vol. 42 (1995), pp. 57–80.

Example

Two items are to be jointly ordered from the same vendor. The following data are available:

	Item	
	A	B
Demand forecast, units/day	25	50
Error of the forecast, units/day	7	11
Lead time, days	14	14
Inventory carrying cost, %/year	30	30
Procurement cost, dollars/order/item	10	10
with a common cost of, dollars/order	30	
In-stock probability during order cycle plus lead time	70%	75%
Product value, dollars/unit	150	75
Stockout cost, dollars/unit	10	15
Selling days per year	365	365

Review Time The common review time for these items according to Equation (9-23) is

$$T^* = \sqrt{\frac{2[30 + (10 + 10)]}{[0.30/365][150(25) + 75(50)]}} = 4.03, \text{ or 4 days}$$

Note that we have taken care to make demand and carrying cost conform to the same period.

Maximum Level From Equation (9-24), the maximal order quantity for item A can be found. First,

$$(s_d')_A = (s_d)_A\sqrt{T^* + LT} = 7\sqrt{4 + 14} = 29.70 \text{ units}$$

Then for $z_{P=0.70} = 0.52$ (see Appendix A), M_A^* is

$$M_A^* = 25(4 + 14) + 0.52(29.70) = 465 \text{ units}$$

The maximum level for item B can be found similarly. First,

$$(s_d')_B = 11\sqrt{4 + 14} = 46.67 \text{ units}$$

Then for $z_{P=0.75} = 0.67$, M_B^* is

$$M_B^* = 50(4 + 14) + 0.67(46.67) = 931 \text{ units}$$

Average Inventory Level The average inventory level for item A according to Equation (9-22) is

$$AIL_A = 25\frac{4}{2} + 0.52(29.70) = 65 \text{ units}$$

And for item B, it is

$$AIL_B = 50\frac{4}{2} + 0.67(46.67) = 131 \text{ units}$$

Total Relevant Cost Using Equation (9-25), the total annual cost for items A and B is

$$\begin{aligned} TC &= \frac{30 + 2(10)}{4/365} + \frac{[4/365][0.30][150(25) + 75(50)][365]}{2} \\ &\quad + 0.30[150(0.52)(29.70) + 75(0.67)(46.67)] \\ &\quad + \frac{1}{4/365}[10(29.70)(0.1917) + 15(46.67)(0.1503)] \\ &= 4{,}563 + 4{,}500 + 1{,}399 + 14{,}796 \\ &= \$25{,}258 \text{ per year} \end{aligned}$$

Service Level The service level actually achieved for item A according to Equation (9-15) is

$$SL_A = 1 - \frac{29.70(0.1917)}{Q^*}$$

Applying a little algebra to Equation (9-8), $Q^* = T^* d = 4.03(25) = 101$. Thus,

$$SL_A = 1 - \frac{29.70(0.1917)}{101} = 0.944$$

For item B,

$$SL_B = 1 - \frac{46.67(0.1503)}{4.03(50)} = 0.9665$$

Practical Pull Inventory Control Methods

The models discussed so far in this chapter serve as a theoretical basis for the inventory control methods found in practice. Several realistic examples can be given.

A Min-Max System

The min-max system of inventory control is probably the most popular of all pull inventory control procedures. Historically, it has been implemented using manual

Figure 9-15 Min-Max Inventory Control Using a Kardex Record Card for an Office Supply Distributors' Standard Paper Item

Date	In/ Customer	Sales	On hand	Date	In/ Customer	Sales	On hand	Date	In/ Customer	Sales	On hand
10/26	Bal Fwd		80,500	2/2	Copies	50,000	35,000	3/30	Sup Meats	25,000	20,000
10/26	100M		180,500	2/5	Bel-Gar	5,000	30,000	3/30	Copies	50	19,950
10/30	Progression	20,000	160,500	2/6	Bel-Gar	15,000	15,000	3/30	Ptrs Dvl	5,000	14,950
10/30	Ogleby	25,000	135,500	2/6	Superior	25,000	0*	3/30	Belmont	10,000	4,950
11/2	Mid Ross	15,000	120,500	2/6	Unt Sply	15,000	0*	4/2	Berea Prtg	4,950	0
11/9	Unt Sply	50,000	70,500	2/6	Berea Prtg	15,000	0*	4/2	Berea Prtg	15,050	0*
11/29	Berea-Lit	25,000	45,500	2/8	Sagamore	5,000	0*	4/9	REM	500	0*
12/1	Dol Fed	10,000	35,500	2/14	100M		100,000	4/12	Mid Ross	5,000	0*
12/13	Card Fed	20,000	15,500	2/15	50M		150,000	5/7	Ohio Ost	5,000	0*
12/14	Belmont	15,000	500	2/16	Bel-Gar	5,000	145,000	5/8	Inkspots	5,000	0*
12/15	Shkr Sav	5,000	500*	2/21	Bel-Gar	15,000	130,000	5/8	Prts Dvl	2,500	0*
1/8	BFK	500	0	2/26	Inkspot	5,000	125,000	5/11	100M		100,000
1/8	100M		100,000	2/27	Lcl 25UAW	50,000	75,000	5/14	BVR	5,000	95,000
1/8	Card Fed	30,000	70,000	2/28	Ptrs Dvl	2,500	72,500	5/15	Guswold	10,000	85,000
1/9	Pt of View	10,000	60,000	2/28	Shkr Sav	25,000	47,500	5/16	ESB	15,000	70,000
1/17	Am Safety	5,000	55,000	3/1	Copies	35,000	12,500	5/16	Superior	50,000	20,000
1/23	Foster	15,000	40,000	3/2	Untd Tor	10,000	2,500	5/16	J Stephen	5,000	15,000
1/24	Gib Prtg	5,000	35,000	3/8	Sagamore	2,500	0	5/16	Am Aster	15,000	0*
1/26	Bel-Gar	5,000	30,000	3/8	Sagamore	12,500	0*	5/16	Am Aster	10,000	0*
1/26	Copies	20,000	10,000	3/12	150M		150,000	5/22	Sagamore	15,000	0*
1/29	Slvr Lake	5,000	5,000	3/12	Untd Tor	40,000	110,000	**Coding** 21200			
1/29	100 M		105,000	3/12	Preston	50,000	60,000	**M. Base Cost**	**Date**	**Min**	125M
2/2	Sagamore	20,000	85,000	3/20	Midland	15,000	45,000	2.64	4/2	**Max**	250M

Size	M/Wgt	Basis	Grain	Color	Finish	Grade	Location	Ctn. Skid Cont.	Att.
8½ x 14	12.72	20	L	White	RmSeal	Advantage Bond	F 14	5M	

*No stock or insufficient stock to meet demand control procedures and record keeping by ledger card (Kardex system), but it is also found in many computerized inventory control procedures. An example of such record keeping and control is shown in Figure 9-15.

The min-max inventory control procedure is a variant of the reorder point model; however, there are two differences. From Figure 9-16, we see that when an order is placed, it is for the amount determined by the difference between the target quantity, *M* (max level), and the quantity on hand, *q*, when the inventory level reaches the reorder point. Do not confuse this min-max control with the periodic review method. The max level *M* is simply the reorder point quantity (*ROP*) *plus* the economic order quantity (Q^*) found by the reorder point model. The reorder quantity is not always the same because the amount that the quantity on hand drops below the reorder point is added to Q^*. This extra amount is needed since the inventory level frequently drops in an amount greater than one unit, due to multiple units of the item being demanded from inventory between record updates. Q^* and *ROP* are approximated from the reorder point system as previously described. Although an exact computational procedure is available for min-max control,[12] this approximate approach results in a total cost of only 3.5 percent above optimum on the average.[13]

The Kardex card shown in Figure 9-15 is a record of the transactions for a particular grade of bond paper sold by an office supply distributor. Note the min and max values in the lower right-hand corner of the card. When the on-hand quantity drops to 125,000

Figure 9-16 A Min-Max System of Inventory Control, a Variant of the Reorder Point System

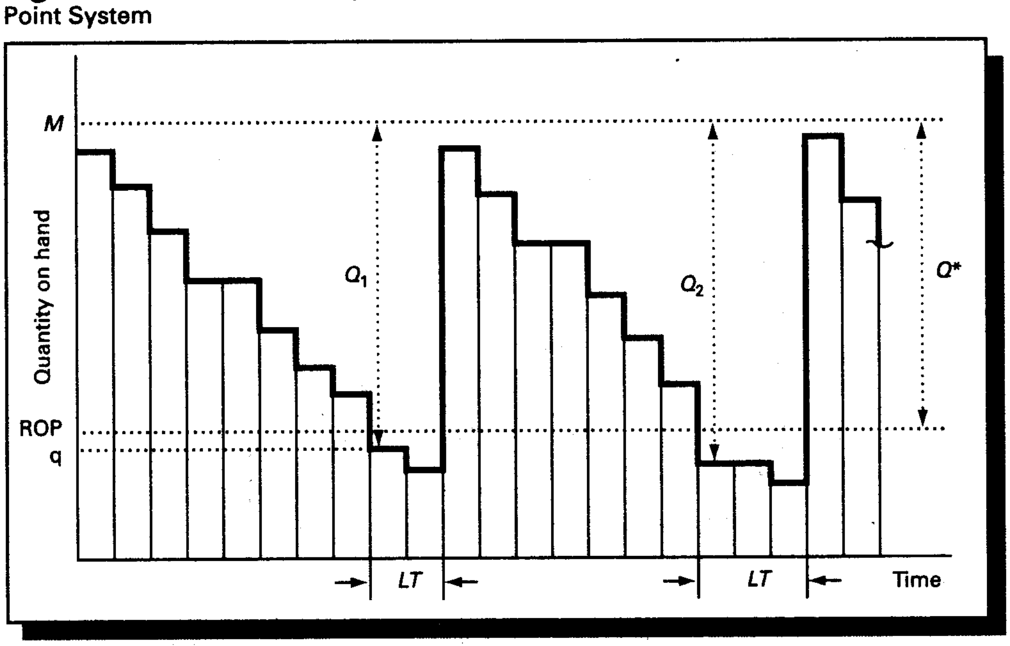

[12]Rein Peterson and Edward Silvers, *Decision Systems for Inventory Management and Production Planning* (New York: John Wiley & Sons, 1979), pp. 540–543.
[13]B. Archibald, "Continuous Review (s,S) Policies for Discrete Compound Poisson Demand Processes" (unpublished Ph.D. diss., University of Waterloo, 1976).

units, an order should be placed for 250,000 – 125,000 = 125,000 units. Note in the record that reorder quantity is not the 125,000 expected. Why? The company is ordering this item jointly with others from the same paper mill. Order minimum sizes are typically required such that when one item reaches its reorder point, the joint order may be "filled out" with other items that have not yet reached their *ROP*. In this way, the company forces single item, reorder point control to operate in a joint-ordered environment.

Although not necessarily better than just-in-time or quick response, the min-max approach to inventory control is an appropriate method to be used when demand is lumpy, or erratic. Lumpy demand is often associated with slow-moving items, but not necessarily limited to them. Actually, the lumpy demand characteristic may be seen in up to 50 percent of the product line items for many firms. Using what we have learned, the min-max approach can be slightly modified as follows to apply to items with lumpy demand:

1. Forecast demand by simply averaging the demand by period over at least 30 periods, if that much information is available. Compute the demand standard deviation over these same periods. If the standard deviation is greater than the average demand, declare demand to be lumpy and proceed to the next step.
2. Calculate the order quantity in any of the appropriate ways as previously discussed.
3. Because the on-hand quantity can drop significantly below the reorder point at the time that the order is placed, we adjust the *ROP* to compensate for it. That is, in addition to the demand during the lead time plus the safety stock that usually make up the *ROP*, we now add the *expected deficit* to *ROP*, which is the average amount that the quantity on hand is likely to fall before a replenishment order is placed. Refer to Figure 9-17.

Figure 9-17 Min-Max Inventory Control Under Lumpy Demand

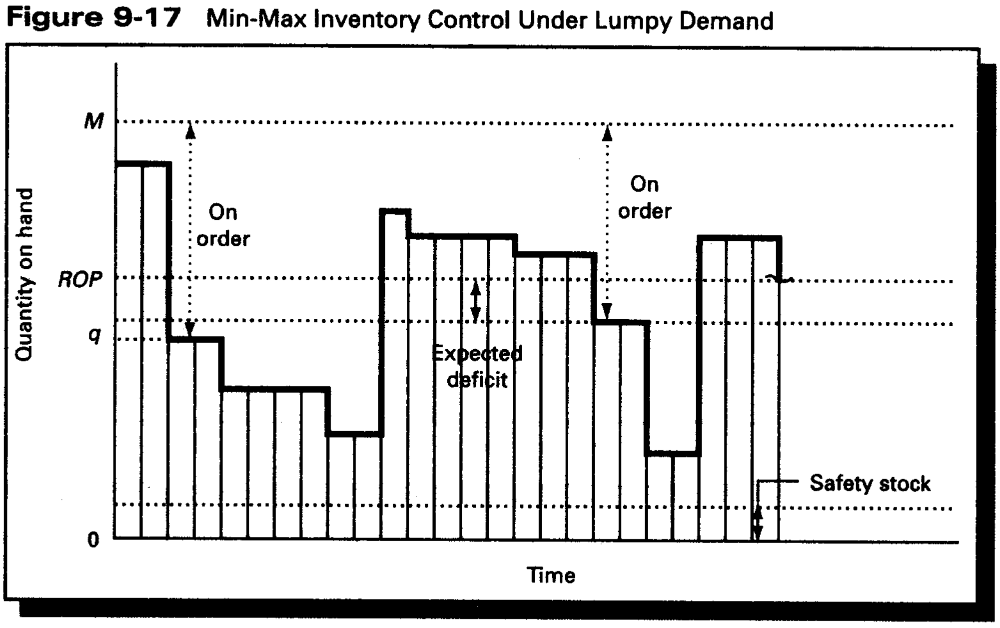

4. Approximate the expected deficit (average period sales) as one-half of the beginning and ending quantity on hand between quantity on hand record updates.
5. Set the max level as the *ROP* quantity *plus* the order quantity *less* the expected deficit.
6. Execute the min-max control system in the normal fashion as previously described. That is, when the effective inventory level falls to the *ROP* quantity, place an order for an amount equal to the difference between the max level (M^*) and the quantity on hand (q).

Example

The weekly requests for an item in inventory show a demand rate of $d = 100$ units and a standard deviation of $s_d = 100$ units. The item costs $1.45, procurement costs are $12 per order, annual carrying costs are 25 percent, and order lead time is 1 week. The in-stock probability during the lead time is to be at least 85 percent. The quantity on hand is updated daily, and the average daily sales quantity is 10 units, where an approximation for the expected deficit is $ED = 10$ units.

Since $s_d \geq d$, the item is believed to have a lumpy demand pattern. The order quantity can be found as

$$Q^* = \sqrt{\frac{2DS}{IC}} = \sqrt{\frac{2(100)(52)(12)}{0.25(1.45)}} = 587 \text{ units}$$

The ROP is

$$ROP = dLT + z(s_d') + ED = 100(1) + 1.04(10) + 10 = 214 \text{ units}$$

where

$$z_{@0.85} = 1.04 \text{ from Appendix A}$$

$$s_d' = s_d\sqrt{LT} = 100\sqrt{1} = 100 \text{ units}$$

The maximum level is

$$M^* = ROP + Q^* - ED = 224 + 587 - 10 = 801 \text{ units}$$

Many times there are specific reasons why lumpiness occurs. Occasional spikes of high customer demand may be predicted with a high degree of accuracy. Hence, a great deal of inventory may be avoided. R. G. Brown provides an excellent illustration of this idea.

Example

In the U.S. Navy, there was an O-ring seal used in the boiler tubes on a particular class of carrier. The period demand history looked something like 0 0 1 3 2 0 0 1 307 0 1 0 0 4 3 5 307 0 3 1 0 0 3 307. . . . The demand certainly appeared to be lumpy.

However, most of the demand was in the single digits, with an occasional demand for 307 pieces. This large demand occurred when an overhaul was carried out in a shipyard, and overhauls were scheduled up to two years in advance.[14]

Stock-to-Demand

Sometimes companies prefer methods that are inherently simple to understand and easy to implement. Overall, such methods may provide better control if they are diligently followed, compared with the more elegant statistical methods of control. The stock-to-demand method is one such practical approach to pull inventory management.

The stock-to-demand method may be paraphrased as follows. At a specified time, a forecast is made for the item's demand rate. The forecast is multiplied by a factor that represents the review interval, the lead time for replenishment, and a time increment representing uncertainty in the demand forecast and lead time to obtain a target quantity. The on-hand quantity is noted at the time of the forecast and an order is placed for the difference between the target quantity and the on-hand quantity. Stock-to-demand inventory control is a periodic review system type.

Example

A materials manager for a large insurance company makes a forecast every month of the paper supplies needed by the office staff. For a particular month, copy machine paper usage is forecasted to be 2,000 reams. Inventory records show that 750 reams are currently on hand with none on order and none committed to users. It takes one week to receive an order placed with the paper distributor. The manager likes to have the equivalent of an extra week's demand on hand as safety stock.

The forecasted demand is multiplied by a factor of 6/4 that is calculated as follows:

Forecast/review interval	4 weeks
Lead time	1 week
Safety stock	1 week
Total	6 weeks

Since the forecast represents four weeks demand, total time is divided by the forecast interval. The order quantity is 2,000(6/4) – 750 = 2,250 reams.

Multiple Item, Multiple Location Control

The problem of inventory control in practice is truly large scale, often involving hundreds of products located at numerous stocking points that are served from multiple

[14]Brown, *Materials Management Systems*, p. 250.

plants. Different modes of transportation may be used to move product between plants and stocking points. Although inventory control may be handled as a number of single item, single location problems, an integrated approach may be used to focus on some important economic concerns, such as shipping in full truckload quantities or producing in economic lot sizes. Consider how a chemical company dealt with its inventory control problem.

Application

A manufacturer of industrial cleaning compounds used by restaurants, hospitals, car washes, manufacturers, and schools sold throughout the country over 200 products represented by more than 750 line items. The items were stocked in nearly 40 warehouses, but not all items were stocked in all warehouses. Of the $220 million annual soap products sales, 70 percent were handled through the warehousing system. A computerized inventory control system was developed that attempted to control the inventory levels in a manner as shown in Figure 9-18. Consider how it worked.

Each item in a warehouse was forecasted on a monthly basis using exponential smoothing. The forecasts were staggered throughout the month to balance the workload on the computing system. The item amount on hand from warehouse computerized inventory records was checked daily.

Figure 9-18 Control of an Item in a Multiple-Item, Multiple-Location Inventory Control System for a Specialty Chemical Company

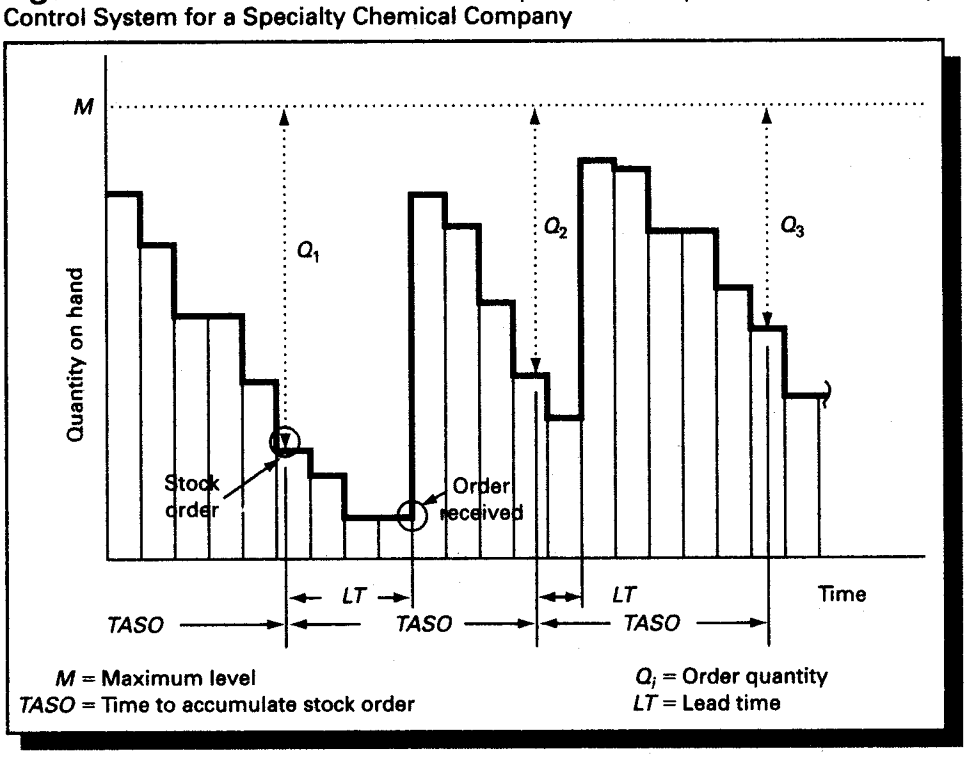

Accumulating a truckload quantity was the main economic force in the design of the inventory control system. For all the items in the warehouse collectively, the *T*ime to *A*ccumulate a *S*tock *O*rder (*TASO*) was computed as a truckload weight divided by the demand rate for the items stocked in the warehouse. Using this average review time, a maximum level was determined for each item.

Once a month, as the item forecast for a warehouse was made and the item stock level checked, a summation of the deficits between the item maximum level and its quantity on hand was determined. If the accumulated differences were greater than or equal to a truckload quantity, a replenishment order was placed on the serving plant. Although there was not precise control of each item, nevertheless major economies were achieved.

There were some side rules on the control process that helped it run smoothly. First, to prevent very small quantities of a particular item being included in an order, the deficit of an item had to be more than 10 percent of its max level. Second, to prevent an item from being out of stock in a warehouse when it alone might not have a deficit that would fill a truck, the inventory manager was issued a low-stock report, which showed that if the current demand rate continued, the item would be out of stock before the next expected warehouse replenishment shipment. The manager could then take action to replenish the item outside of normal ordering procedures, if it was desirable. Third, new items in the warehouse were not forecasted using exponential smoothing until they had accumulated at least a six-month sales history. Salespersons provided the interim forecasts. Stock-status reports, out-of-stock reports, forecasts, and shipping reports are examples of the report types that this system can produce.

Multi-Echelon Control

Recall in Figure 9-1 that inventories were located throughout the supply channel. These inventories are rarely independent of each other. That is, inventories at retail are backed up by inventories at their serving warehouses. In turn, warehouse inventories are backed up with inventories at the plants. If substantial amounts of inventory are maintained in field warehouses, then less may be needed at the next downstream echelon in the channel, namely, retail outlets, to maintain the same overall level of product availability. Managing the inventories throughout the *entire* channel becomes the important issue rather than the management at individual, independent stocking points.

Approximation to good multi-echelon inventory control is made using a base stock control system. The basis for this system is for any echelon in the supply channel to plan its stocking level on its inventory position *plus* the inventory from *all* downstream echelons. That is, planning the inventory level for a particular echelon is not determined from the demand information derived from just the next downstream echelon, but rather on the demand from the end customer. There is less demand variability for the particular echelon when end demand can be used in an upstream echelon inventory planning process. The demand characteristic throughout a supply channel: There is greater lumpiness at each echelon the further the echelon is upstream from the end customers. Inventory planning based only on orders

Figure 9-19
A Frequently Occurring Multi-Echelon Supply Channel

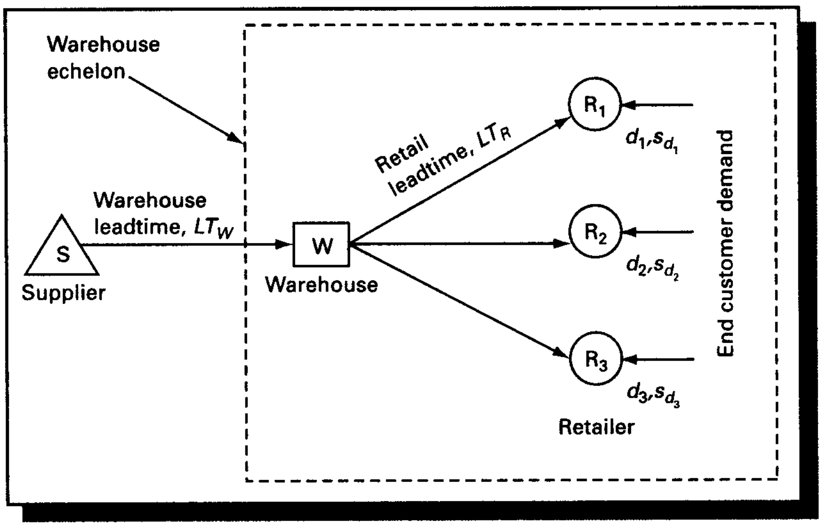

from the next downstream echelon results in more safety stock than from planning using end customer demand.

A simple two-echelon supply channel might be like that shown for a warehouse-retailer channel as shown in Figure 9-19. Retailers serve the end customers from their inventories and the warehouse replenishes retailer stocks. In a base stock control system, retailer inventory levels are controlled using any appropriate method, such as reorder point control. Demand information for a retailer is derived from the end customers in the retailer's territory. The inventory position for a retailer is the quantity on hand plus the quantity on order from the warehouse.

For one echelon upstream from the retailers (the warehouse echelon), the demand for planning purposes is derived from the aggregation of the end customer demand on all retailers. The inventory position for the warehouse echelon, but not the warehouse itself, is the sum of the inventory at the retailers, the inventory in the warehouse, and the inventory in transit (on order) to and from the warehouse. The reorder point and reorder quantities are determined for the inventory position of the echelon and not for the warehouse itself. The average stocking level in the warehouse is found by subtracting the retailers' average inventory levels from the echelon inventory, assuming there is negligible inventory in transit.

The base stock system approach to inventory planning can be continued for additional echelons within the supply chain. Remember to plan inventory levels for any echelon based on end item demand and not on the orders from the next echelon downstream.

Example

Suppose a portion of a distribution network is illustrated in Figure 9-19. Retailers forecast their end customer demand for their particular territories. For a particular item, the retailers' monthly demand (normally distributed) is shown in Table 9-5.

	JAN.	FEB.	MAR.	APR.	MAY	JUNE	JULY	AUG.	SEPT.	OCT.	NOV.	DEC.	AVG.	STD. DEV.
Retailer 1	218	188	225	217	176	187	221	212	210	203	188	185	202.5	16.8
Retailer 2	101	87	123	101	95	97	93	131	76	101	87	114	100.5	15.6
Retailer 3	268	296	321	312	301	294	285	305	289	303	324	332	302.5	18.0
Combined	587	571	669	630	572	578	599	648	575	607	599	631	605.5	32.4

Table 9-5 Typical Retailer Monthly Demand and Combined Demand for Warehouse Echelon

The item has a value of $C_R = \$10$ per unit at the retail level and $C_W = \$5$ per unit at the warehouse level. Carrying costs are $I = 20\%$ per year. The cost to process a replenishment order for a retailer is $S_R = \$40$ per order and $S_W = \$75$ per order at the warehouse. The retailers' lead times are all one week ($LT_R = 0.25$ months) and the warehouse lead time is two weeks ($LT_W = 0.5$ months). An in-stock probability during the lead time of 90 percent is used for both warehouse and retailers. Using a reorder point inventory control method, find the reorder points and order quantities for both retail and warehouse echelons. How much inventory is needed at the warehouse?

First, compute the inventory policy for each retailer. For retailer 1, the order quantity (Q) is

$$Q_1 = \sqrt{\frac{2D_{R_1}S_R}{IC_R}} = \sqrt{\frac{2(202.5 \times 12)(40)}{0.20(10)}} = 311.8, \text{ or } 312 \text{ units}$$

The reorder point (ROP) is

$$ROP_1 = d_1 \times LT_R + zs_{d_1}\sqrt{LT_R} = 202.5 \times 0.25 + 1.28 \times 16.8\sqrt{0.25} = 61.38, \text{ or } 61 \text{ units}$$

The average inventory (AIL) is

$$AIL_1 = \frac{Q_1}{2} + zs_{d_1}\sqrt{LT_R} = \frac{311.8}{2} + 1.28 \times 16.8\sqrt{0.25} = 166.65, \text{ or } 167 \text{ units}$$

The inventory control rule: When the inventory level at retailer 1 falls to 61 units, place a replenishment order for 312 units.

Repeat the previous calculations for the remaining two retailers. The results are summarized in Table 9-6. The retailer echelon inventory is 167 + 120 + 202 = 489 units.

Next, compute the warehouse's inventory policy. Find the demand properties for the warehouse echelon by combining the retailers' demand, as shown in Table 9-5. The warehouse *echelon* order quantity is

$$Q_W = \sqrt{\frac{2D_W S_W}{IC_W}} = \sqrt{\frac{2(605.5 \times 12)(75)}{0.20(5)}} = 1{,}043.98, \text{ or } 1{,}044 \text{ units,}$$

Table 9-6 Inventory Statistics for Retailers

	RETAILER 1	RETAILER 2	RETAILER 3
Reorder quantity, *Q*	312	220	381
Reorder point, *ROP*	61	35	87
Average inventory, *AIL*	167	120	202

the *ROP* is

$$ROP_W = d_W \times LT_W + zs_W\sqrt{LT_W} = 605.5 \times 0.5 + 1.28 \times 32.4\sqrt{0.5} = 332.03, \text{ or } 332 \text{ units}$$

and *AIL* for the warehouse echelon is

$$AIL_W = \frac{Q_W}{2} + zs_W\sqrt{LT_W} = \frac{1,043.98}{2} + 1.28 \times 32.4\sqrt{0.5} = 551.32, \text{ or } 551 \text{ units}$$

However, the expected warehouse inventory is the warehouse echelon inventory less the retail echelon inventory, or 551 − 489 = 62 units. There is no inventory assumed to be in the pipeline.

The warehouse inventory control policy is to monitor the warehouse echelon inventory, which is the total of the inventory at each retailer, the inventory held in the warehouse, the inventory on order by the warehouse, and the inventory on order by the retail outlets less any inventory committed to end customers but not yet deducted from retail inventory. When this inventory echelon position falls to 332 units, place an order with the supplier for 1,044 units.

When the multi-echelon problems become too complex for the previous type of mathematical analysis, especially when more than two echelons are involved, computer simulation is an alternative. Simulations of this type are constructed from general simulation languages such a SLAM, DYNAMO, or SIMSCRIPT, or they may be conducted using custom packages such as Long Range Environmental Planning Simulator (LREPS)[15] or PIPELINE MANAGER.[16] The SCSIM module in LOGWARE software that accompanies this text illustrates this capability. The action of these simulators is to generate demand over time in a manner similar to that actually experienced by the operating channel. The product flows that take place to serve the demand are replicated. The product movement through the channel is observed, and statistics relating to product movement, inventory levels, stockouts, production rates, and transport shipments are reported. Alternative inventory policies can be tested through rerunning the simulation with different inventory stocking rules and service levels. Costs of alternatives can then be compared.

---

[15]Donald J. Bowersox, Omar K. Helferich, Edward J. Marien, Peter Gilmour, Michael L. Lawrence, Fred W. Morgan, Jr., and Richard T. Rogers, *Dynamic Simulation of Physical Distribution Systems* (East Lansing, MI: Division of Research, Graduate School of Business Administration, Michigan State University, 1972).
[16]Developed by Arthur Andersen & Company.

PIPELINE INVENTORIES

Pipeline inventories are in-transit stocks that reside in transportation equipment *moving* between inventory holding points. Management of these is a matter of controlling the time in transit, mainly through transport service selection. The inventories in transit can be surprisingly high, and good management can yield impressive cost reductions.

Example

A manufacturer of auto parts has assembly operations in the United States. It purchases components from companies located in the Pacific Rim and distributes them primarily in the United States. A diagram of the supply channel is shown in Figure 9-20. The average unit value of the product flowing in the channel is $50. Sales are 1,000 units per day. Carrying cost is 30 percent per year. The current in-transit inventories can be summarized as follows:

Pipeline	Days	Pipeline Inventory
Vendors to factory	21	21,000 units
In process at factory	14	—
Factory to warehouses	7	7,000
Warehouse storage	42	—
Warehouse to customers	3	3,000
Totals	87	31,000 units

Figure 9-20 A Typical Supply Channel Showing In-Transit Times

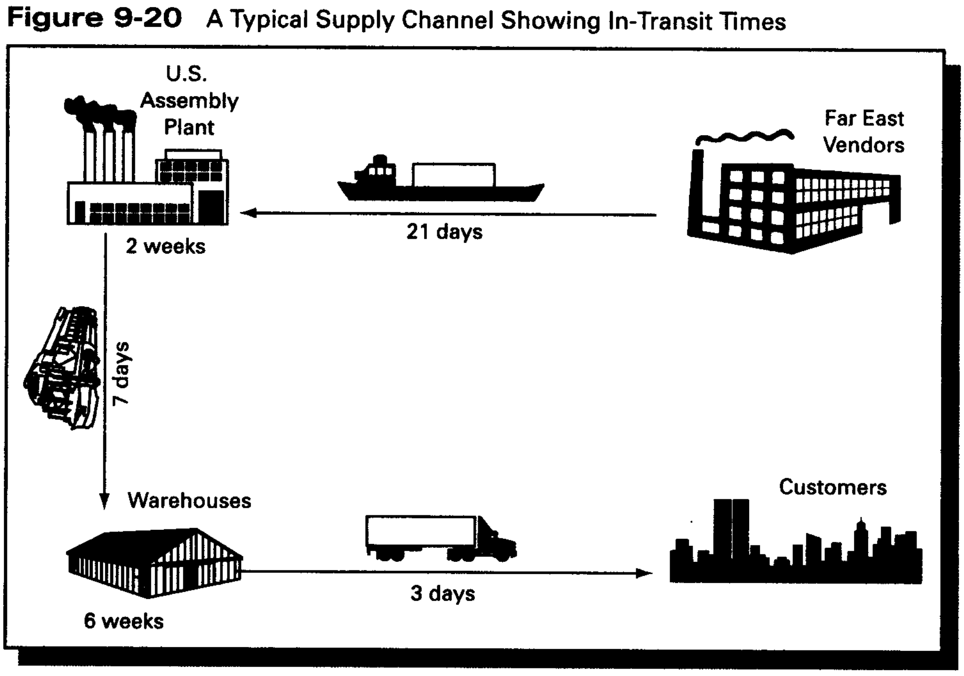

The total value of inventory in transit is 50 × 31,000 = \$1,550,000 and a carrying cost of 0.30 × 1,550,000 = \$465,000 per year.

If airfreight is used, the transit time between Far East vendors and the factory can be reduced to four days, most of which is ground handling. This would save 21 − 4 = 17 days in the pipeline, \$50 × 17,000 = \$850,000 in inventory value, and 0.30 × 850,000 = \$255,000 in annual carrying charges. This potential cost savings must be weighed against increased costs of using airfreight.

Reducing the average transit time in the pipeline usually has a coincident effect of reducing the transit time variability as well. Since transit time is a significant component of order lead time, the safety stocks in inventories will be lowered as an indirect benefit of reducing transit time uncertainty.

The annual in-transit inventory-carrying cost associated with a single link in the supply channel is calculated from

$$\text{In-transit inventory carrying cost} = \frac{ICDt}{365} \qquad (9\text{-}26)$$

where

I = annual carrying for product in transit, %/year
C = value of the product at the pipeline point in the supply channel, dollars/unit
D = annual demand, units
t = time in transit, days
365 = number of calendar days in a year

Note that I may be different from that for a stocking point since it does not need to include operating costs associated with storage. On the other hand, there may be operating costs for the transport of product in the pipeline, especially if private transportation is used. The costs in C should be those for holding the product and not for transporting it.

Example

Automobiles are imported to the United States through Boston, Massachusetts via Emden, Germany. The value of the automobile at the exit port in Germany is \$9,000. Carrying cost is primarily the cost of the capital tied up in the vehicles, or 20 percent per year. The average sailing time to the United States is 10 days. The in-transit inventory cost per vehicle is found from $ICt/365 = (0.20)(9{,}000)(10)/365 = \$49.32$ per automobile.

AGGREGATE CONTROL OF INVENTORIES

Top management is frequently more interested in the total amount of money tied up in inventories and the service levels for broad item groups than in the control of individual items. Although carefully setting the policy for each item does provide precise control of individual item inventories as well as inventories taken together, management at this level of detail for general planning purposes becomes too cumbersome. Therefore, methods that collectively control items in groups have had a place among inventory control procedures. Turnover ratios, ABC product classification, and risk pooling are a few of the methods used to control aggregated inventories.

Turnover Ratios

Perhaps the most popular aggregate inventory control procedure is the turnover ratio. It is a ratio of the annual sales on inventory to the average investment in inventory for the same sales period, where sales and inventory investment are valued at the echelon in the logistics channel where the items are held. That is,

$$\text{Turnover ratio} = \frac{\text{Annual sales at inventory cost}}{\text{Average inventory investment}} \tag{9-27}$$

The popularity of the measure undoubtedly stems from the ready availability of data (the company's financial statements) and the simplicity of the measure itself. Different turnover ratios may be specified for different product classes of, or for the entire, inventory. As a point of reference, the inventory turnover ratios for manufacturers, wholesalers, and retailers are 9:1, 9:1, and 8:1, respectively.[17]

By specifying the turnover ratio to be achieved, the overall inventory investment is controlled relative to the level of sales. It is appealing to have inventory investment change with the sales level; however, using the turnover ratio causes inventories to vary *directly* with sales. This is a disadvantage since we normally expect that inventories increase at a decreasing rate due to economies of scale. There is a price to be paid for simplicity!

A study of 100 major manufacturing companies in India with sales above Rs. 500 crore during 2000–2005 reveals that the turnover ratio of these companies has increased to 8.9 in 2004–2005 from 8.5 in 2000–2001.[18] This implies better management of inventories by the companies. The total inventories of these 100 companies has increased by 62.64 percent to Rs. 97,289 crore in 2004–2005 (Rs. 59,819 crore in 2000–2001), while the sales have increased by 70.5 percent during 2000–2005. Of the 100 companies, 69 have recorded an increase in the ratio, while the remaining 31, a decline. Companies in the consumer goods and engineering industries have registered a significant increase in the ratio. The worst performers in 2004–2005 include Indian telephone industries, Oswal Chemicals, Cipla, and Arvind Mills.

[17] *Statistical Abstract of the United States: 2001*, pp. 623, 644, and 657.
[18] *Financial Express* (March 21, 2006), p. 18.

ABC Product Classification

A common practice in aggregate inventory control is to differentiate products into a limited number of categories and then to apply a separate inventory control policy to each category. This makes sense since not all products are of equal importance to a firm in terms of sales, profit margin, market share, or competitiveness. By selectively applying inventory policy to these different groups, inventory service goals can be achieved with lower inventory levels than with a single policy applied collectively to all products.

It is well known that product sales display a life-cycle phenomenon where sales begin at product introduction with low levels, increase rapidly at some point, level off, and finally decline. A firm's products are usually in various stages of their life cycles and, therefore, are contributing disproportionately to sales and profits. That is, a few items may be contributing a high proportion of the sales volume. This disproportionate relationship between the percent of items in inventory and the percent of sales has generally been referred to as the 80-20 principle, although rarely does exactly 20 percent of the items in a product line represent 80 percent of the sales. The 80-20 principle serves as a basis for the *ABC* classification of items. *A* items are typically the fast movers, *B* items the medium movers, and *C* items the slow movers. There is no precise way that the items are grouped into one category or another, or even of determining the number of categories to use. However, rank ordering the items by sales and then dividing them into a few categories is a start. Some of the items are reassigned to other categories as their importance dictates. Inventory service levels are then assigned to each category. The development of the *ABC* product classification scheme is more thoroughly discussed in Chapter 3.

Table 9-7 Annual Sales Data for Sorensen Research Company

	Number of Units	Volume, Dollars	Product Type
INTRASET	1,000,000	$ 2,500,000	Catheter
SUBCLAVIAN II	250,000	137,000	Catheter
SUBVLAVIAN	150,000	975,000	Catheter
JUGULAR II	300,000	300,000	Catheter
CATHASPEC	100,000	150,000	Catheter
IV-SET	700,000	1,000,000	Catheter
CENTRI-CATH	500,000	3,500,000	Catheter
IV-12	15,000	74,700	Catheter
CSP	1,000,000	750,000	Catheter
Pressure Cuff	600,000	972,000	Catheter support
Pressure Tubing	25,000	825,000	Catheter support
EZE-FLO	4,200	65,100	Catheter support
REGUFLO	1,000,000	5,000,000	Catheter support
TRUSET	2,850,000	7,115,000	Catheter support
INTRAVAL	10,000	8,300	Catheter support
VACUFLO	355,000	350,000	Fluid suction
COLLECTAL Canisters	40,000	54,800	Fluid suction
COLLECTAL Liners	393,000	727,000	Fluid suction
	9,292,200	$24,503,900	

	Item Number	Cum. percent of Items	Volume, Dollars	Cum. percent of Sales	Item Class
TRUSET	1	5.56%	$ 7,115,000	29.04%	A
REGUFLO	2	11.11	5,000,000	49.44	
CENTRI-CATH	3	16.67	3,500,000	63.72	B
INTRASET	4	22.22	2,500,000	73.93	
IV-SET	5	27.78	1,000,000	78.01	
SUBVLAVIAN	6	33.33	975,000	81.99	
Pressure Cuff	7	38.89	972,000	85.95	
Pressure Tubing	8	44.44	825,000	89.32	
CSP	9	50.00	750,000	92.38	
COLLECTAL Lin.	10	55.56	727,000	95.35	C
VACUFLO	11	61.11	350,000	96.78	
JUGULAR II	12	66.67	300,000	98.00	
CATHASPEC	13	72.22	150,000	98.61	
SUBCLAVIAN II	14	77.78	137,000	99.17	
IV-12	15	83.33	74,700	99.48	
EZE-FLO	16	88.89	65,100	99.74	
COLLECTAL Can.	17	94.44	54,800	99.97	
INTRAVAL	18	100.00	8,300	100.00	
			$24,503,900		

Table 9-8 Items Sorted in Descending Order Using Sales

Example

The Sorensen Research Company produces a limited line of high technology products for hospital use. The main products are arterial catheters (INTRASET); catheter-support devices (REGUFLO); and fluid-suction systems (VACUFLO, COLLECTAL). Annual sales data are summarized in Table 9-7.[19]

For inventory control reasons, suppose these items are classified into three groups. The *A* items are to represent approximately the top 10 percent of items, *B* items are to be about the next 40 percent, and the *C* items are the remaining 50 percent. Table 9-7 is sorted in descending order according to item dollar sales. Computing the cumulative percent of items and the cumulative percent of sales on the sorted data yields Table 9-8:

Scanning down the cumulative percent of items column until approximately 10 percent of the items are accumulated will represent the *A* item category. Due to the small number of items, we cannot find exactly 10 percent. We may choose to round up. Next is the break point for *B* items, which is where the cumulative percent of items is 50 percent. We can now see that *A* items, or 11 percent of the items, account

[19]Disguised data as reported in Sorensen Research Company, Harvard Business School Case 9-677-257 prepared under the direction of Steven C. Wheelwright.

for 49 percent of the sales. *B* items, or 50% – 11% = 39%, of the items account for 92% – 49% = 43% of the sales. *C* items, representing 50 percent of the items, account for only 100% – 92% = 8% of the sales. Service levels can be set for these categories according to the importance of each to the company and to its customers.

Risk Pooling

Aggregate inventory level planning often involves projecting how stocking point inventory levels will change with changes in the number of stocking locations and their throughputs. In planning a logistical network, it is common to expand or contract the number of stocking points to meet customer service and cost objectives. As the number of locations is changed or even as the sales are reassigned among existing locations, the inventory in the system does not remain constant due to the risk pooling, or consolidation, effect. Risk pooling suggests that if inventories are consolidated into fewer locations, their levels will be reduced. Expanding the number of inventory locations has the opposite effect. System inventory levels are a result of balancing regular stock, which is affected by inventory policy, and safety stock, which is affected the degree of uncertainty in demand and lead time.

Illustration

Suppose that a product is stocked in two warehouses. Average territory monthly demand in warehouse 1 is d_1 = 41 units with a standard deviation s_{d_1} = 11 units/month. And for warehouse 2, d_2 = 67 and s_{d_2} = 9. Inventory replenishment quantities are determined using the economic order quantity formula. The replenishment lead time for both warehouses is 0.5 months, and the product value is $75 per unit. The replenishment order cost is $50 and the inventory-carrying cost is 2 percent per month. The probability of being in stock during the lead time is set at 95 percent. What inventory benefit would there be to consolidating the inventories into one warehouse?

First we estimate the regular and safety stock in the two warehouses.

Regular Stock. Compute the average amount of regular stock.

$$RS = \frac{Q}{2} = \frac{\sqrt{\frac{2dS}{IC}}}{2}$$

$$RS_1 = \frac{\sqrt{\frac{2(41)(50)}{0.02(75)}}}{2} = 26 \text{ units}$$

$$RS_2 = \frac{\sqrt{\frac{2(67)(50)}{0.02(75)}}}{2} = 33 \text{ units}$$

Regular system inventory for two warehouses is $RS_S = RS_1 + RS_2 = 26 + 33 = 59$ units.

Now compute the regular stock if held in one central warehouse. The average demand for the central warehouse is $d_C = d_1 + d_2 = 41 + 67 = 108$. Now,

$$RS_C = \frac{\sqrt{\frac{2(108)(50)}{0.02(75)}}}{2} = 42 \text{ units}$$

Safety Stock. Safety stock in two warehouses is found as follows.

$$SS = zs_d\sqrt{LT}$$
$$SS_A = 1.96(11)\sqrt{0.5} = 15.25 \text{ units}$$
$$SS_B = 1.96(9)\sqrt{0.5} = 12.47 \text{ units}$$

System safety stock in two warehouses is $SS_S = SS_A + SS_B = 15.25 + 12.47 = 27.72$, or 28 units.

For the safety stock in the central warehouse, estimate the demand standard deviation from

$$s_C = \sqrt{s_1^2 + s_2^2} = \sqrt{11^2 + 9^2} = 14.21.$$

Now, the safety stock is

$$SS_C = 1.96(14.21)\sqrt{0.5} = 19.69, \text{ or } 20 \text{ units}$$

Total inventory is the sum of regular and safety stocks. For two warehouses $AIL_2 = 59 + 28 = 87$ units. In the central warehouse, $AIL_C = 42 + 20 = 62$ units. Note that regular and safety stock have decreased through consolidation.

Square Root Rule. The square-root rule is a well-known method for determining the consolidation effect on inventories. However, it measures only the regular stock reduction, not both regular and safety stock effects as described in the previous section. Assuming that an inventory control policy based on the *EOQ* formula is being followed and that all stocking points carry the same amount of inventory, the square-root rule can be stated as follows:

$$AIL_T = AIL_i\sqrt{n} \tag{9-28}$$

where

AIL_T = the optimal amount of inventory to stock, if consolidated into one location in dollars, pounds, cases, or other units

AIL_i = the amount of inventory in each of n locations in the same units as AIL_T

n = the number of stocking locations before consolidation

Note that inventory varies with the number of stocking points in the logistics network.

Example

Sorensen Research Company operated 16 regional public warehouses. Each warehouse carried \$165,000 of inventory on the average. If all stocks are consolidated into one location at the plant, how much inventory can be expected?

Using Equation (9-28), we calculate

$$AIL_T = \$165,000\sqrt{16} = \$660,000$$

Note that the previous system of stocking points had a total of 16 × 165,000 = \$2,630,000 in inventory investment.

Example

Suppose that Sorensen wishes to consolidate inventories into two locations that equally divide the stock. How much inventory can be expected in each warehouse?

We already know that one location should have \$660,000 in inventory investment. Now, we simply need to estimate from this value the amount to stock in two warehouses. By algebraically manipulating Equation (9-28), the inventory in a multiple warehouse system would be

$$AIL_i = \frac{AIL_T}{\sqrt{n}}$$

Hence, for two warehouses the inventory in each would be

$$AIL_i = \frac{\$660,000}{\sqrt{2}} = \$466,690$$

Systemwide inventory is 2 × 466,690 = \$993,381. Thus, reducing from 16 to two warehouse saves \$2,630,000 – 933,381 = \$1,696,619 in inventory investment.

Inventory Throughput Curve. Although the square-root rule of inventory consolidation is generally useful, the assumptions that there are equal amounts of inventory in all warehouses, inventories consolidate precisely as the square root of the number of warehouses, demand and lead time are known for sure, and order quantity is determined from the *EOQ* formula may be too limiting. Using a slightly different approach, these limitations can be dropped. First, from the company's stock status reports, construct a plot of average inventory (AIL_i) level against annual warehouse shipments (D_i), as shown in Figure 9-21. Each point on the plot represents a single warehouse. It is the turnover ratio for the warehouse. From a family of curves of the form $AIL = aD^b$, fit the best curve possible to the data. The data in Figure 9-21 are for a specialty chemical company, and give

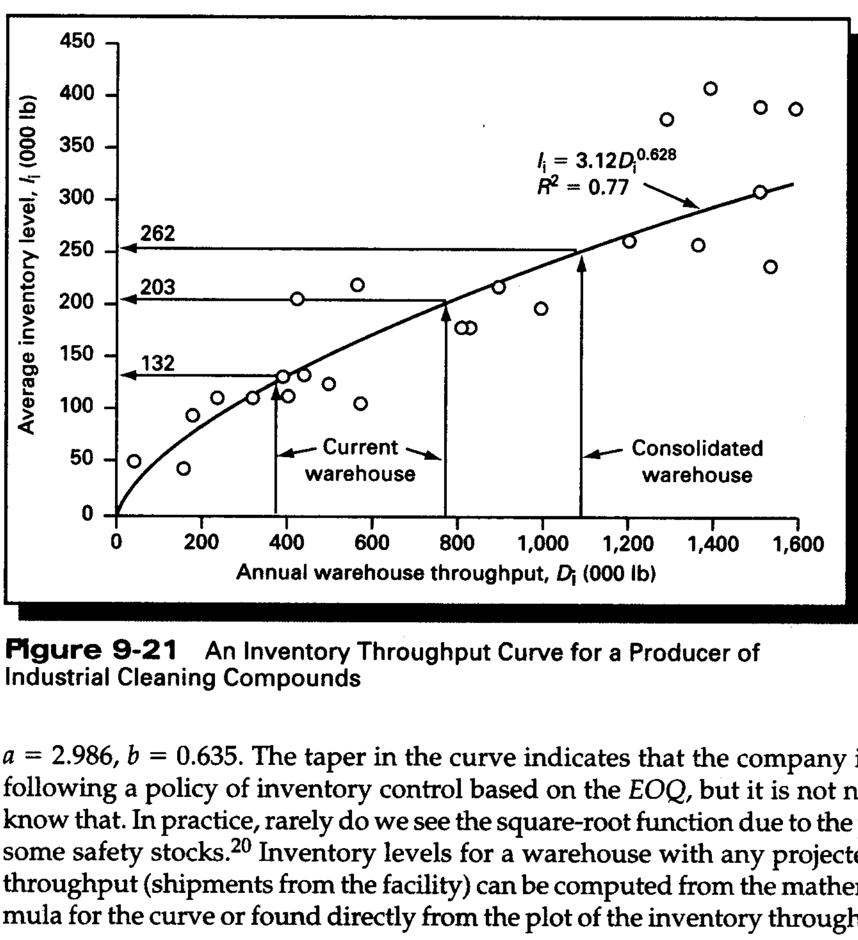

Figure 9-21 An Inventory Throughput Curve for a Producer of Industrial Cleaning Compounds

$a = 2.986$, $b = 0.635$. The taper in the curve indicates that the company is probably following a policy of inventory control based on the *EOQ*, but it is not necessary to know that. In practice, rarely do we see the square-root function due to the presence of some safety stocks.[20] Inventory levels for a warehouse with any projected demand throughput (shipments from the facility) can be computed from the mathematical formula for the curve or found directly from the plot of the inventory throughput curve.

Example

The specialty chemical company, whose data are represented in Figure 9-21, has 25 public warehouses and plants from which it distributes product. Suppose that two warehouses with 390,000 lb and 770,000 lb of throughput, respectively, are to be consolidated into a single warehouse with 390,000 + 770,000 = 1,160,000 lb of annual throughput. How much inventory should be stocked in the single warehouse?

From the plot in Figure 9-21, the inventory for the two current warehouses is 132,000 + 203,000 = 335,000 lb. Combining the throughput and reading the inventory from the plot gives 262,000 lb.

[20]Ronald H. Ballou, "Estimating and Auditing Aggregate Inventory Levels at Multiple Stocking Points," *Journal of Operations Management*, Vol. 1, No. 3 (February 1981), pp. 143–153; and Ronald H. Ballou, "Evaluating Inventory Management Performance Using a Turnover Curve" *International Journal of Physical Distribution and Logistics Management*, Vol. 30, No. 1 (2000), pp. 72–85.

Total Investment Limit

Inventories represent a substantial capital investment for many firms. Because of this, managers frequently will place a limit on the amount of inventory to be carried. The inventory control policy must then be adjusted to meet this goal, if total average inventory investment exceeds it. Suppose that inventory is controlled by a reorder point control policy under conditions of demand and lead time certainty. If a monetary limit is placed on all items carried at an inventory location, we can state that

$$\sum_i C_i \frac{Q_i}{2} \leq L \tag{9-29}$$

where

L = investment limit for i items in inventory, dollars
C_i = value of item i in inventory
Q_i = order quantity for item i in inventory

The order quantity can be determined from a modified Equation (9-7). When the average inventory value for all items exceeds the investment limit (L), the order quantities for the items need to be reduced to decrease the average item inventory levels and meet the investment limit. A reasonable way to do this is to artificially inflate the carrying cost I to a value $I + \alpha$; large enough to shrink the stock levels an appropriate amount. The basic economic order quantity formula is modified to be

$$Q_i = \sqrt{\frac{2D_iS_i}{C_i(I + \alpha)}} \tag{9-30}$$

where α is a constant to be determined. Equation (9-30) is substituted into Equation (9-29) and reworked to give a formula for α .

$$\alpha = \left(\frac{\sum_i \sqrt{2D_iS_iC_i}}{2L} \right)^2 - I \tag{9-31}$$

Once α is found, it is substituted into Equation (9-30) to find the revised Q_i.

Example

Suppose that an inventory contains three items. Management has placed a total dollar limit on the average inventory investment of $10,000 for these items. The annual carrying cost is 30 percent per year. Other relevant data are

Item i	Procurement Cost, S_i	Purchase Cost, C_i	Annual Demand, D_i
1	$50/order	$20/unit	12,000 units
2	50	10	25,000
3	50	15	8,000

We first compute each Q. That is,

$$Q_1 = \sqrt{\frac{2DS}{IC}} = \sqrt{\frac{2(12,000)(50)}{0.30(20)}} = 447.21 \text{ units}$$

Similarly,

$$Q_2 = \sqrt{\frac{2(25,000)(50)}{0.30(10)}} = 912.87 \text{ units}$$

$$Q_3 = \sqrt{\frac{2(8,000)(50)}{0.30(15)}} = 421.64 \text{ units}$$

Checking the total inventory investment by solving the left side of Equation (9-29) with the above computed values gives

$$\begin{aligned} \text{Inventory investment} &= C_1(Q_1/2) + C_2(Q_2/2) + C_3(Q_3/2) \\ &= 20(447.21/2) + 10(912.87/2) + 15(421.64/2) \\ &= \$12,199 \end{aligned}$$

Since the investment limit of $10,000 is exceeded, solve for α in Equation (9-31). That is,

$$\begin{aligned} \alpha &= \left(\frac{\sqrt{2(12,000)(50)(20)} + \sqrt{2(25,000)(50)(10)} + \sqrt{2(8,000)(50)(15)}}{2(10,000)}\right)^2 - 0.30 \\ &= 0.146 \end{aligned}$$

We now can substitute $\alpha = 0.146$ into Equation (9-30) and solve for the revised order quantity for each item. That is,

$$Q_1 = \sqrt{\frac{2D_iS_i}{C_i(I + \alpha)}} = \sqrt{\frac{2(12,000)(50)}{20(0.30 + 0.146)}} = 366.78 \text{ units}$$

Similarly, the order quantities for the other items are computed to be $Q_2 = 748.69$ units and $Q_3 = 345.81$ units. The average investment now comes to $10,004. Close enough!

Supply-Driven Inventory Control

There are situations where the previously discussed methodology does not seem appropriate. It assumes that supply and demand can be reasonably matched; however, there are situations where, in spite of management's best forecasting efforts, supply cannot be aligned well with demand. That is, supply may be so valuable that the producer will obtain all that is available. This can cause over- and understock in the supply channel. Little can be done when demand exceeds supply. On the other hand, when the producer pushes excess supply into the distribution channel, the producer has but one option to control excess inventory levels that can occur—increase demand to lower the inventory to more acceptable levels. Price discounting is commonly used as the variable to increase demand.

Applications

- StarKist operates a unique aggregate inventory-control system for its tuna products. Since the company is committed to purchase and pack all the tuna it can, the distribution system can become overloaded with finished product. To manage overstocking, the company will run sales on its products. Customers readily buy extra quantities of these highly prized products, thus reducing StarKist's inventory levels.
- The American Red Cross: Blood Services plans up to a year in advance for blood collections. Donors are highly valued and are not turned away, even if collections exceed anticipated levels or current blood needs. If inventory levels of specific blood types are high and outdating may occur, the Red Cross will either convert the whole blood to another blood product or reduce the price to their hospital customers. Price discounting is effective because hospitals obtain their blood supply from multiple sources, not just the Red Cross.
- KLG Systel Limited, headquartered in Gurgaon, has entered into an alliance with Sage ACCPAC India to offer e-point of sale (e-PoS) solution, which is likely to give the clients a complete view of their activities and inventory management. The e-PoS solution is priced Rs. 3.50 lakh for a single location setup. The solution is capable of creating quotations, order confirmations, invoices, credit notes, and receipts.[21]

VIRTUAL INVENTORIES[22]

With improvements in corporate information systems, it has become an increasing practice to serve customer demand from more than one stocking location. Although customers may be assigned to a primary stocking site, it is rare that enough inventory will be maintained to meet all demand requests from the primary location all of the time. Such an inventory policy is reasonable since the inventory cost to ensure never being out of stock is exceedingly high. Alternatively, demand will be cross filled from other sites holding the same items, as illustrated in Figure 9-22 for an inventory system of two stocking points. The combination of inventory sites is referred to as a *virtual inventory*. Cross filling demand from the multiple stocking locations in the virtual inventory leads to the expectation that demand fill rates will be increased, systemwide inventory levels will be lower, or both, as compared with meeting demand from only the customer's primary stocking location and incurring some stockouts and order filling delays.

The logistician's problem is to decide the items to cross fill and those that should only be filled from the primary location. The solution requires a balance between the

[21]*Business Line* (June 29, 2006), p. 4.
[22]Ronald H. Ballou and Apostolos Burenetas, "Planning Multiple Location Inventories," *Journal of Business Logistics*, Forthcoming.

Figure 9-22
Illustration of Inventory Cross Filling

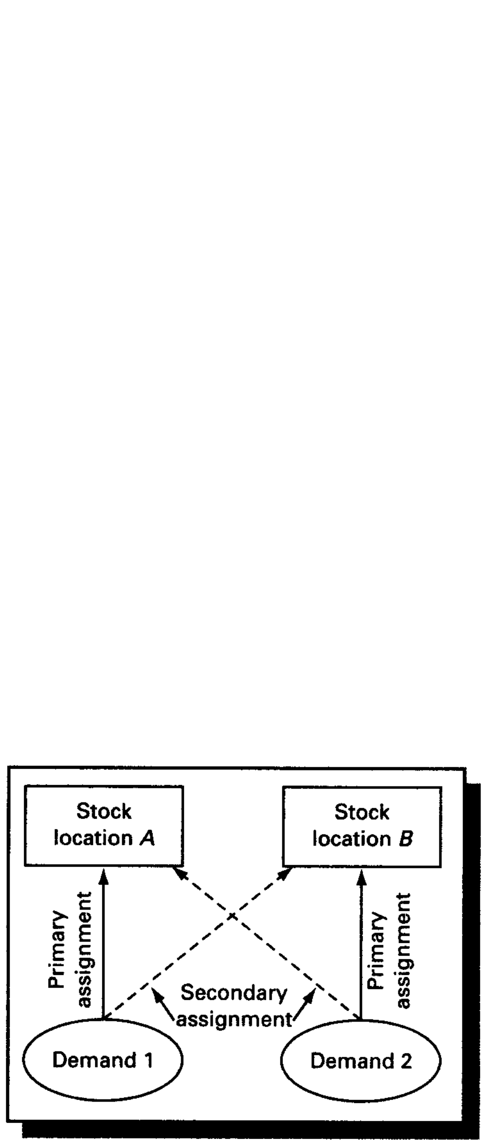

costs associated with regular stock versus safety stock. Recall that regular stock is the inventory to meet average demand and average lead time. Safety stock, on the other hand, is the extra inventory needed to meet uncertainty in demand and lead time. When it comes to cross filling, the economic forces associated with these two inventory types are opposed. That is, on a systemwide basis, regular stock increases with cross filling, whereas safety stock decreases. Consider how this happens.

Regular Stock

In a system of multiple stocking points, the maximum amount of inventory held system wide will occur when the demand on the stocking points is equally divided among them. On the other hand, when there is unequal demand dispersion, cross filling causes the effective demand throughout the system to be more balanced than was the primary demand. Effective demand refers to the demand actually realized on a stocking point through cross filling from other stocking points rather than the primary demand assigned to it.

To illustrate, suppose that there are two stocking points, as in Figure 9-22, with demand of 50 and 150 units per week, respectively, and an inventory fill rate of 90 percent. A 90 percent fill rate means that only 90 percent of each point's demand on the average can be filled from the primary location, with 10 percent being filled from the alternate location. Thus, the effective demand on location 1 is $50 \times 0.9 + 0.1 \times 150 = 60$ units per week. For location 2, the effective demand is $0.1 \times 50 + 0.9 \times 150 = 140$ units per week. Note that the effective demand has moved closer to the 100/100 demand split from the original 50/150 dispersion. If regular stock is determined from the EOQ formula in Equation (9-7), then the inventory at each stocking location is $AIL = K\sqrt{D}$, where K is a constant derived from the costs for a particular item. Therefore, the systemwide inventory (AIL_S) without cross filling is $AIL_S = \sqrt{50} + \sqrt{150} = 19.3$ units, where K is taken as one for illustration purposes. With cross filling, $AIL_S = \sqrt{60} + \sqrt{140} = 19.6$ units. Regular stock has increased.

Safety Stock

In a system of multiple stocking points, safety stock levels are affected by the fill rate and the dispersion of demand among the stocking locations. In contrast to regular

Table 9-9
Effective Demand Standard Deviation at Two Location, Units per Week

	STOCK, LOCATION A	STOCK, LOCATION B
Std. dev.	4.5[a]	0.5[b]
Std. dev.	1.5[c]	13.5
Combined std. dev.	4.7[d]	13.5

[a] $\sqrt{[FR(1-FR)^{N-1}]^2 s^2} = \sqrt{[0.9(1-0.9)^{1-1}]^2 5^2} = 4.5$

[b] $\sqrt{[FR(1-FR)^{N-1}]^2 s^2} = \sqrt{[0.9(1-0.9)^{2-1}]^2 5^2} = 0.45$, round to 0.5

[c] Round to 1.5

[d] $\sqrt{4.5^2 + 1.5^2} = 4.7$

stocks, minimum safety stocks occur when demand is balanced among the stocking locations.

To illustrate, suppose that the standard deviations (s) for regular stock demand given previously are 5 and 15 units per week, respectively, lead time (LT) is one week, and the fill rate is 90 percent. Recall that for a reorder point control system, safety stock (ss) can be estimated from $ss = zs\sqrt{LT}$, where $z = 1.28$ is from the normal distribution at 90 percent. (*Note:* It is assumed that the fill rate and the probability during the lead time or the probability during the lead time plus order review time depending on method of control are approximately the same.) With cross filling, the standard deviation for a particular location among multiple stocking points is $s_N = \sqrt{[FR(1-FR)^{N-1}]^2 s^2}$, for the Nth stocking point. The effective standard deviations for the two stocking locations are shown in Table 9-9.

The safety stock at location A without cross filling is $ss = zs\sqrt{LT} = 1.28(5)\sqrt{1} =$ 6.4 units, at B is 19.2 units, and for the entire system is 6.4 + 1.28 = 25.6 units. With cross filling, the safety stock at A is $1.28(4.7)\sqrt{1} = 6.0$ units, at B is 17.3 units, and for the system is 23.3 units. Cross filling gives a reduction of 25.6 − 23.3 = 2.3 units.

Now, the decision to cross fill or not cross fill an inventoried item is a result of balancing these opposing inventory-holding costs. In addition, transportation costs may be included for shipping to a customer from distant stocking points. This cost does not encourage cross filling. However, if the stockout costs incurred due to demand not being served on request from the primary stocking location are included, cross filling is encouraged. Computing these costs for each item maintained in inventory identifies those items whose demand should be filled only from the primary stocking location and those that should be filled from the virtual inventory.

Example

Suppose that a firm has two options to serve its customers in order to maintain a high level of product availability. First, the customers can be served from a designated warehouse in their vicinity. If there is a stockout, the sale may be lost or a back order will occur. Second, when a stockout occurs, the order may be filled from a secondary warehouse, with the company paying the extra transportation costs. For any given item in inventory, which of these designs should be selected?

A representative item is selected from inventory to be tested. The distribution system is similar to that illustrated in Figure 9-22. The item has a cost in inventory of \$200 per unit, a carrying cost of 25 percent per year, a stock level of six weeks of demand, a replenishment lead time of eight weeks, and a target fill rate of 95 percent. The transport rate for cross hauling from a secondary warehouse outside of a customer's territory is \$10 per unit. The item's weekly demand characteristics are as follows:

Location	Mean Demand	Standard Deviation
1	300	138
2	100	80
System	400	160

The inventory-control policy is not known for sure.

The cross filling decision curves of Figure 9-23 help with this type of decision making. It is necessary to make some preliminary estimates. Since there is no inventory throughput curve of the type described in a previous section, we will assume that the company is operating a well-run inventory control system with an $\alpha = 0.7$. We know that the average inventory level can be described as $AIL = KD^{\alpha}$ and K can be found from a manipulation of this formula. That is, $K = D^{1-\alpha}/TO$. AIL can be approximated as D/TO, where TO is the turnover ratio or 52 weeks per year/6 weeks of demand in inventory = 8.67. Hence, $K = (400 \times 52)^{1-0.7}/8.67 = 2.28$. Also, we assume that the fill rate and the in-stock probability during the lead time are approximately the same so that z can be found from the normal distribution (Appendix A) for the fill rate percentage, or $z_{@0.95} = 1.96$.

The X, Y parameters for Figure 9-23 can now be computed.

$$X = \frac{tD^{1-\alpha}}{ICK} = \frac{10([400 \times 52])^{1-0.7})}{0.25(200)(2.28)} = 1.73$$

and

$$Y = \frac{zs\sqrt{LT}}{KD^{\alpha}} = \frac{1.96(160)\sqrt{8}}{2.28([400 \times 52])^{0.7}} = 0.4$$

The demand ratio r is $100/400 = 0.25$.

In Figure 9-23, we now see that the intersection of X and r falls below the decision curve of $Y = 0.4$ (use $Y = 0.5$). Since the value falls below the curve Y, do not cross fill the item.

Additional curves of the type shown in Figure 9-23 are available for various FR and α values.[23] They permit virtual inventory planning problems to be handled for a variety of items.

[23]Ibid

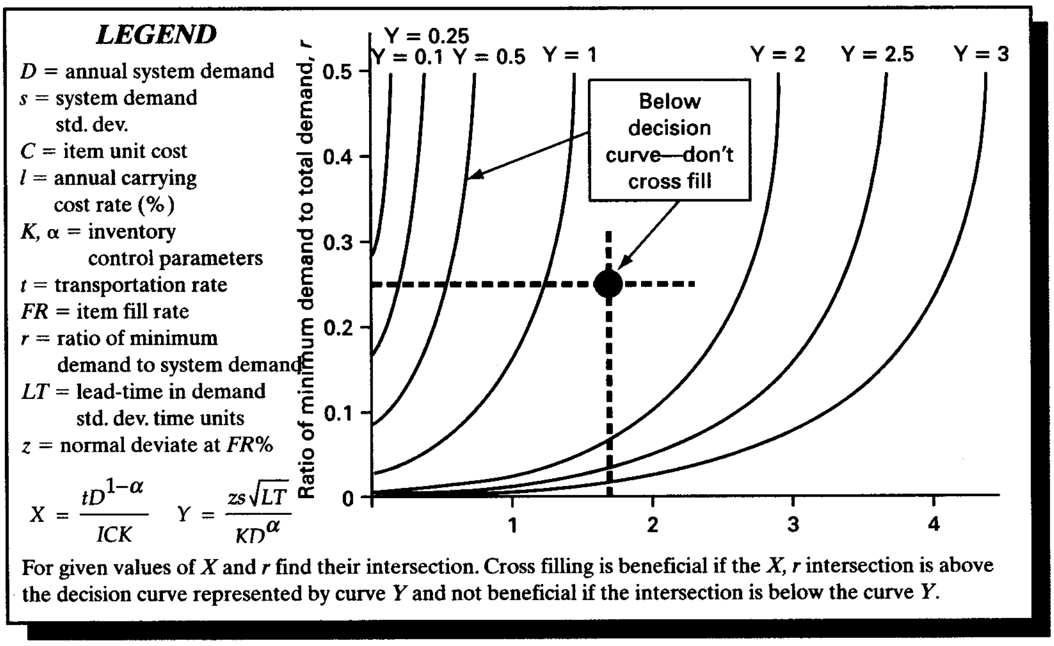

Figure 9-23 Cross filling Decision Curves for *FR* = 0.95 and α = 0.7.

Concluding Comments

Inventories continue to represent a major use of capital in the supply channel. Good management of them means keeping them at the lowest possible level consistent with a balance of direct and indirect costs attributed to their level and with the need to maintain a desired level of product availability. Extensive research has been conducted on optimally managing inventories, and this chapter summarizes the key inventory control methods that have proven to be useful in practice. The differences between pull, push, supply driven, and aggregate approaches to inventory planning and control have been noted. Specific mathematical methods are illustrated for a variety of circumstances, such as certainty and uncertainty of demand and lead time, perpetual and seasonal demand patterns, single and multiple echelons, single and multiple stock locations, and inventories at rest and in transit. They are all useful in establishing sound policies for managing an expensive asset.

Glossary of Terms

Q^* = optimal order quantity, units
Q = an order quantity, units
EOQ = economic order quantity, usually Q^*
CP_n = cumulative frequency of selling at least n units
D = average annual demand, units/year

S = item procurement cost, dollars/order
I = annual inventory carrying cost, percent of item value per year
C = item value in inventory, dollars/unit
T^* = optimal order interval expressed in units of time
T = a particular order interval expressed in units of time
N^* = optimal number of order intervals per year
ROP = reorder point quantity, units
d = average daily demand rate, units/day
LT = average lead time expressed in units of time
Q_p^* = optimum production run quantity, units
p = production output rate expressed in the same dimensions as d
s_d = standard deviation of demand, units
s_d' = standard demiation of the $DDLT$ or $DD[T + LT]$ demand distributions, units
X' = mean of the $DDLT$ or $DD[T + LT]$ demand distributions
s_{LT} = standard deviation of the lead time expressed in units of time
P = probability of being in stock during an order cycle (reorder point system) or during an order cycle plus lead time (periodic review system), expressed as a fraction or percent
z = normal deviate on the standardized normal distribution
SL = customer service level or item fill rate, expressed as a fraction or percent
AIL = average inventory level, units
$E_{(z)}$ = unit normal loss integral
TC = total relevant inventory cost, dollars/year
M^* = optimum max level for the periodic review system or min-max system
O = common order-processing cost for joint orders, $/order
k = stockout cost, dollars/unit
i = subscript to denote item number
Y = cumulative fraction of sales
X = cumulative fraction of items
A = a constant
n = number of stocking points in multiple stocking point system
L = inventory investment limit, dollars
K, α = inventory-throughput curve constants
ED = expected deficit, units

QUESTIONS

1. Why do inventories cost so much to maintain?
2. What are the reasons that inventories are held throughout a supply channel? Why should they be avoided?
3. Contrast a *push* inventory philosophy with a *pull* philosophy. When would each be most appropriately applied? Similarly, what is the difference between push and pull methods and aggregate control methods? When should each type be applied?
4. Identify the costs that are relevant to the control of inventories. Where do you think they might be obtained within a company?
5. Explain what safety stock is and why it is needed.
6. Explain the difference between the probability of a stockout during an order cycle and the service level, or item fill rate.
7. How might you decide which items in a company's product line are to be classified as *A*, *B*, and *C* items?
8. Explain what you think an executive meant by this statement: "Every management mistake ends up in inventory."

9. What is the square-root law in inventory planning, and to what problem types do you think it applies?
10. Why is the economic order quantity not very sensitive to inaccurate input data?
11. Where are pipeline inventories located in the supply channel, and how do we best control them?
12. How do you think one should go about setting the stock availability service level and determining stockout costs?
13. If the demand for an item in inventory showed the pattern of 0 1 2 5 150 0 1 0 3 4 150 1 0 0 5 1 150, what suggestions can you make on how the inventory level should be controlled?
14. Describe a supply-driven inventory system. How are inventory levels controlled compared with a pull system?
15. Contrast the stock to demand approach to inventory control with the *EOQ*-based period review method of control. Why is it simpler? Is there a price to be paid for this simplicity?
16. Explain the inventory pooling effect if the number of stocking points is varied.
17. What is the inventory throughput curve? How can it be determined? How is it useful?
18. What is a "virtual inventory"? What is the planning problem associated with such inventories?

PROBLEMS

A number of the following problems and case studies in this chapter can be solved or partially solved with the aid of computer software. The software packages in LOGWARE most important for this chapter are INPOL (*I*), and MULREG (*MR*). The CD icon will appear with the software package designation where the problem analysis is assisted by one of these software programs. A database may be prepared for the problem if extensive data input is required. Where the problem can be solved without the aid of the computer (by hand), the hand icon is shown. If no icon appears, hand calculation is assumed.

1. A shopper goes to a drug store in search of six items. The store stocks these items with the following probabilities of their being in stock:

Item	In-stock Probability, %
Toothpaste	95
Mouthwash	93
Batteries	87
Shaving cream	85
Aspirin tablets	94
Deodorant	90

Assuming that only one item of each of these products is purchased, what is the probability that the shopper will fill his or her order completely?

2. Central Hospital Supply has a policy that a hospital can expect its orders to be filled directly from stock 92 percent of the time. If any one item on an order is out of stock, the entire order will be held as a back order to avoid additional shipping costs. Orders typically include up to ten items. A sampling of the orders over the last year shows that six combinations of products frequently appear on orders, as follows:

Order Item	Combination	Frequency of Order
1	A,C,F,G,I	0.20
2	B,D,E	0.15
3	E,F,I,J	0.05
4	A,B,C,D,F,H,J	0.15
5	D,F,G,H,I,J	0.30
6	A,C,D,E,F	0.15
		1.00

Inventory levels have been set so that products A, B, C, D, E, and F have a common service level of 0.95 each. The remaining products are set at 0.90 each.

a. Is the firm meeting its inventory service target?

b. If not, at what item service levels would these two groups of products have to be set to meet the 92 percent order fill rate?

3. An importer of television sets from the Far East distributes them throughout the European Union from four warehouses. Shipments are received monthly, and this month's shipment is 120,000 sets. Due to the long lead time, demand and supply for the sets are hard to match. Therefore, an allocation to the warehouses is based on a monthly demand forecast and the service level for each warehouse. The inventory records and forecast for the coming month show the following situation:

Warehouse	Quantity on hand, sets	Demand forecast, sets[a]	Forecast error, sets[b]	Service level[c]
1	700	10,000	1,000	90%
2	0	15,000	1,200	85
3	2,500	35,000	2,000	88
4	1,800	25,000	3,000	92

[a] Projected to the time of a stock replenishment based on the current sales rate.
[b] A standard deviation. Forecast errors are normally distributed.
[c] Probability of being in stock during the month.

If transportation to the warehouses takes one week and import handling requires one week after a shipment arrives, how should the allocation of the sets be made to the warehouses?

4. A computer-supply mail-order house has a memory chip in inventory that it sells to customers around the country. A Japanese manufacturer supplies the item using airfreight. It has the following characteristics:

Average annual demand	= 3,200 units
Replenishment lead time	= 1.5 weeks
Carrying cost	= 15% per year
Purchase price, delivered	= $55 per unit
Procurement ordering cost	= $35 per order

a. Design a reorder point method of control for this item.
b. What are the annual ordering and annual carrying costs if your design is used?
c. Suppose that the lead time stretches to three weeks, so that the $ROP > Q^*$. What adjustments would you suggest making in the control policy?

5. Helen's Secretarial School trains young people in word processing and other secretarial skills. Tuition for the course is $8,500, but she will rebate up to 10 percent of the fee annually until the graduate receives a job. Average annual demand for her graduates is 300 per year. (Note: Product value and setup costs are the same.)
   a. How many potential secretaries should Helen admit per class?
   b. How many times per year should she offer the course?
6. A retail store purchases computer software from a distributor for resale. For an upcoming promotion, the retailer needs to determine the best order size for a one-time purchase. One of the products is a word processing software program that will have a special sale price of $350. The retailer estimates the probabilities for selling various quantities as follows:

Quantity	Probability
50	0.10
55	0.20
60	0.20
65	0.30
70	0.15
75	0.05
	1.00

   The program can be purchased from the distributor for $250 each, but there is a restocking charge of 20 percent of the purchase price for the return to the distributor of any unsold programs.

   What size of a purchase order should the retailer commit to?
7. An automatic teller machine (ATM) is being installed at a branch of MetroBank. From the bank's research, it figures to indirectly benefit from offering this service. Estimates are that the bank will generate revenues at the rate of 1 percent of the money passed through the machine in the form of new customer accounts for checking services, loans, savings accounts, and the like. The average withdrawal from the teller machine is $75, and the bank figures its cost of money to be 10 percent per year.

   Stocking the machine for the two-day weekend is its most difficult planning problem. From historical records for other ATMs, the bank estimates the average number of withdrawals to be 120 with a standard deviation of 20, with the distribution being normal.

How much money should the bank stock in the machine for the weekend? (*Hint*: Consider as a single order problem.)

8. Cabot Appliances, a retail chain, is trying to decide what size order it should place with its supplier of room air conditioners. Room air conditioner sales are highly seasonal, and the number of units sold is very dependent on summer weather patterns. Cabot places one order per year. Reorders are impractical after the selling season begins to develop. Although the actual sales level cannot be known for sure, Cabot analyzes past selling seasons, long-term weather reports, and the general state of the economy. The following probabilities of various sales levels are then estimated:

Sales, units	Probability
500	0.2
750	0.2
1,000	0.3
1,250	0.2
1,500	0.1
	1.0

A unit has a delivered price to Cabot of $320 and it is priced to customers at $400. Air conditioners unsold at the end of the season are discounted to $300, which clears them from inventory. Purchases can be made only in increments of 250 units, with a 500-unit minimum order.

a. Assuming no inventory is to be carried to the next year, what single order size should be placed?
b. Would you modify the order quantity in part a if Cabot can borrow money to support inventory at 20 percent per year? Excess units can be carried over to the next selling season.

9. Suppose that an auto part in a manufacturer's inventory has the following characteristics:

Forecast of demand	= 1,250 cases per week
Forecast error, std. dev.	= 475 cases per week
Lead time	= 2.5 weeks
Carrying cost	= 30% per year
Purchase price, delivered	= $56 per case
Replenishment order cost	= $40 per order
Stockout cost	= $10 per case
Probability of being in stock during the lead time, P	= 80%

a. Design a reorder point control system for this part, given the assigned P. How would you state the inventory control policy if the $ROP > Q^*$?

b. Design a periodic review system for this part. Now assume that the probability of being in stock extends to the order interval plus lead time.
c. Determine and compare the relevant costs of each approach.
d. What service level (fill rate) do you actually achieve with both designs?
e. Find the in-stock probability during the lead time that will optimize a reorder point system design. How does the total cost compare with that in part a?

10. Repeat Question 9, but include that the lead time is normally distributed with a standard deviation of 0.5 weeks.
11. A manufacturer of fractional horsepower motors for use in industrial sweepers and floor polishers produces its own wiring harnesses. These wiring harnesses are used in final assembly at the rate of 100 per day for 250 working days per year. It costs $250 to start up the wiring harness production. Production is at the rate of 300 per day when operating. The standard cost of a wiring harness is $75, and the company's inventory carrying cost is 25 percent per year.
    a. What should the production run quantity be?
    b. How long should each production run cycle be?
    c. How many times per year should the part be produced?
12. A Japanese appliance manufacturer uses a valve in its refrigerator final-assembly operation. The valve is obtained from a local supplier in any quantity needed within one hour of an order request. The working day is eight hours. The production schedule calls for this part to be used at the constant rate of 2,000 per day, 250 working days a year. The company pays ¥35 for this valve delivered to the assembly line. Inventory-carrying costs are 30 percent per year. Due to contractual arrangements with the supplier, procurement costs amount to only ¥1.00 per order placed.
    a. Design a reorder point method of inventory control for this item.
    b. Suggest how a two-bin system could be used as a way of implementing this control method.
13. A large chemical company in Green River, Wyoming, mines soda ash used in glass manufacturing. Soda ash is sold to a number of manufacturers through annual contracts. The glass companies release their requests for soda ash against their contracts. The mining company sees demand in the form of rail-car quantities. A typical week shows demand to be normally distributed at 40 rail-car loads plus or minus 10 cars. They estimate the standard deviation at (max cars – min cars)/6 = (50 – 30)/6 = 3.33 cars.

    Soda ash is valued at $30 per ton, and an average rail-car load is 90,000 lb of product. Annual carrying cost for the company is 25 percent per year. Setup costs at the mine are estimated to be $500 per order. It takes one week to produce the product and/or secure the rail-cars for shipment. A 90 percent in-stock probability during the lead time is desired.
    a. The company must call for cars from the railroad to fill orders. How many cars should be requested at a time? (*Remember:* One ton is 2,000 lb.)
    b. At what quantity of soda ash remaining in inventory should the request for cars be made?

14. A large hospital uses a certain intravenous solution that it maintains in inventory. Pertinent data about this item are as follows:

Forecasted daily usage[a]	= 50 units
Forecast error std. dev.[b]	= 15 units
Average lead time	= 7 days
Lead time std. dev.[b]	= 2 days
Annual carrying cost	= 30%
Procurement cost per order	= $50
Stockout cost	= $15 per unit
Product value	= $45 per unit
In-stock probability[c]	= 85%

[a]365 days per year
[b]Normally distributed
[c]During the lead time or during the order interval plus lead time, depending on the inventory control design

a. Design a reorder point system of control for this item.
b. Design a periodic review system of control for this item.
c. Do you think the in-stock probability is correctly specified to minimize costs? Appraise with reference to the reorder point system design.

15. A periodic review method of inventory control is to be used for two products that are to be purchased from the same supplier at the same time. The following data have been collected for these items:

	Products	
	A	B
Weekly demand forecast, units	2,000	500
Forecast error[a] (std. dev.), units	100	70
Lead time, weeks	1.5	1.5
Purchase price, dollar/unit	$2.25	$1.90
In-stock probability during lead time plus order cycle	90%	80%
Out-of-stock cost	Unknown	
Carrying cost, percent/year	30%	30%
Common purchase order cost, dollar/order	$100	

[a]Normally distributed

a. Design the control system for these products. State how the control system will work.
b. What is the average inventory level for each of these items?
c. What is the customer service level that can be expected for these items?
d. Suppose that the review time is set at four weeks. How will your answers to the previous questions change?

16. A company imports parts from Taiwan through the Port of Seattle on the West Coast. The parts are destined for its assembly operations on the East Coast. Shipments are by rail and require 21 days transit time. The parts are worth $250 each at the port, and 40,000 of them are used annually in assembly operations.

Inventory carrying costs are 25 percent per year. The rail rate to the East Coast is \$6 per 100 lb, and crated parts weigh 125 lb each.

As an alternative, trucking can be used to cross the country in seven days. Truck rates are \$11 per cwt. Do the savings from reduced in-transit inventories justify the higher cost of trucking?

17. At a point in Ohio, a manufacturer of hydraulic equipment (hoses, cylinders, and controls) consolidates the items on orders produced at various points in the United States. Consolidated orders are destined for Brazil and may be shipped via an ocean freight forwarder or using airfreight. The average order size is 292 lb. Ocean shipping (\$4.94/lb) is less expensive than airfreight (\$9.04/lb), but takes longer. Ocean movements from the consolidation center require transportation to the Port of Baltimore, vessel-waiting time for loading at the port, stop offs at Savannah and Miami for pickups, and sailing to São Paulo, Brazil. The total transit time averages 20 days. On the other hand, air shipments require only two days for handling and transit.

    The manufacturer owns the goods in transit until they arrive at the destination port and is concerned with the cost of inventory while in transit. The product in transit is valued at \$185/lb and 20,000 lb are shipped per year. The company's cost of capital is 17 percent per year.

    From purely an inventory-transportation viewpoint, which transport mode should be used?
18. A distributor of truck and bus parts has a tie-down strap (B2162H) in inventory. The item has a monthly demand of 169 units with a standard deviation of 327 units per month, making the demand pattern quite lumpy. The lead time for the item is four months with a standard deviation of 0.8 months. The item costs \$0.96 each at the factory with a \$0.048 transportation charge from supplier to distributor. Carrying costs are 20 percent per year, and order-processing costs are \$10 per order. The desired in-stock probability during the lead time is 85 percent. The inventory records are updated daily, and the average daily sales quantity is eight units.

    Develop a min-max (reorder point system) inventory control policy for this lumpy-demand item.
19. Acme Computer maintains a stock of spare parts for the entire country at one warehouse in Austin, Texas. To provide improved customer service, the company will expand the number of warehouses to ten, and they will be of equal size. The total inventory investment in the current warehouse is \$5,000,000.
    a. Using the square-root law, estimate the amount of inventory investment that the distribution system is likely to contain with ten warehouses.
    b. Suppose that nine warehouses are operating with \$1,000,000 of inventory investment in each. If the company were to consolidate the inventory into three equal-size warehouses, how much inventory would be in each of them?
20. The California Fruit Growers' Association is a consortium of fruit farmers on the West Coast for product distribution. The association currently operates 24 warehouses throughout the country. For the most recent calendar year, the statistics on average inventory levels and warehouse throughputs were compiled as given in Table 9-10.

WAREHOUSE	ANNUAL WAREHOUSE THROUGHPUT	AVERAGE INVENTORY LEVEL
1	$ 21,136,032	$ 2,217,790
2	16,174,988	2,196,364
3	78,559,012	9,510,027
4	17,102,486	2,085,246
5	88,226,672	11,443,489
6	40,884,400	5,293,539
7	43,105,917	6,542,079
8	47,136,632	5,722,640
9	24,745,328	2,641,138
10	57,789,509	6,403,076
11	16,483,970	1,991,016
12	26,368,290	2,719,330
13	$ 6,812,207	$1,241,921
14	28,368,270	3,473,799
15	28,356,369	4,166,288
16	48,697,015	5,449,058
17	47,412,142	5,412,573
18	25,832,337	3,599,421
19	75,266,622	7,523,846
20	6,403,349	1,009,402
21	2,586,217	504,355
22	44,503,623	2,580,183
23	22,617,380	3,001,390
24	4,230,491	796,669
Totals	$818,799,258	97,524,639

Table 9-10 California Fruit Growers' Association's Inventory Versus Throughput Statistics

a. What overall turnover ratio is the association able to achieve? Compare the turnover ratio for the three smallest warehouses with the three largest ones in terms of throughput. Suggest why there is a difference.
b. Construct the inventory throughput curve by fitting a straight line to the data by hand or use a simple, linear regression model.
c. Warehouses 1, 12, and 23 are to be consolidated into one warehouse. How much inventory would you expect in the one warehouse using the curve from part b?
d. Warehouse 5 is to be expanded into two warehouses. Thirty percent of the throughput will be assigned to one warehouse, and the remainder to the other. How much inventory would you estimate to be in each warehouse using the curve from part b?