CHAPTER

12

Managing Uncertainty in a Supply Chain

Safety Inventory

LEARNING OBJECTIVES

After reading this chapter, you will be able to

1. Describe different measures of product availability.
2. Understand the role of safety inventory in a supply chain.
3. Identify factors that influence the required level of safety inventory.
4. Use available managerial levers to lower safety inventory without hurting product availability.

In this chapter, we discuss how safety inventory can help a supply chain improve product availability in the presence of supply and demand variability. We discuss various measures of product availability and how managers can set safety inventory levels to provide the desired product availability. We also explore what managers can do to reduce the amount of safety inventory required while maintaining or even improving product availability.

12.1 THE ROLE OF SAFETY INVENTORY IN A SUPPLY CHAIN

Safety inventory is inventory carried to satisfy demand that exceeds the amount forecast. Safety inventory is required because demand is uncertain, and a product shortage may result if actual demand exceeds the forecast demand. Consider, for example, Bloomingdale's, a high-end department store. Bloomingdale's sells purses purchased from Gucci, an Italian manufacturer. Given the high transportation cost from Italy, the store manager at Bloomingdale's orders in lots of 600 purses. Demand for purses at Bloomingdale's averages 100 a week. Gucci takes three weeks to deliver the purses to Bloomingdale's in response to an order. If there is no demand uncertainty and exactly 100 purses are sold each week, the store manager at Bloomingdale's can place an order when the store has exactly 300 purses remaining. In the absence of demand uncertainty, such a policy ensures that the new lot arrives just as the last purse is being sold at the store.

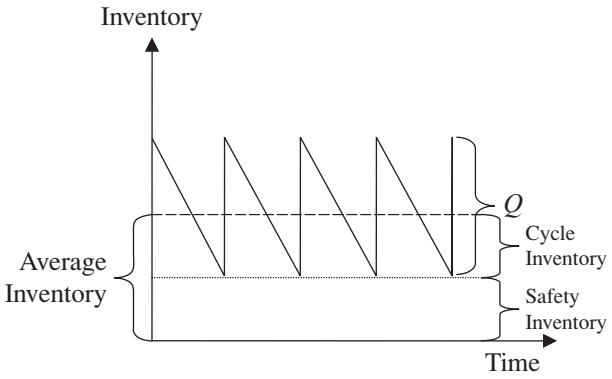

FIGURE 12-1 Inventory Profile with Safety Inventory

However, given demand fluctuations and forecast errors, actual demand over the three weeks may be higher or lower than the 300 purses that were forecast. If the actual demand at Bloomingdale's is higher than 300, some customers will be unable to purchase purses, resulting in a potential loss of margin for Bloomingdale's. The store manager thus decides to place an order with Gucci when the store still has 400 purses. This policy improves product availability for the customer because the store now runs out of purses only if the demand over the three weeks exceeds 400. Given an average weekly demand of 100 purses, the store will have an average of 100 purses remaining when the replenishment lot arrives. Safety inventory is the average inventory remaining when the replenishment lot arrives. Thus, Bloomingdale's carries a safety inventory of 100 purses.

Given a lot size of $Q = 600$ purses, the cycle inventory, the focus of the previous chapter, is $Q/2 = 300$ purses. The inventory profile at Bloomingdale's in the presence of safety inventory is shown in Figure 12-1, which illustrates that the average inventory at Bloomingdale's is the sum of the cycle and safety inventories.

This example illustrates a trade-off that a supply chain manager must consider when planning safety inventory. On one hand, raising the level of safety inventory increases product availability, and thus the margin captured from customer purchases. On the other hand, raising the level of safety inventory increases inventory holding costs. This issue is particularly significant in industries in which product life cycles are short and demand is volatile. Carrying excessive inventory can help counter demand volatility but can really hurt if new products come onto the market and demand for the product in inventory dries up. The inventory on hand then becomes worthless.

In today's business environment, it has become easier for customers to search across stores for product availability. If Amazon is out of a book, for example, a customer can easily check to see whether barnesandnoble.com has the title available. The increased ease of searching puts pressure on firms to improve product availability. Simultaneously, product variety has grown with increased customization. As a result, markets have become increasingly heterogeneous and demand for individual products is unstable and difficult to forecast. Both the increased variety and the greater pressure for availability push firms to raise the level of safety inventory they hold. Given the product variety and high demand uncertainty in most high-tech supply chains, a significant fraction of the inventory carried is safety inventory.

As product variety has grown, however, product life cycles have shrunk. Thus, it is more likely that a product that is "hot" today will be obsolete tomorrow, which increases the cost to firms of carrying too much inventory. Thus, a key to the success of any supply chain is to figure out ways to decrease the level of safety inventory carried without hurting the level of product availability.

The importance of reduced safety inventories is emphasized by the experience of Nordstrom, Macy's, and Saks during the 2008–2009 recession. Nordstrom outperformed the other

two chains by moving its inventories about twice as fast as its competitors. In 2008 (2009), Nordstrom carried an average of about 2 (2) months, Macy's carried about 4 (4.15) months, and Saks carried about 4.24 (4.67) months of inventory. A key to Nordstrom's success has been its ability to provide a high level of product availability to customers while carrying low levels of safety inventory in its supply chain. This fact has also played an important role in the success of Zara, Walmart, and Seven-Eleven Japan.

For any supply chain, three key questions need to be considered when planning safety inventory:

1. What is the appropriate level of product availability?
2. How much safety inventory is needed for the desired level of product availability?
3. What actions can be taken to reduce safety inventory without hurting product availability?

The first question is discussed in detail in Chapter 13. The remainder of this chapter focuses on answering the second and third questions, assuming a desired level of product availability. Next, we consider factors that influence the appropriate level of safety inventory.

12.2 FACTORS AFFECTING THE LEVEL OF SAFETY INVENTORY

The appropriate level of safety inventory is determined by the following two factors:

- The uncertainty of both demand and supply
- The desired level of product availability

As the uncertainty of supply or demand grows, the required level of safety inventories increases. Demand for milk at a supermarket is quite predictable. As a result, supermarkets can operate with low levels of safety inventory relative to demand. In contrast, demand for spices at the same supermarket is much harder to predict. Thus the supermarket needs to carry high levels of safety inventory for spices relative to demand. Whereas most of the milk inventory at a supermarket is cycle inventory (with very little being safety inventory), most of the spice inventory is safety inventory carried to deal with uncertainty of demand.

As the desired level of product availability increases, the required level of safety inventory also increases. If the supermarket targets a higher level of product availability for a certain spice, it must carry a higher level of safety inventory for that spice.

Next, we discuss some measures of demand uncertainty.

Measuring Demand Uncertainty

As discussed in Chapter 7, demand has a systematic as well as a random component. The random component is a measure of demand uncertainty. The goal of forecasting is to predict the systematic component and estimate the random component. The random component is usually estimated as the standard deviation of forecast error. We illustrate our ideas using uncertain demand for a smartphone at B&M Office Supplies as the context. We assume that periodic demand for the phone at B&M is normally distributed with the following inputs:

D: Average demand per period

σ_D: Standard deviation of demand (forecast error) per period

Even though standard deviation of demand is not necessarily the same as forecast error, we treat the two to be interchangeable in our discussion. Safety inventory calculations should really be based on forecast error.

Lead time is the gap between the time an order is placed and when it is received. In our discussion, we denote the lead time by L. In the B&M example, L is the time between when B&M orders phones and when they are delivered. In this case, B&M is exposed to the uncertainty of demand during the lead time. Whether B&M is able to satisfy all demand from inventory depends

on the demand for phones experienced during the lead time and the inventory B&M has when a replenishment order is placed. Thus, B&M must estimate the uncertainty of demand during the lead time, not just in a single period. We now evaluate the distribution of demand over L periods, given the distribution of demand during each period.

EVALUATING DEMAND DISTRIBUTION OVER *L* PERIODS Assume that demand for each period i, $i = 1, \ldots, L$, is normally distributed with a mean D_i and standard deviation σ_i. Let ρ_{ij} be the correlation coefficient of demand between periods i and j. In this case, the total demand during L periods is normally distributed with a mean of D_L and a standard deviation of σ_L, where the following is true:

$$D_L = \sum_{i=1}^{L} D_i, \quad \sigma_L = \sqrt{\sum_{i=1}^{L} \sigma_i^2 + 2\sum_{i>j} \rho_{ij}\sigma_i\sigma_j} \tag{12.1}$$

Demand in two periods is *perfectly positively correlated* if $\rho_{ij} = 1$. Demand in two periods is *perfectly negatively correlated* if $\rho_{ij} = -1$. Demand in two periods is *independent* if $\rho_{ij} = 0$. If demand during each of L periods is independent and normally distributed with a mean of D and a standard deviation of σ_D, Equation 12.1 can be used to show that total demand during the L periods is normally distributed with a mean D_L and a standard deviation of σ_L, where the following is true:

$$D_L = D \times L, \quad \sigma_L = \sqrt{L}\sigma_D \tag{12.2}$$

Another important measure of uncertainty is the *coefficient of variation* (cv), which is the ratio of the standard deviation to the mean. Given demand with a mean of μ and a standard deviation of σ, we have

$$cv = \sigma / \mu$$

The coefficient of variation measures the size of the uncertainty relative to demand. It captures the fact that a product with a mean demand of 100 and a standard deviation of 100 has greater demand uncertainty than a product with a mean demand of 1,000 and a standard deviation of 100. Considering the standard deviation alone cannot capture this difference.

Next, we discuss some measures of product availability.

Measuring Product Availability

Product availability reflects a firm's ability to fill a customer order out of available inventory. A *stockout* results if a customer order arrives when product is not available. There are several ways to measure product availability. Some of the important measures are listed next.

1. *Product fill rate (fr)* is the fraction of product demand that is satisfied from product in inventory. Fill rate is equivalent to the probability that product demand is supplied from available inventory. Fill rate should be measured over specified amounts of demand rather than over time. Thus, it is more appropriate to measure fill rate over every million units of demand rather than every month. Assume that B&M provides smartphones to 90 percent of its customers from inventory, with the remaining 10 percent lost to a neighboring competitor because of a lack of available inventory. In this case, B&M achieves a fill rate of 90 percent.

2. *Order fill rate* is the fraction of orders that are filled from available inventory. Order fill rate should also be measured over a specified number of orders rather than over time. In a multiproduct scenario, an order is filled from inventory only if all products in the order can be supplied from the available inventory. In the case of B&M, a customer may order a phone along with a laptop. The order is filled from inventory only if both the phone and the laptop are available through the store. Order fill rates tend to be lower than product fill rates because all products must be in stock for an order to be filled.

3. ***Cycle service level (CSL)*** is the fraction of replenishment cycles that end with all the customer demand being met. A *replenishment cycle* is the interval between two successive replenishment deliveries. The CSL is equal to the probability of not having a stockout in a replenishment cycle. CSL should be measured over a specified number of replenishment cycles. If B&M orders replenishment lots of 600 phones, the interval between the arrival of two successive replenishment lots is a replenishment cycle. If the manager at B&M manages inventory such that the store does not run out of inventory in 6 out of 10 replenishment cycles, the store achieves a CSL of 0.6 or 60 percent. Observe that a CSL of 0.6 typically results in a much higher fill rate. In the 60 percent of cycles in which B&M does not run out of inventory, all customer demand is satisfied from available inventory. In the 40 percent of cycles in which a stockout does occur, most of the customer demand is satisfied from inventory. Only the small fraction toward the end of the cycle that arrives after B&M is out of inventory is lost. As a result, the fill rate is much higher than 0.6.

The distinction between product fill rate and order fill rate is usually not significant in a single-product situation. When a firm is selling multiple products, however, this difference may be significant. For example, if most orders include 10 or more products that are to be shipped, an out-of-stock situation of one product results in the order not being filled from stock. The firm in this case may have a poor order fill rate even though it has good product fill rates. Tracking order fill rates is important when customers place a high value on the entire order being filled at one time.

Next, we describe two replenishment policies that are often used in practice.

Replenishment Policies

A replenishment policy consists of decisions regarding when to reorder and how much to reorder. These decisions determine the cycle and safety inventories along with the fill rate *fr* and the cycle service level CSL. Replenishment policies may take any of several forms. We restrict attention to two types:

1. ***Continuous review:*** Inventory is continuously tracked, and an order for a lot size Q is placed when the inventory declines to the reorder point (ROP). As an example, consider the store manager at B&M who continuously tracks the inventory of phones. She orders 600 phones when the inventory drops below ROP $=$ 400. In this case, the size of the order does not change from one order to the next. The time between orders may fluctuate, given variable demand.

2. ***Periodic review:*** Inventory status is checked at regular periodic intervals, and an order is placed to raise the inventory level to a specified threshold. As an example, consider the purchase of flash drives at B&M. The store manager does not track flash drive inventory continuously. Every Thursday, employees check flash drive inventory, and the manager orders enough so that the total of the available inventory and the size of the order equals 1,000 flash drives. In this case, the time between orders is fixed. The size of each order, however, can fluctuate given variable demand.

These inventory policies are not comprehensive, but they suffice to illustrate the key managerial issues concerning safety inventories.

12.3 DETERMINING THE APPROPRIATE LEVEL OF SAFETY INVENTORY

We now discuss the relationship between safety inventory and the CSL and *fr*. In this section, we restrict our attention to the continuous review policy. The periodic review policy is discussed in detail in Section 12.6. The continuous review policy consists of a lot size Q ordered when the inventory on hand declines to the ROP. Assume that weekly demand is normally distributed, with mean D and standard deviation σ_D. Assume replenishment lead time of L weeks.

Linking Safety Inventory and Cycle Service Level

We first show how cycle service levels can be evaluated given a replenishment policy (and thus the corresponding safety inventory). We then show how to determine the required safety inventory given a desired cycle service level.

EVALUATING SAFETY INVENTORY GIVEN A REPLENISHMENT POLICY In the case of B&M, safety inventory corresponds to the average number of phones on hand when a replenishment order arrives. Given the lead time of L weeks and a mean weekly demand of D, using Equation 12.2, we have

$$\text{Expected demand during lead time} = D \times L$$

Given that the store manager places a replenishment order when ROP phones are on hand, we have

$$\text{Safety inventory, } ss = ROP - D \times L \quad \textbf{(12.3)}$$

This is because, on average, $D \times L$ phones will sell over the L weeks between when the order is placed and when the lot arrives. The average safety inventory when the replenishment lot arrives is thus $ROP - D \times L$. The evaluation of safety inventory for a given inventory policy is described in Example 12-1 (see spreadsheet *Chapter 12-examples* worksheet *Example 12-1*).

EXAMPLE 12-1 Evaluating Safety Inventory Given an Inventory Policy

Assume that weekly demand for phones at B&M Office Supplies is normally distributed, with a mean of 2,500 and a standard deviation of 500. The manufacturer takes two weeks to fill an order placed by the B&M manager. The store manager currently orders 10,000 phones when the inventory on hand drops to 6,000. Evaluate the safety inventory and the average inventory carried by B&M. Also evaluate the average time a phone spends at B&M.

Analysis:

Under this replenishment policy, we have

Average demand per week, $D = 2{,}500$
Standard deviation of weekly demand, $\sigma_D = 500$
Average lead time for replenishment, $L = 2$ weeks
Reorder point, $ROP = 6{,}000$
Average lot size, $Q = 10{,}000$

Using Equation 12.3, we thus have

$$\text{Safety inventory, } ss = ROP - D \times L = 6{,}000 - 5{,}000 = 1{,}000$$

B&M thus carries a safety inventory of 1,000 phones. From Chapter 11, recall that

$$\text{Cycle inventory} = Q/2 = 10{,}000/2 = 5{,}000$$

We thus have

$$\text{Average inventory} = \text{cycle inventory} + \text{safety inventory} = 5{,}000 + 1{,}000 = 6{,}000$$

B&M thus carries an average of 6,000 phones in inventory. Using Little's law (Equation 3.1), we have

$$\text{Average flow time} = \text{average inventory/throughput} = 6{,}000/2{,}500 = 2.4 \text{ weeks}$$

Each phone thus spends an average of 2.4 weeks at B&M.

Next, we discuss how to evaluate the CSL given a replenishment policy.

EVALUATING CYCLE SERVICE LEVEL GIVEN A REPLENISHMENT POLICY Given a replenishment policy, our goal is to evaluate the CSL, the probability of not stocking out in a replenishment cycle. We return to B&M's continuous review replenishment policy of ordering Q units when the inventory on hand drops to the ROP. The lead time is L weeks and weekly demand is normally distributed, with a mean of D and a standard deviation of σ_D. Observe that a stockout occurs in a cycle if demand during the lead time is larger than the ROP. Thus, we have

$$CSL = \text{Prob}(\text{demand during lead time of } L \text{ weeks} \leq ROP)$$

To evaluate this probability, we need to obtain the distribution of demand during the lead time. From Equation 12.2, we know that demand during lead time is normally distributed, with a mean of D_L and a standard deviation of σ_L. Using the notation for the normal distribution from Appendix 12A and the equivalent Excel function from Equation 12.22 in Appendix 12B, the CSL is

$$CSL = F(ROP, D_L, \sigma_L) = NORMDIST(ROP, D_L, \sigma_L, 1) \quad \textbf{(12.4)}$$

We now illustrate this evaluation in Example 12-2 (see worksheet *Example 12-2*).

EXAMPLE 12-2 Evaluating Cycle Service Level Given a Replenishment Policy

Weekly demand for phones at B&M is normally distributed, with a mean of 2,500 and a standard deviation of 500. The replenishment lead time is two weeks. Assume that the demand is independent from one week to the next. Evaluate the CSL resulting from a policy of ordering 10,000 phones when there are 6,000 phones in inventory.

Analysis:

In this case, we have

$$Q = 10{,}000,\ ROP = 6{,}000,\ L = 2 \text{ weeks}$$
$$D = 2{,}500/\text{week},\ \sigma_D = 500$$

Observe that B&M runs the risk of stocking out during the lead time of two weeks between when an order is placed and when the replenishment arrives. Thus, whether or not a stockout occurs depends on the demand during the lead time of two weeks.

Because demand across time is independent, we use Equation 12.2 to obtain demand during the lead time to be normally distributed with a mean of D_L and a standard deviation of σ_L, where

$$D_L = D \times L = 2 \times 2{,}500 = 5{,}000, \quad \sigma_L = \sqrt{L}\sigma_D = \sqrt{2} \times 500 = 707$$

Using Equation 12.4, the CSL is evaluated as

$$CSL = F(ROP, D_L, \sigma_L) = NORMDIST(ROP, D_L, \sigma_L, 1)$$
$$= NORMDIST(6000, 5000, 707, 1) = 0.92$$

A CSL of 0.92 implies that in 92 percent of the replenishment cycles, B&M supplies all demand from available inventory. In the remaining 8 percent of the cycles, stockouts occur and some demand is not satisfied because of the lack of inventory.

We now discuss how the appropriate level of safety inventory may be obtained given a desired CSL.

Evaluating Safety Inventory Given Desired Cycle Service Level

In many practical settings, firms have a desired level of product availability and want to design replenishment polices that achieve this level. For example, Walmart has a desired level of product availability for each product sold in a store. The store manager must design a replenishment policy with the appropriate level of safety inventory to meet this goal. The desired level of product availability may be determined by trading off the cost of holding inventory with the cost of a stockout. This trade-off is discussed in detail in Chapter 13. In other instances, the desired level of product availability (in terms of CSL or fill rate) is stated explicitly in contracts, and management must design replenishment policies that achieve the desired target.

EVALUATING REQUIRED SAFETY INVENTORY GIVEN DESIRED CYCLE SERVICE LEVEL Our goal is to obtain the appropriate level of safety inventory given the desired CSL. We assume that a continuous review replenishment policy is followed. Consider the store manager at Walmart responsible for designing replenishment policies for all products in the store. He has targeted a CSL for the basic box of Lego building blocks. Given a lead time of L, the store manager wants to identify a suitable reorder point ROP and safety inventory that achieves the desired service level. Assume that demand for Legos at Walmart is normally distributed and independent from one week to the next. We assume the following inputs:

Desired cycle service level $= CSL$

Mean demand during lead time $= D_L$

Standard deviation of demand during lead time $= \sigma_L$

From Equation 12.3, recall that $ROP = D_L + ss$. The store manager needs to identify safety inventory ss such that the following is true:

$$\text{Probability}(\text{demand during lead time} \leq D_L + ss) = CSL$$

Given that demand is normally distributed, the store manager must identify safety inventory ss such that the following is true (using Equation 12.4):

$$F(D_L + ss, D_L, \sigma_L) = CSL$$

Given the definition of the inverse normal in Appendix 12A and the equivalent Excel function from Appendix 12B, we obtain

$$D_L + ss = F^{-1}(CSL, D_L, \sigma_L) = NORMINV(CSL, D_L, \sigma_L)$$
$$\text{or } ss = F^{-1}(CSL, D_L, \sigma_L) - D_L = NORMINV(CSL, D_L, \sigma_L) - D_L$$

Using the definition of the standard normal distribution and its inverse from Appendix 12A, and the equivalent Excel function from Appendix 12B, it can also be shown that the following is true:

$$ss = F_S^{-1}(CSL) \times \sigma_L = F_S^{-1}(CSL) \times \sqrt{L}\sigma_D = NORMSINV(CSL) \times \sqrt{L}\sigma_D \quad \textbf{(12.5)}$$

In Example 12-3 (see worksheet *Example 12-3*), we illustrate the evaluation of safety inventory given a desired CSL.

EXAMPLE 12-3 Evaluating Safety Inventory Given a Desired Cycle Service Level

Weekly demand for Legos at a Walmart store is normally distributed, with a mean of 2,500 boxes and a standard deviation of 500. The replenishment lead time is two weeks. Assuming a continuous-review replenishment policy, evaluate the safety inventory that the store should carry to achieve a CSL of 90 percent.

Analysis:

In this case we have

$$D = 2{,}500/\text{week},\ \sigma_D = 500,\ CSL = 0.9,\ L = 2 \text{ weeks}$$

Because demand across time is independent, we use Equation 12.2 to find demand during the lead time to be normally distributed with a mean of D_L and a standard deviation of σ_L, where

$$D_L = D \times L = 2 \times 2{,}500 = 5{,}000;\ \ \sigma_L = \sqrt{L}\sigma_D = \sqrt{2} \times 500 = 707$$

Using Equation 12.5, we obtain

$$ss = F_s^{-1}(CSL) \times \sigma_L = NORMSINV(CSL) \times \sigma_L = NORMSINV(0.90) \times 707 = 906$$

Thus, the required safety inventory to achieve a CSL of 90 percent is 906 boxes.

Linking Safety Inventory and Fill Rate

We now show how fill rates can be evaluated given a replenishment policy (and thus the corresponding safety inventory). We then show how to determine the required safety inventory given a desired fill rate.

EVALUATING FILL RATE GIVEN A REPLENISHMENT POLICY Recall that fill rate measures the proportion of customer demand that is satisfied from available inventory. Fill rate is generally a more relevant measure than cycle service level because it allows the retailer to estimate the fraction of demand that is turned into sales. The two measures are closely related, as raising the cycle service level also raises the fill rate for a firm. Our discussion focuses on evaluating fill rate for a continuous review policy under which Q units are ordered when the quantity on hand drops to the ROP.

To evaluate the fill rate, it is important to understand the process by which a stockout occurs during a replenishment cycle. A stockout occurs if the demand during the lead time exceeds the ROP. We thus need to evaluate the average amount of demand in excess of the ROP in each replenishment cycle.

The *expected shortage per replenishment cycle* (ESC) is the average units of demand that are not satisfied from inventory in stock per replenishment cycle. Given a lot size of Q (which is also the average demand in a replenishment cycle), the fraction of demand lost is thus ESC/Q.

The product fill rate fr is thus given by

$$fr = 1 - ESC/Q = (Q - ESC)Q \tag{12.6}$$

A shortage occurs in a replenishment cycle only if the demand during the lead time exceeds the ROP. Let $f(x)$ be the density function of the demand distribution during the lead time. The ESC is given by

$$ESC = \int_{x=ROP}^{\infty} (x - ROP) f(x)\,dx \tag{12.7}$$

When demand during the lead time is normally distributed with mean D_L and standard deviation σ_L, given a safety inventory ss, Equation 12.7 can be simplified to

$$ESC = -ss\left[1 - F_s\left(\frac{ss}{\sigma_L}\right)\right] + \sigma_L f_s\left(\frac{ss}{\sigma_L}\right) \tag{12.8}$$

where F_s is the standard normal cumulative distribution function and f_s is the standard normal density function. The standard normal distribution has a mean of 0 and a standard deviation of 1. A detailed description of the normal distribution is given in Appendix 12A. Details of the simplification in Equation 12.8 are described in Appendix 12C. Using Excel functions

	A	B	C	D	E
1	*Inputs*				
2	Q	D	σ_D	L	ss
3	10,000	2,500	500	2	1,000
4	*Distribution of demand during lead time*				
5	D_L	σ_L			
6	5,000	707			
7	*Cycle Service Level and Fill Rate*				
8	*CSL*	*ESC*	*fr*		
9	0.92	25.13	0.9975		

Cell	Cell Formula	Equation
A6	=B3*D3	12.2
B6	=SQRT(D3)*C3	12.2
A9	=NORMDIST(A6+E3, A6, B6, 1)	12.4
B9	=-E3*(1-NORMDIST(E3/B6, 0, 1, 1)) + B6*NORMDIST(E3/B6, 0, 1, 0)	12.8
C9	=(A3-B9)/A3	12.5

FIGURE 12-2 Excel Solution of Example 12-4

(Equations 12.25 and 12.26) discussed in Appendix 12B, ESC may be evaluated (using Equation 12.8) as

$$ESC = -ss[1 - NORMDIST(ss/\sigma_L, 0, 1, 1)] + \sigma_L NORMDIST(ss/\sigma_L, 0, 1, 0) \quad \textbf{(12.9)}$$

Given the ESC, we can use Equation 12.6 to evaluate the fill rate *fr*. Next, we illustrate this evaluation in Example 12-4 (see worksheet *Example 12-4* and Figure 12-2).

EXAMPLE 12-4 Evaluating Fill Rate Given a Replenishment Policy

From Example 12-2, recall that weekly demand for phones at B&M is normally distributed, with a mean of 2,500 and a standard deviation of 500. The replenishment lead time is two weeks. Assume that the demand is independent from one week to the next. Evaluate the fill rate resulting from the policy of ordering 10,000 phones when there are 6,000 phones in inventory.

Analysis:

From the analysis of Example 12-2, we have

Lot size, $Q = 10{,}000$

Average demand during lead time, $D_L = 5{,}000$

Standard deviation of demand during lead time, $\sigma_L = 707$

Using Equation 12.3, we obtain

$$\text{Safety inventory, } ss = ROP - D_L = 6{,}000 - 5{,}000 = 1{,}000$$

From Equation 12.9, we thus have

$$ESC = -1{,}000[1 - NORMDIST(1{,}000/707, 0, 1, 1)] + 707\, NORMDIST(1{,}000/707, 0, 1, 0) = 25$$

Thus, on average, in each replenishment cycle, 25 phones are demanded by customers but not available in inventory. Using Equation 12.6, we thus obtain the following fill rate:

$$fr = (Q - ESC)/Q = (10{,}000 - 25)/10{,}000 = 0.9975$$

In other words, 99.75 percent of the demand is filled from inventory in stock. This is much higher than the CSL of 92 percent that resulted in Example 12-2 for the same replenishment policy.

A few key observations should be made. First, observe that the fill rate (0.9975) in Example 12-4 is significantly higher than the CSL (0.92) in Example 12-2 for the same replenishment policy. Next, by rerunning the examples with a different lot size (in worksheet *Example 12-4*), we can observe the impact of lot-size changes on the service level. Increasing the lot size of phones from 10,000 to 20,000 has no impact on the CSL (which stays at 0.92). The fill rate, however, now increases to 0.9987. This happens because an increase in lot size results in fewer replenishment cycles. In the case of B&M, an increase in lot size from 10,000 to 20,000 results in replenishment occurring once every eight weeks instead of once every four weeks. With a 92 percent CSL, a lot size of 10,000 results in, on average, one cycle with a stockout per year. With a lot size of 20,000, we have, on average, one stockout every two years. Thus, the fill rate is higher.

Key Point

Both fill rate and cycle service level increase as the safety inventory is increased. For the same safety inventory, an increase in lot size increases the fill rate but not the cycle service level.

EVALUATING REQUIRED SAFETY INVENTORY GIVEN DESIRED FILL RATE For a continuous review replenishment policy, we now evaluate the required safety inventory given a desired fill rate *fr*. Consider the store manager at Walmart targeting a fill rate *fr* for Lego building blocks. The current replenishment lot size is *Q*. The first step is to obtain the ESC using Equation 12.6.

The next step is to obtain a safety inventory *ss* that solves Equation 12.8 (and its Excel equivalent, Equation 12.9) given the ESC evaluated earlier. It is not possible to give a formula that provides the answer. The appropriate safety inventory that solves Equation 12.9 can be obtained easily using Excel and trying different values of *ss*. In Excel, the safety inventory may also be obtained directly using the tool *GOALSEEK,* as illustrated in Example 12-5 (use worksheet *Example 12-5*).

EXAMPLE 12-5 Evaluating Safety Inventory Given Desired Fill Rate

Weekly demand for Legos at a Walmart store is normally distributed, with a mean of 2,500 boxes and a standard deviation of 500. The replenishment lead time is two weeks. The store manager currently orders replenishment lots of 10,000 boxes from Lego. Assuming a continuous-review replenishment policy, evaluate the safety inventory the store should carry to achieve a fill rate of 97.5 percent.

Analysis:

In this case, we have

Desired fill rate, $fr = 0.975$

Lot size, $Q = 10{,}000$ boxes

Standard deviation of demand during lead time, $\sigma_L = \sqrt{2} \times 500 = 707$

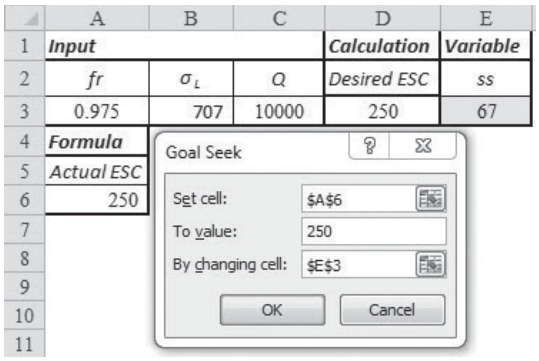

Cell	Cell Formula	Equation
A6	-E3*(1-NORMSDIST(E3/B3, 0, 1,1)) + B3*NORMDIST(E3/B3, 0, 1, 0)	12.10

FIGURE 12-3 Spreadsheet to Solve for *ss* Using *GOALSEEK*

From Equation 12.6, we thus obtain an ESC as

$$ESC = (1 - fr)\,Q = (1 - 0.975)10{,}000 = 250$$

Now we need to solve Equation 12.8 for the safety inventory *ss*, where

$$ESC = 250 = -ss\left[1 - F_s\left(\frac{ss}{\sigma_L}\right)\right] + \sigma_L f_s\left(\frac{ss}{\sigma_L}\right) = -ss\left[1 - F_s\left(\frac{ss}{707}\right)\right] + 707 f_s\left(\frac{ss}{707}\right)$$

Using Equation 12.9, this equation may be restated with Excel functions as follows:

$$250 = -ss[1 - NORMDIST(ss/707,0,1,1)] + 707NORMDIST(ss/707,0,1,0) \quad \textbf{(12.10)}$$

Equation 12.10 may be solved in Excel by trying different values of *ss* until the equation is satisfied. A more elegant approach for solving Equation 12.10 is to use the Excel tool *GOALSEEK*, as follows.

In the worksheet *Example 12-5,* invoke *GOALSEEK* using Data | What-If Analysis | Goal Seek. In the *GOALSEEK* dialog box, enter the data as shown in Figure 12-3 and click the OK button. In this case, cell D3 is changed until the value of the formula in cell A6 equals 250.

Using *GOALSEEK*, we obtain a safety inventory of $ss = 67$ boxes, as shown in Figure 12-3. Thus, the store manager at Walmart should target a safety inventory of 67 boxes to achieve the desired fill rate of 97.5 percent.

Impact of Desired Product Availability and Uncertainty on Safety Inventory

The two key factors that affect the required level of safety inventory are the desired level of product availability and uncertainty. We now discuss the impact that each factor has on the safety inventory.

As the desired product availability goes up, the required safety inventory also increases because the supply chain must now be able to accommodate uncommonly high demand or uncommonly low supply. For the Walmart situation in Example 12-5, we evaluate the required safety inventory for varying levels of fill rate as shown in Table 12-1.

Observe that raising the fill rate from 97.5 percent to 98.0 percent requires an additional 116 units of safety inventory, whereas raising the fill rate from 99.0 percent to 99.5 percent requires an additional 268 units of safety inventory. Thus, the marginal increase in safety inventory grows as product availability rises. This phenomenon highlights the importance of selecting

TABLE 12-1 Required Safety Inventory for Different Values of Fill Rate

Fill Rate	Safety Inventory
97.5%	67
98.0%	183
98.5%	321
99.0%	499
99.5%	767

suitable product availability levels. It is important for a supply chain manager to be aware of the products that require a high level of availability and hold high safety inventories only for those products. It is not appropriate to select a high level of product availability and require it arbitrarily for all products.

Key Point

The required safety inventory grows rapidly with an increase in the desired product availability.

From Equation 12.5, we see that the required safety inventory *ss* is also influenced by the standard deviation of demand during the lead time, σ_L. The standard deviation of demand during the lead time is influenced by the duration of the lead time L and the standard deviation of periodic demand σ_D, as shown in Equation 12.2. The relationship between safety inventory and σ_D is linear, in that a 10 percent increase in σ_D results in a 10 percent increase in safety inventory. Safety inventory also increases with an increase in lead time L. The safety inventory, however, is proportional to the square root of the lead time (if demand is independent over time) and thus grows more slowly than the lead time itself.

Key Point

The required safety inventory increases with an increase in the lead time and the uncertainty of periodic demand.

A goal of any supply chain manager is to reduce the level of safety inventory required in a way that does not adversely affect product availability. The previous discussion highlights two key managerial levers that may be used to achieve this goal:

1. ***Reduce the supplier lead time L:*** If lead time decreases by a factor of k, the required safety inventory decreases by a factor of $\sqrt{k}$. The only caveat here is that reducing the supplier lead time requires significant effort from the supplier, whereas reduction in safety inventory occurs at the retailer. Thus, it is important for the retailer to share some of the resulting benefits, as discussed in Chapter 10. Walmart, Seven-Eleven Japan, and many other retailers apply tremendous pressure on their suppliers to reduce the replenishment lead time. Apparel retailer Zara has built its entire strategy around using local flexible production to reduce replenishment lead times. In each case, the benefit has manifested itself in the form of reduced safety inventory while maintaining the desired level of product availability.

2. ***Reduce the underlying uncertainty of demand (represented by σ_D):*** If uncertainty represented by σ_D is reduced by a factor of k, the required safety inventory also decreases by a factor of k. A reduction in uncertainty may be achieved by better market intelligence, increased

supply chain visibility, and the use of more sophisticated forecasting methods. Seven-Eleven Japan provides its store managers with detailed data about prior demand along with weather and other factors that may influence demand. This market intelligence allows the store managers to make better forecasts, reducing uncertainty. In most supply chains, however, the key to reducing the underlying forecast uncertainty is to link all forecasts throughout the supply chain to customer demand data. A lot of the demand uncertainty exists only because each stage of the supply chain plans and forecasts independently. This distorts demand throughout the supply chain, increasing uncertainty. Improved visibility and coordination, as discussed in Chapter 10, can often reduce the demand uncertainty significantly. Zara plans its production and replenishment based on actual sales at its retail stores to ensure that no unnecessary uncertainties are introduced. Both Walmart and Seven-Eleven Japan share demand information with their suppliers, reducing uncertainty and thus safety inventory within the supply chain.

We illustrate the benefits of reducing lead time and demand uncertainty in Example 12-6 (see worksheet *Example 12-6*).

EXAMPLE 12-6 Benefits of Reducing Lead Time and Demand Uncertainty

Weekly demand for white shirts at a Target store is normally distributed, with a mean of 2,500 and a standard deviation of 800. The replenishment lead time from the current supplier is nine weeks. The store manager aims for a cycle service level of 95 percent. What savings in safety inventory can the store expect if the supplier reduces lead time to one week? What savings in safety inventory can the store expect if reduced demand uncertainty results in a standard deviation of demand of 400?

Analysis:

For the base case, we have

$$D = 2{,}500/\text{week},\ \sigma_D = 800,\ CSL = 0.95$$

From Equation 12.5, we thus obtain the base case safety inventory to be

$$ss = NORMSINV(CSL) \times \sqrt{L}\sigma_D = NORMSINV(.95) \times \sqrt{9} \times 800 = 3{,}948$$

If the suppler reduces the lead time L to one week, the required safety inventory is given by

$$ss = NORMSINV(CSL) \times \sqrt{L}\sigma_D = NORMSINV(.95) \times \sqrt{1} \times 800 = 1{,}316$$

Thus, reducing the lead time from nine weeks to one week reduces the required safety inventory by 2,632 shirts.

We now consider the benefits of reducing forecast error. If Target reduces the standard deviation from 800 to 400 (for the nine-week lead time), the required safety inventory is obtained as follows:

$$ss = NORMSINV(CSL) \times \sqrt{L}\sigma_D = NORMSINV(.95) \times \sqrt{9} \times 400 = 1{,}974$$

Thus, reducing the standard deviation (equal to forecast error) of demand from 800 to 400 reduces the required safety inventory by 1,974 shirts.

12.4 IMPACT OF SUPPLY UNCERTAINTY ON SAFETY INVENTORY

In our discussion to this point, we have focused on situations with demand uncertainty in the form of a forecast error. In many practical situations, supply uncertainty also plays a significant role. The impact of supply uncertainty is well illustrated by the impact of the grounding of MSC

Napoli on the south coast of Britain in January 2007. The container ship was carrying more than 1,000 tons of nickel, a key ingredient of stainless steel. Given that 1,000 tons was almost 20 percent of the 5,052 tons of nickel then stored in warehouses globally, this delay in bringing nickel to market resulted in significant shortages and raised the price of nickel by about 20 percent in the first 3.5 weeks of January 2007. Supply uncertainty arises because of many factors, including production delays, transportation delays, and quality problems. Supply chains must account for supply uncertainty when planning safety inventories.

In this section, we incorporate supply uncertainty by assuming that lead time is uncertain and identify the impact of lead time uncertainty on safety inventories. Assume that the customer demand per period for tablets at Amazon and the replenishment lead time from the supplier are normally distributed. We are provided the following inputs:

D: Average demand per period

σ_D: Standard deviation of demand per period

L: Average lead time for replenishment

s_L: Standard deviation of lead time

We consider the safety inventory requirements given that Amazon follows a continuous review policy to manage tablet inventory. Amazon experiences a stockout of product if demand during the lead time exceeds the ROP—that is, the quantity on hand when Amazon places a replenishment order. Thus, we need to identify the distribution of customer demand during the lead time. Given that both lead time and periodic demand are uncertain, demand during the lead time is normally distributed with a mean of D_L and a standard deviation σ_L, where

$$D_L = D \times L; \quad \sigma_L = \sqrt{L\sigma_D^2 + D^2 s_L^2} \tag{12.11}$$

Given the distribution of demand during the lead time in Equation 12.11 and a desired CSL, Amazon can obtain the required safety inventory using Equation 12.5. If product availability is specified as a fill rate, Amazon can obtain the required safety inventory using the procedure outlined in Example 12-5. In Example 12-7, we illustrate the impact of lead time uncertainty on the required level of safety inventory at Amazon (see worksheet *Example 12-7*).

EXAMPLE 12-7 Impact of Lead Time Uncertainty on Safety Inventory

Daily demand for tablets at Amazon is normally distributed, with a mean of 2,500 and a standard deviation of 500. The tablet supplier takes an average of $L = 7$ days to replenish inventory at Amazon. Amazon is targeting a CSL of 90 percent (providing a fill rate close to 100 percent) for its tablet inventory. Evaluate the safety inventory of tablets that Amazon must carry if the standard deviation of the lead time is seven days. Amazon is working with the supplier to reduce the standard deviation to zero. Evaluate the reduction in safety inventory that Amazon can expect as a result of this initiative.

Analysis:

In this case, we have

Average demand per period, $D = 2{,}500$

Standard deviation of demand per period, $\sigma_D = 500$

Average lead time for replenishment, $L = 7$ days

Standard deviation of lead time, $s_L = 7$ days

TABLE 12-2 Required Safety Inventory as a Function of Lead Time Uncertainty

s_L	σ_L	*ss* (units)	*ss* (days)
6	15,058	19,298	7.72
5	12,570	16,109	6.44
4	10,087	12,927	5.17
3	7,616	9,760	3.90
2	5,172	6,628	2.65
1	2,828	3,625	1.45
0	1,323	1,695	0.68

We first evaluate the distribution of demand during the lead time. Using Equation 12.11, we have

$$\text{Mean demand during lead time, } D_L = D \times L = 2{,}500 \times 7 = 17{,}500$$

$$\text{Standard deviation of demand during lead time, } \sigma_L = \sqrt{L\sigma_D^2 + D^2 s_L^2}$$
$$= \sqrt{7 \times 500^2 + 2{,}500^2 \times 7^2} = 17{,}550$$

The required safety inventory is obtained using Equations 12.5 and 12.27, as follows:

$$ss = NORMSINV(CSL) \times \sigma_L = NORMSINV(0.90) \times 17{,}550 = 22{,}491 \text{ tablets}$$

If the standard deviation of lead time is seven days, Amazon must carry a safety inventory of 22,491 tablets. This is equivalent to about nine days of demand for tablets.

In Table 12-2, we provide the required safety inventory as Amazon works with the supplier to reduce the standard deviation of lead time (s_L) from six down to zero. From Table 12-2, observe that the reduction in lead time uncertainty allows Amazon to reduce its safety inventory of tablets by a significant amount. As the standard deviation of lead time declines from seven days to zero, the amount of safety inventory declines from about nine days of demand to less than a day of demand.

The preceding example emphasizes the impact of lead time variability on safety inventory requirements (and thus material flow time) and the large potential benefits from reducing lead time variability or improving on-time deliveries. Often, safety inventory calculations in practice do not include any measure of supply uncertainty, resulting in levels that may be lower than required. This hurts product availability.

Key Point

A reduction in supply uncertainty can help to dramatically reduce the required safety inventory without hurting product availability.

In practice, variability of supply lead time is caused by practices at both the supplier and the party receiving the order. Suppliers sometimes have poor planning tools that do not allow them to schedule production in a way that can be executed. Today, most supply chain planning software suites have good production planning tools that allow suppliers to promise lead times that can be met. This helps reduce lead time variability. The lack of visibility for a supplier into future customer plans is also a significant factor that increases supply chain uncertainty. W.W. Grainger was able to get its suppliers to reduce both lead time and lead time variability by sharing its future plans with them. This allowed suppliers to schedule Grainger orders into

production without waiting for the orders to actually arrive. The quantity produced was finalized closer to actual production. In other instances, the behavior of the party placing the order often increases lead time variability. In one instance, a distributor placed orders to all suppliers on the same day of the week. As a result, all deliveries arrived on the same day of the week. The surge in deliveries made it impossible for all of them to be recorded into inventory on the day they arrived. This led to a perception that supply lead times were long and variable. Just by leveling out the orders over the week, the lead time and the lead time variability were significantly reduced, allowing the distributor to reduce its safety inventory.

Next, we discuss how aggregation can help reduce the amount of safety inventory in the supply chain.

12.5 IMPACT OF AGGREGATION ON SAFETY INVENTORY

In practice, supply chains have varying degrees of inventory aggregation. For example, Barnes & Noble sells books from retail stores with inventory geographically distributed across the country. Amazon, in contrast, ships all its books from a few facilities. Seven-Eleven Japan has many small convenience stores densely distributed across Japan. In contrast, supermarkets tend to be much larger, with fewer outlets that are not as densely distributed. Redbox rents its movies from tens of thousands of kiosks distributed across the United States. In contrast, Netflix centralizes its DVD inventory at fewer than fifty distribution centers.

Our goal is to understand how aggregation in each of these cases affects forecast accuracy and safety inventories. Consider k regions, with demand in each region normally distributed with the following characteristics:

D_i: Mean periodic demand in region i, $i = 1, \ldots, k$

σ_i: Standard deviation of periodic demand in region i, $i = 1, \ldots, k$

ρ_{ij}: Correlation of periodic demand for regions i, j, $1 \leq i \neq j \leq k$

There are two ways to serve demand in the k regions. One is to have local inventories in each region and the other is to aggregate all inventories into one centralized facility. Our goal is to compare safety inventories in the two cases. With a replenishment lead time of L and a desired cycle service level CSL, the total safety inventory in the decentralized case is (using Equation 12.5):

$$\text{Total safety inventory in decentralized option} = \sum_{i=1}^{k} F_S^{-1}(CSL) \times \sqrt{L} \times \sigma_i \quad \textbf{(12.12)}$$

If all inventories are aggregated in a central location, we need to evaluate the distribution of aggregated demand. The aggregate demand is normally distributed, with a mean of D^C, standard deviation of σ_D^C, and a variance of $\text{var}(D^C)$, as follows:

$$D^C = \sum_{i=1}^{k} D_i; \quad \text{var}(D^C) = \sum_{i=1}^{k} \sigma_i^2 + 2\sum_{i>j} \rho_{ij}\sigma_i\sigma_j; \quad \sigma_D^C = \sqrt{\text{var}(D^C)} \quad \textbf{(12.13)}$$

Observe that Equation 12.13 is like Equation 12.1 except that we are aggregating across k regions rather than L periods. If all k regions have demand that is identically distributed, with mean D and standard deviation σ_D, and have the same correlation ρ, Equation 12.13 can be simplified as

$$D^C = kD; \sigma_D^C = \sqrt{k\sigma^2 + k(k-1)\rho\sigma^2} \quad \textbf{(12.14)}$$

If all k regions have demand that is independent ($\rho_{ij} = 0$) and identically distributed, with mean D and standard deviation σ_D, Equation 12.13 can be simplified as

$$D^C = kD; \sigma_D^C = \sqrt{k}\sigma_D \quad \textbf{(12.15)}$$

Using Equations 12.5 and 12.13, the required safety inventory at the centralized location is given as

$$\text{Required safety inventory on aggregation} = F_S^{-1}(CSL) \times \sqrt{L} \times \sigma_D^C \quad \textbf{(12.16)}$$

The holding cost savings on aggregation per unit sold are obtained by dividing the savings in holding cost by the total demand kD. If H is the holding cost per unit, using Equations 12.12 and 12.16, the savings per units are

$$\text{Holding - cost savings on aggregation per unit sold} = \frac{F_S^{-1}(CSL) \times \sqrt{L} \times H}{D^C} \times \left(\sum_{i=1}^{k} \sigma_i - \sigma_D^C \right) \quad \textbf{(12.17)}$$

From Equation 12.13, it follows that the difference $(\sum_{i=1}^{k} \sigma_i - \sigma_D^C)$ is influenced by the correlation coefficients ρ_{ij}. This difference is large when the correlation coefficients are close to -1 (negative correlation) and shrinks as they approach $+1$ (positive correlation). Inventory savings on aggregation are always positive as long as the correlation coefficients are less than 1. From Equation 12.17, we thus draw the following conclusions regarding the value of aggregation:

- The safety inventory savings on aggregation increase with the desired cycle service level *CSL*.
- The safety inventory savings on aggregation increase with the replenishment lead time L.
- The safety inventory savings on aggregation increase with the holding cost H.
- The safety inventory savings on aggregation increase with the coefficient of variation (σ_D/D) of demand.
- The safety inventory savings on aggregation decrease as the correlation coefficients increase.

In Example 12-8 (see worksheet *Example 12-8*), we illustrate the inventory savings on aggregation and the impact of the correlation coefficient on these savings.

EXAMPLE 12-8 Impact of Correlation on Value of Aggregation

A BMW dealership has $k = 4$ retail outlets serving the entire Chicago area (disaggregate option). Weekly demand at each outlet is normally distributed, with a mean of $D = 25$ cars and a standard deviation of $\sigma_D = 5$. The lead time for replenishment from the manufacturer is $L = 2$ weeks. Each outlet covers a separate geographic area, and the correlation of demand across any pair of areas is ρ. The dealership is considering the possibility of replacing the four outlets with a single large outlet (aggregate option). Assume that the demand in the central outlet is the sum of the demand across all four areas. The dealership is targeting a CSL of 0.90. Compare the level of safety inventory needed in the two options as the correlation coefficient ρ varies between 0 and 1.

Analysis:

We provide a detailed analysis for the case when demand in each area is independent (i.e., $\rho = 0$). For each retail outlet we have

Standard deviation of weekly demand, $\sigma_D = 5$

Replenishment lead time, $L = 2$ weeks

Using Equation 12.12, the required safety inventory in the decentralized option for $CSL = 0.90$ is

$$\text{Total required safety inventory, } ss = k \times F_s^{-1}(CSL) \times \sqrt{L} \times \sigma_D$$
$$= 4 \times F_s^{-1}(0.9) \times \sqrt{2} \times 5 = 4 \times NORMSINV(0.9) \times \sqrt{2} \times 5 = 36.25 \text{ cars}$$

Now, consider the aggregate option. Using Equation 12.14, the standard deviation of aggregate weekly demand is

$$\text{Standard deviation of weekly demand at central outlet, } \sigma_D^C = \sqrt{4 \times 5^2 + 4 \times 3 \times 5^2 \times \rho}$$

TABLE 12-3 Safety Inventory in the Disaggregate and Aggregate Options

ρ	Disaggregate Safety Inventory	Aggregate Safety Inventory
0	36.25	18.12
0.2	36.25	22.93
0.4	36.25	26.88
0.6	36.25	30.33
0.8	36.25	33.42
1.0	36.25	36.25

For a CSL of 0.90 and $\rho = 0$, safety inventory required for the aggregate option (using Equation 12.16) is given as

$$ss = F_S^{-1}(0.90) \times \sqrt{L} \times \sigma_D^C = NORMSINV(0.90) \times \sqrt{2} \times 10 = 18.12$$

Using Equations 12.12 to 12.16, the required level of safety inventory for the disaggregate as well as the aggregate option can be obtained for different values of ρ as shown in Table 12-3 using worksheet *Example 12-8*. Observe that the safety inventory for the disaggregate option is higher than for the aggregate option except when all demands are perfectly positively correlated. The benefit of aggregation decreases as demand in different areas is more positively correlated.

Example 12-8 and the previous discussion demonstrate that aggregation reduces demand uncertainty—and, thus, the required safety inventory—as long as the demand being aggregated is not perfectly positively correlated. Demand for most products does not show perfect positive correlation across different geographic regions. Products such as heating oil are likely to have demand that is positively correlated across nearby regions. In contrast, products such as milk and sugar are likely to have demand that is much more independent across regions. If demand in different geographic regions is about the same size and independent, aggregation reduces safety inventory by the square root of the number of regions aggregated. In other words, if the number of independent stocking locations decreases by a factor of n, the average safety inventory is expected to decrease by a factor of $\sqrt{n}$. This principle is referred to as the *square-root law* and is illustrated in Figure 12-4.

Most online retailers exploit the benefits of aggregation in terms of reduced inventories. For example, Blue Nile sells diamonds online and serves the entire United States out of one

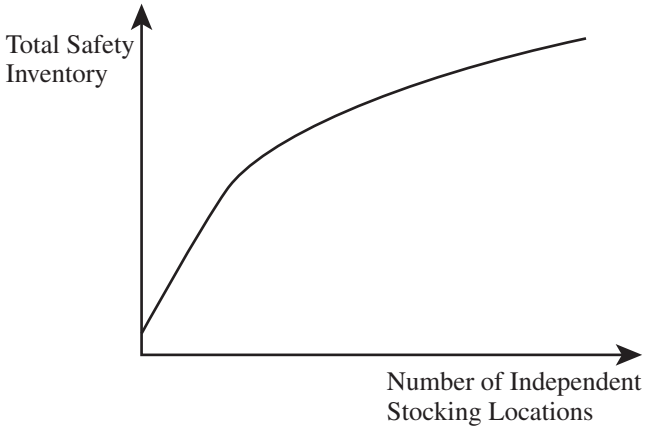

FIGURE 12-4 Square-Root Law

warehouse. As a result, it has lower levels of diamond inventories than jewelry chains such as Tiffany and Zales, which must keep inventory in every retail store.

There are situations, however, in which physical aggregation of inventories in one location may not be optimal. There are two major disadvantages of aggregating all inventories in one location:

1. Increase in response time to customer order
2. Increase in transportation cost to customer

Both disadvantages result because the average distance between the inventory and the customer increases with aggregation. Either the customer must travel farther to reach the product or the product must be shipped over longer distances to reach the customer. A retail chain such as Gap has the option of building many small retail outlets or a few large ones. Gap tends to have many smaller outlets distributed evenly in a region because this strategy reduces the distance that customers travel to reach a store. If Gap had one large centralized outlet, the average distance that customers need to travel would increase, and thus the response time would increase. A desire to decrease customer response time is thus the impetus for the firm to have multiple outlets. Another example is McMaster-Carr, a distributor of MRO supplies. McMaster-Carr uses UPS for shipping product to customers. Because shipping charges are based on distance, having one centralized warehouse increases the average shipping cost as well as the response time to the customer. Thus, McMaster-Carr has five warehouses that allow it to provide next-day delivery to a large fraction of the United States. Next-day delivery by UPS would not be feasible at a reasonable cost if McMaster-Carr had only one warehouse. Even Amazon, which started with one warehouse in Seattle, has added more warehouses in other parts of the United States in an effort to improve response time and reduce transportation cost to the customer. We illustrate the trade-offs of centralization in Example 12-9 (see worksheet *Example 12-9*).

EXAMPLE 12-9 Trade-Offs of Physical Centralization

The Shanghai branch office of an Italian coffee machine company is considering setting up either one distribution center for each of its east, south, west, and north regions or simply one center in Ningbo for the whole of China. The weekly demand for the automatic espresso coffee machine is normally distributed with a mean of 1,000 units and a standard deviation of 300 units. Although the demand for each region is independent, the supply lead time is more or less the same—four weeks. Each machine costs $1,000 and the holding cost is 20 percent. With the next-day delivery promise, the branch office needs to bear an inland trucking cost of $10/unit for all four regional centers. However, if a single national distribution center is decided upon, a more expensive transport fleet is needed and that will cost $13/unit for next-day service. Setting up and operating four regional DCs costs $150,000 per year more than building and operating the single Ningbo national distribution center.

Assume that the Italian company would like a CSL of 0.95. What should the Shanghai branch office decide based on the cost considerations?

Analysis:

Observe that using only one distribution center would decrease facility and inventory costs but increase transportation costs. We therefore have to evaluate the change in each cost category on aggregation.

We start with inventory costs. For each region we have

$$D = \frac{1{,}000}{\text{week}}, \sigma_D = 300, L = 4 \text{ weeks}$$

Given the desired $CSL = 0.95$, the required safety inventory across all four regional distribution centers is obtained using Equation 12.9 to be

$$\begin{aligned}\text{Total required safety inventory, } ss &= 4 \times F_s^{-1}(CSL) \times \sqrt{L} \times \sigma_D \\ &= 4 \times NORMSINV(0.95) \times \sqrt{4} \times 300 = 3{,}948\end{aligned}$$

Now consider the aggregate option. Because demand in all four areas is independent, $\rho = 0$. Using Equation 12.14, the standard deviation of aggregate weekly demand is

$$\text{Standard deviation of weekly demand at national distribution center, } \sigma_D^C = \sqrt{4} \times 300 = 600$$

For a CSL of 0.95, safety inventory required for the aggregate option (using Equation 12.15) is given as

$$ss = F_S^{-1}(0.95) \times \sqrt{L} \times \sigma_D^C = NORMSINV(0.95) \times \sqrt{4} \times 600 = 1{,}974$$

We can now evaluate the changes in inventory, transportation, and facility costs upon aggregation as follows:

$$\text{Decrease in annual inventory holding cost on aggregation} = (3{,}948 - 1{,}974) \times \$1{,}000 \times 0.2 = \$394{,}765$$

$$\text{Decrease in annual facility costs on aggregation} = \$150{,}000$$

$$\text{Increase in annual transportation costs on aggregation} = 52 \times 1{,}000 \times (13 - 10) \times 4 = \$624{,}000$$

Observe that in this case, the annual costs for the Shanghai branch office will be increased by $\$624{,}000 - \$394{,}765 - \$150{,}000 = \$79{,}235$ upon centralization. It is clearly better to run the four centers in the east, south, west, and north regions instead of the Ningbo center.

Example 12-9 and the previous discussion highlight instances in which physical aggregation of inventory at one location may not be optimal. However, aggregating safety inventory has clear benefits. We now discuss various methods by which a supply chain can extract the benefits of aggregation without having to physically centralize all inventories in one location.

Information Centralization

Redbox uses *information centralization* to virtually aggregate its inventories of DVDs despite having tens of thousands of vending machines. The company has set up an online system that allows customers to locate nearby vending machines with the DVD they are searching for in stock. This allows Redbox to provide a much higher level of product availability than would be possible if a customer found out about availability only by visiting a vending machine. The benefit of information centralization derives from the fact that most customers get their DVD from the vending machine closest to their house. In case of a stockout at the closest vending machine, the customer is served from another vending machine, thus improving product availability without adding to inventories.

Retailers such as Gap also use information centralization effectively. If a store does not have the size or color that a customer wants, store employees can use their information system to inform the customer of the closest store with the product in inventory. Customers can then either go to this store or have the product delivered to their house. Gap thus uses information centralization to virtually aggregate inventory across all retail stores even though the inventory is physically separated. This allows Gap to reduce the amount of safety inventory it carries while providing a high level of product availability.

Walmart has an information system in place that allows store managers to search other stores for an excess of items that may be hot sellers at their stores. Walmart provides transportation that allows store managers to exchange products so they arrive at stores where they are in high demand. In this case, Walmart uses information centralization with a responsive transportation system to reduce the amount of safety inventory carried while providing a high level of product availability.

Specialization

Most supply chains provide a variety of products to customers. When inventory is carried at multiple locations, a key decision for a supply chain manager is whether all products should be

stocked at every location. Clearly, a product that does not sell in a geographic region should not be carried in inventory by the warehouse or retail store located there. For example, it does not make sense for a Sears retail store in southern Florida to carry a wide variety of snow boots in inventory.

Another important factor that must be considered when making stocking decisions is the reduction in safety inventory that results from aggregation. If aggregation reduces the required safety inventory for a product by a large amount, it is better to carry the product in one central location. If aggregation reduces the required safety inventory for a product by a small amount, it may be best to carry the product in multiple decentralized locations to reduce response time and transportation cost.

The reduction in safety inventory due to aggregation is strongly influenced by the demand's coefficient of variation. For a product with a low coefficient of variation, disaggregate demand can be forecast with accuracy. As a result, the inventory benefit from aggregation is minimal. For a product with a high coefficient of variation of demand, disaggregate demand is difficult to forecast. In this case, aggregation improves forecast accuracy significantly, providing great benefits. We illustrate this idea in Example 12-10 (see worksheet *Example 12-10*).

EXAMPLE 12-10 Impact of Coefficient of Variation on Value of Aggregation

Assume that W.W. Grainger, a supplier of MRO products, has 1,600 stores distributed throughout the United States. Consider two products—large electric motors and industrial cleaner. Large electric motors are high-value items with low demand, whereas the industrial cleaner is a low-value item with high demand. Each motor costs $500 and each can of cleaner costs $30. Weekly demand for motors at each store is normally distributed, with a mean of 20 and a standard deviation of

TABLE 12-4 Value of Aggregation at W.W. Grainger

	Motors	Cleaner
Inventory is stocked in each store		
Mean weekly demand per store	20	1,000
Standard deviation	40	100
Coefficient of variation	2.0	0.1
Safety inventory per store	132	329
Total safety inventory	211,200	526,400
Value of safety inventory	$105,600,000	$15,792,000
Inventory is aggregated at the DC		
Mean weekly aggregate demand	32,000	1,600,000
Standard deviation of aggregate demand	1,600	4,000
Coefficient of variation	0.05	0.0025
Aggregate safety inventory	5,264	13,159
Value of safety inventory	$2,632,000	$394,770
Savings		
Total inventory saving on aggregation	$102,968,000	$15,397,230
Total holding cost saving on aggregation	$25,742,000	$3,849,308
Holding cost saving per unit sold	$15.47	$0.046
Savings as a percentage of product cost	3.09%	0.15%

40. Weekly demand for cleaner at each store is normally distributed, with a mean of 1,000 and a standard deviation of 100. Demand experienced by each store is independent, and supply lead time for both motors and cleaner is four weeks. W.W. Grainger has a holding cost of 25 percent. For each of the two products, evaluate the reduction in safety inventories that will result if they are removed from retail stores and carried only in a centralized DC. Assume a desired CSL of 0.95.

Analysis:

The evaluation of safety inventories and the value of aggregation for each of the two products is shown in Table 12-4. All calculations use the approach discussed earlier and illustrated in Example 12-8. As Table 12-4 shows, the inventory reduction benefit from centralizing motors is much larger than the benefit from centralizing cleaner. From this analysis, W.W. Grainger should stock cleaner at the stores and motors in the DC. Given that cleaner is a high-demand item, customers will be able to pick it up on the same day at the stores. Given that motors are a low-demand item, customers may be willing to wait the extra day that shipping from the DC will entail.

Key Point

The higher the coefficient of variation of an item, the greater is the reduction in safety inventories as a result of centralization.

Items with low demand are referred to as *slow-moving items* and typically have a high coefficient of variation, whereas items with high demand are referred to as *fast-moving items* and typically have a low coefficient of variation. For many supply chains, specializing the distribution network with fast-moving items stocked at decentralized locations close to the customer and slow-moving items stocked at a centralized location can significantly reduce the safety inventory carried without hurting customer response time or adding to transportation costs. The centralized location then specializes in handling slow-moving items.

Of course, other factors also need to be considered when deciding on the allocation of products to stocking locations. For example, an item that is considered an emergency item because the customer needs it urgently may be stocked at stores even if it has a high coefficient of variation. In this case the customer will be willing to pay a premium for having the item available at a store. One also needs to consider the cost of the item. High-value items provide a greater benefit from centralization than do low-value items.

The insights from Example 12-10 and the above discussion are summarized in Figure 12-5. In general, decentralized networks like Costco provide a low-cost supply chain for fast-moving, predictable, low-value products like detergent. Centralized networks like Blue Nile provide a low-cost supply chain for slow-moving, unpredictable, high-value products like diamonds. A decentralized supply chain like Tiffany may carry slow-moving items like diamonds as long as customers are willing to pay a premium for this choice. Similarly, a centralized supply chain like Amazon may carry a fast-moving item like detergent, but only if customers are willing to pay a premium. It can be argued that Amazon's inability to extract a significant enough premium from its customers for the fast-moving items it sells has hurt its profitability.

Item Type	Centralized Inventories	Decentralized Inventories
Fast Moving Predictable {Low value}	Customer willing to pay premium?	Low cost
Slow Moving Unpredictable {High value}	Low cost	Customer willing to pay premium?

FIGURE 12-5 Specialization of Inventory Based on Product Type

It is important for firms with bricks-and-mortar stores to take the idea of specialization into account when incorporating the online channel into an omni-channel strategy. Consider, for example, a bookstore chain such as Barnes & Noble, which carries about 100,000 titles at each retail store. The titles carried can be divided into two broad categories—best sellers with high demand and other books with much lower demand. Barnes & Noble can design an omni-channel strategy under which the retail stores carry primarily best sellers in inventory. They may also carry one, or at most two, copies of each of the other titles, to allow customers to browse. Customers should be able to access all titles that are not in the store via electronic kiosks in the store, which provide access to barnesandnoble.com inventory. This strategy allows customers to access an increased variety of books from Barnes & Noble stores. Customers could place orders for low-volume titles with barnesandnoble.com while purchasing high-volume titles at the store itself. This strategy of specialization would allow Barnes & Noble to aggregate all slow-moving items to be sold by the online channel. All best sellers would be decentralized and carried close to the customer. The supply chain thus reduces inventory costs for slow-moving items at the expense of somewhat higher transportation costs. For the fast-moving items, the supply chain provides a lower transportation cost and better response time by carrying the items at retail stores close to the customer.

Home Depot follows a similar strategy and integrates its online channel with its retail stores. The retail stores carry fast-moving items, and the customer is able to order slow-moving variants online. Home Depot is thus able to increase the variety of products available to customers while keeping supply chain inventories down. Walmart.com has also employed a strategy of carrying slower-moving items online.

Product Substitution

Substitution refers to the use of one product to satisfy demand for a different product. Substitution may occur in two situations:

1. ***Manufacturer-driven substitution:*** The manufacturer or supplier makes the decision to substitute. Typically, the manufacturer substitutes a higher-value product for a lower-value product that is not in inventory. For example, Dell may install a 1.2-terabyte hard drive into a customer order requiring a 1-terabyte drive if the smaller drive is out of stock.
2. ***Customer-driven substitution:*** Customers make the decision to substitute. For example, a customer walking into a Walmart store to buy a gallon of detergent may buy the half-gallon size if the gallon size is not available. The customer substitutes the half-gallon size for the gallon size.

In both cases, exploiting substitution allows the supply chain to satisfy demand using aggregate inventories, which permits the supply chain to reduce safety inventories without hurting product availability. In general, given two products or components, substitution may be one-way (i.e., only one of the products [components] substitutes for the other) or two-way (i.e., either product [component] substitutes for the other). We briefly discuss one-way substitution in the context of manufacturer-driven substitution and two-way substitution in the context of customer-driven substitution.

MANUFACTURER-DRIVEN ONE-WAY SUBSTITUTION Consider a server manufacturer selling direct to customers that offers drives that vary in size from 0.8 to 1.2 terabytes. Customers are charged according to the size of drive that they select, with larger sizes being more expensive. If a customer orders a 1-terabyte drive and the manufacturer is out of drives of this size, there are two possible choices: (1) delay or deny the customer order or (2) substitute a larger drive that is in stock (say, a 1.2-terabyte drive) and fill the customer order on time. The first case is potentially a lost sale or loss of future sales because the customer experiences a delayed delivery. In the second case, the manufacturer installs a higher-cost component, reducing the company's profit margin. These factors, along with the fact that only larger drives can substitute for

smaller drives, must be considered when the manufacturer makes inventory decisions for individual drive sizes.

Substitution allows the server manufacturer to aggregate demand across the components, reducing safety inventories required. The value of substitution increases as demand uncertainty increases. Thus, the manufacturer should consider substitution for components displaying high demand uncertainty.

The desired degree of substitution is influenced by the cost differential between the higher-value and lower-value component. If the cost differential is very small, the manufacturer should aggregate most of the demand and carry most of its inventory in the form of the higher-value component. As the cost differential increases, though, the benefit of substitution decreases. In this case, the manufacturer will find it more profitable to carry inventory of each of the two components and decrease the amount of substitution.

The desired level of substitution is also influenced by the correlation of demand between the products. If demand between two components is strongly positively correlated, there is little value in substitution. As demand for the two components becomes less positively correlated (or even negatively correlated), the benefit of substitution increases.

Key Point

Manufacturer-driven substitution increases overall profitability for the manufacturer by allowing some aggregation of demand, which reduces the inventory requirements for the same level of availability.

CUSTOMER-DRIVEN TWO-WAY SUBSTITUTION Consider W.W. Grainger selling two brands of motors, GE and SE, which have similar performance characteristics. Customers are generally willing to purchase either brand, depending on product availability. If W.W. Grainger managers do not recognize customer substitution, they will not encourage it. For a given level of product availability, they will thus have to carry high levels of safety inventory of each brand. If its managers recognize and encourage customer substitution, they can aggregate the safety inventory across the two brands, thereby improving product availability.

W.W. Grainger does a good job of recognizing customer substitution. When a customer calls or goes online to place an order and the product he or she requests is not available, the customer is immediately told the availability of all equivalent products that may be substituted. Most customers ultimately buy a substitute product in this case. W.W. Grainger exploits this substitution by managing safety inventory of all substitutable products jointly. Recognition and exploitation of customer substitution allows W.W. Grainger to provide a high level of product availability with lower levels of safety inventory.

A good understanding of customer-driven substitution is important in the retail industry. It must be exploited when merchandising to ensure that substitute products are placed near each other, allowing a customer to buy one if the other is out of stock. In the online channel, substitution requires a retailer to present the availability of substitute products if the one the customer requests is out of stock. The supply chain is thus able to reduce the required level of safety inventory while providing a high level of product availability.

Key Point

Recognition of customer-driven substitution and joint management of inventories across substitutable products allows a supply chain to reduce the required safety inventory while ensuring a high level of product availability.

The demand uncertainties and the correlation of demand between the substitutable products influence the benefit to a retailer from exploiting substitution. The greater the demand

uncertainty, the greater is the benefit of substitution. The less positive the correlation of demand between substitutable products, the greater is the benefit from exploiting substitution.

Component Commonality

In any supply chain, a significant amount of inventory is held in the form of components. A single product such as a server contains hundreds of components. When a supply chain is producing a large variety of products, component inventories can easily become very large. The use of common components in a variety of products is an effective supply chain strategy to exploit aggregation and reduce component inventories.

Dell sells thousands of server configurations to customers. An extreme option for Dell is to design distinct components that are suited to the performance of a particular configuration. Under this option, Dell would use different memory, hard drive, and other components for each distinct finished product. The other option is to design products such that common components are used in different finished products.

Without common components, the uncertainty of demand for any component is the same as the uncertainty of demand for the finished product in which it is used. Given the large number of components in each finished product, demand uncertainty will be high, resulting in high levels of safety inventory. When products with common components are designed, the demand for each component is an aggregation of the demand for all the finished products of which the component is a part. Component demand is thus more predictable than the demand for any one finished product. This fact reduces the component inventories carried in the supply chain. This idea has been a key factor for success in the electronics industry and has also started to play a big role in the auto industry. With increasing product variety, component commonality is a key to reducing supply chain inventories without hurting product availability. We illustrate the basic idea behind component commonality in Example 12-11 (see worksheet *Example 12-11*).

EXAMPLE 12-11 Value of Component Commonality

Assume that Dell is to manufacture 27 servers with three distinct components: processor, memory, and hard drive. Under the disaggregate option, Dell designs specific components for each server, resulting in $3 \times 27 = 81$ distinct components. Under the common-component option, Dell designs servers such that three distinct processors, three distinct memory units, and three distinct hard drives can be combined to create 27 servers. Each component is thus used in nine servers. Monthly demand for each of the 27 servers is independent and normally distributed, with a mean of 5,000 and a standard deviation of 3,000. The replenishment lead time for each component is one month. Dell is targeting a CSL of 95 percent for component inventory. Evaluate the safety inventory requirements with and without the use of component commonality. Also evaluate the change in safety inventory requirements as the number of finished products of which a component is a part varies from one to nine.

Analysis:

We first evaluate the disaggregate option, in which components are specific to a server. For each component, we have

$$\text{Standard deviation of monthly demand} = 3{,}000$$

Given a lead time of one month and a total of 81 components across 27 servers, we thus use Equation 12.12 to obtain

$$\text{Total safety inventory required} = 81 \times NORMSINV(0.95) \times \sqrt{1} \times 3{,}000 = 399{,}699 \text{ units}$$

TABLE 12-5 Marginal Benefit of Component Commonality

Number of Finished Products per Component	Safety Inventory	Marginal Reduction in Safety Inventory	Total Reduction in Safety Inventory
1	399,699		
2	282,630	117,069	117,069
3	230,766	51,864	168,933
4	199,849	30,917	199,850
5	178,751	21,098	220,948
6	163,176	15,575	236,523
7	151,072	12,104	248,627
8	141,315	9,757	258,384
9	133,233	8,082	266,466

In the case of component commonality, each component ends up in nine finished products. Therefore, the demand at the component level is the sum of demand across nine products. Using Equations 12.15 and 12.16, the safety inventory required for each component is thus

$$\begin{aligned}\text{Safety inventory per common component} &= NORMSINV(0.95) \times \sqrt{1} \times \sqrt{9} \times 3{,}000 \\ &= 14{,}803.68 \text{ units}\end{aligned}$$

With component commonality, there are a total of nine distinct components. The total safety inventory across all nine components is thus

$$\text{Total safety inventory required} = 9 \times 14{,}803.68 = 133{,}233$$

Thus, having each component common to nine products results in a reduction in safety inventory for Dell from 399,699 to 133,233 units.

In Table 12-5, we evaluate the marginal benefit in terms of reduction in safety inventory as a result of increasing component commonality. Starting with the required safety inventory when each component is used in only one finished product, we evaluate the safety inventory as the number of products in which a component is used increases to nine. Observe that component commonality decreases the required safety inventory for Dell. The marginal benefit of commonality, however, declines as a component is used in more and more finished products.

As a component is used in more finished products, it must be more flexible. As a result, the cost of producing the component typically increases with increasing commonality. Given that the marginal benefit of component commonality decreases as we increase commonality, we need to trade off the increase in component cost and the decrease in safety inventory when deciding on the appropriate level of component commonality.

Key Point

Component commonality decreases the safety inventory required. The marginal benefit, however, decreases with increasing commonality.

Postponement

Postponement is the ability of a supply chain to delay product differentiation or customization until closer to the time the product is sold. The goal is to have common components in the supply chain for most of the push phase and move product differentiation as close to the pull phase of

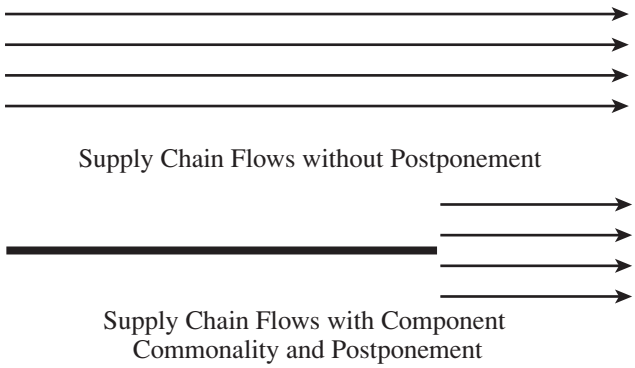

FIGURE 12-6 Supply Chain Flows without and with Postponement

the supply chain as possible. For example, the final mixing of paint today is done at the retail store after the customer has selected the color he or she wants. Thus, paint variety is produced only when demand is known with certainty. Postponement coupled with component commonality allows paint retailers to carry significantly lower safety inventories than in the past, when mixing was done at the paint factory. In the past, the factory manager had to forecast paint demand by color when planning production. Today, a factory manager needs to forecast only aggregate paint demand because mixing has been postponed until after customer demand is known. As a result, each retail store primarily carries aggregate inventory in the form of base paint that is configured to the appropriate color based on customer demand.

Another classic example of postponement is the production process at Benetton to make colored knit garments. The original process called for the thread to be dyed and then knitted and assembled into garments. The entire process required up to six months. Because the color of the final garment was fixed the moment the thread was dyed, demand for individual colors had to be forecast far in advance (up to six months). Benetton developed a manufacturing technology that allowed it to dye knitted garments to the appropriate color. Now, *greige* thread (the term used for thread that has not yet been dyed) can be purchased, knitted, and assembled into garments before dyeing. The dyeing of the garments is done much closer to the selling season. In fact, part of the dyeing is done after the start of the selling season, when demand is known with great accuracy. In this case, Benetton has postponed the color customization of the knit garments. When thread is purchased, only the aggregate demand across all colors needs to be forecast. Given that this decision is made far in advance, when forecasts are least likely to be accurate, there is great advantage to this aggregation. As Benetton moves closer to the selling season, the forecast uncertainty reduces. At the time Benetton dyes the knit garments, demand is known with a high degree of accuracy. Thus, postponement allows Benetton to exploit aggregation and significantly reduce the level of safety inventory carried. Supply chain flows with and without postponement are illustrated in Figure 12-6.

Without component commonality and postponement, product differentiation occurs early on in the supply chain, and most of the supply chain inventories are disaggregate. Postponement allows the supply chain to delay product differentiation. As a result, most of the inventories in the supply chain are aggregate. Postponement thus allows a supply chain to exploit aggregation to reduce safety inventories without hurting product availability. We illustrate the benefits of postponement in Example 12-12 (see worksheet *Example 12-12*). A more nuanced discussion of the value of postponement is given in Chapter 13.

EXAMPLE 12-12 Value of Postponement

Consider a paint retailer that sells 100 different colors of paint. Assume that weekly demand for each color is independent and is normally distributed with a mean of 30 and a standard deviation of 10. The replenishment lead time from the paint factory is two weeks and the retailer

aims for a CSL = 0.95. How much safety stock will the retailer have to hold if paint is mixed at the factory and held in inventory at the retailer as individual colors? How does the safety stock requirement change if the retailer holds base paint (supplied by the paint factory) and mixes colors on demand?

Analysis:

We first evaluate the disaggregate option without postponement, in which the retailer holds safety inventory for each color sold. For each color, we have

$$D = 30/\text{week}, \sigma_D = 10, L = 2 \text{ weeks}$$

Given the desired CSL = 0.95, the required safety inventory across all 100 colors is obtained using Equation 12.12 to be

$$\begin{aligned}\text{Total required safety inventory, } ss &= 100 \times F_S^{-1}(CSL) \times \sqrt{L} \times \sigma_D \\ &= 100 \times NORMSINV(0.95) \times \sqrt{2} \times 10 = 2{,}326\end{aligned}$$

Now, consider the option whereby mixing is postponed until after the customer orders. Safety inventory is held in the form of base paint, whose demand is an aggregate of demand of the 100 colors. Because demand in all 100 colors is independent, $\rho = 0$. Using Equation 12.15, the standard deviation of aggregate weekly demand of base paint is

$$\text{Standard deviation of weekly demand of base paint, } \sigma_D^C = \sqrt{100} \times 10 = 100$$

For a CSL of 0.95, safety inventory required for the aggregate option (using Equation 12.16) is given as

$$ss = F_S^{-1}(0.95) \times \sqrt{L} \times \sigma_D^C = NORMSINV(0.95) \times \sqrt{2} \times 100 = 233$$

Observe that postponement reduces the required safety inventory at the paint retailer from 2,326 units to 233 units.

Postponement can be a powerful concept when customers are willing to wait a little for their orders to arrive. This delay offers the supply chain an opportunity to reduce inventories by postponing product differentiation until after the customer order arrives. It is important that the manufacturing process be designed in a way that enables assembly to be completed quickly. Given that customers are often willing to wait for delivery, several furniture and window manufacturers have postponed some of the assembly processes for their products.

12.6 IMPACT OF REPLENISHMENT POLICIES ON SAFETY INVENTORY

In this section, we describe the evaluation of safety inventories for both continuous and periodic-review replenishment policies. We highlight the fact that periodic review policies require more safety inventory than continuous review policies for the same level of product availability. To simplify the discussion, we focus on the CSL as the measure of product availability. The managerial implications are the same if we use fill rate; the analysis, however, is more cumbersome.

Continuous Review Policies

Given that continuous review policies were discussed in detail in Section 12.2, we reiterate only the main points here. When using a continuous review policy, a manager orders Q units when the inventory drops to the ROP. Clearly, a continuous review policy requires technology that monitors the level of available inventory. This is the case for many firms such as Walmart and Dell, whose inventories are monitored continuously.

Given a desired CSL, our goal is to identify the required safety inventory *ss* and the ROP. We assume that demand is normally distributed, with the following inputs:

D: Average demand per period

σ_D: Standard deviation of demand per period

L: Average lead time for replenishment

The ROP represents the available inventory to meet demand during the lead time *L*. A stockout occurs if the demand during the lead time is larger than the ROP. If demand across periods is independent, demand during the lead time is normally distributed with the following:

$$\text{Mean demand during lead time, } D_L = D \times L$$
$$\text{Standard deviation of demand during lead time, } \sigma_L = \sqrt{L}\sigma_D$$

Given the desired CSL, the required safety inventory (*ss*) obtained using Equation 12.5 and the ROP obtained using Equation 12.3 are

$$ss = F_S^{-1}(CSL) \times \sigma_L = NORMSINV(CSL) \times \sqrt{L}\sigma_D, ROP = D_L + ss$$

A manager using a continuous review policy has to account only for the uncertainty of demand during the lead time. This is because the continuous monitoring of inventory allows the manager to adjust the timing of the replenishment order, depending on the demand experienced. If demand is very high, inventory reaches the ROP quickly, leading to a quick replenishment order. If demand is very low, inventory drops slowly to the ROP, leading to a delayed replenishment order. The manager, however, has no recourse during the lead time once a replenishment order has been placed. The available safety inventory thus must cover for the uncertainty of demand over this period.

Typically, in continuous review policies, the lot size ordered is kept fixed between replenishment cycles. The optimal lot size may be evaluated using the EOQ formula discussed in Chapter 11.

Periodic Review Policies

In periodic review policies, inventory levels are reviewed after a fixed period of time *T* and an order is placed such that the level of current inventory plus the replenishment lot size equals a prespecified level called the *order-up-to level* (OUL). The *review interval* is the time *T* between successive orders. Observe that the size of each order may vary, depending on the demand experienced between successive orders and the resulting inventory at the time of ordering. Periodic review policies are simpler for retailers to implement because they do not require that the retailer have the capability of monitoring inventory continuously. Suppliers may also prefer them because they result in replenishment orders placed at regular intervals.

Let us consider the store manager at Walmart who is responsible for designing a replenishment policy for Lego building blocks. He wants to analyze the impact on safety inventory if he decides to use a periodic review policy. Demand for Legos is normally distributed and independent from one week to the next. We assume the following inputs:

D: Average demand per period

σ_D: Standard deviation of demand per period

L: Average lead time for replenishment

T: Review interval

CSL: Desired cycle service level

To understand the safety inventory requirement, we track the sequence of events over time as the store manager places orders. The store manager places the first order at time 0

such that the lot size ordered and the inventory on hand sum to the OUL. Once an order is placed, the replenishment lot arrives after the lead time L. The next review period is time T, when the store manager places the next order, which then arrives at time $T + L$. The OUL represents the inventory available to meet all demand that arises between periods 0 and $T + L$. The Walmart store will experience a stockout if demand during the time interval between 0 and $T + L$ exceeds the OUL. Thus, the store manager must identify an OUL such that the following is true:

$$\text{Probability}(\text{demand during } T + L \leq OUL) = CSL$$

The next step is to evaluate the distribution of demand during the time interval $T + L$. Using Equation 12.2, demand during the time interval $T + L$ is normally distributed, with

$$\text{Mean demand during } T + L \text{ periods, } D_{T+L} = (T + L)D$$
$$\text{Standard deviation of demand during } T + L \text{ periods, } \sigma_{T+L} = \sqrt{T + L}\sigma_D$$

The safety inventory in this case is the quantity in excess of D_{T+L} carried by Walmart over the time interval $T + L$. The OUL and the safety inventory ss are related as follows:

$$OUL = D_{T+L} + ss \tag{12.18}$$

Given the desired CSL, the safety inventory (ss) required is given by

$$ss = F_S^{-1}(CSL) \times \sigma_{T+L} = NORMSINV(CSL) \times \sigma_{T+L} \tag{12.19}$$

The average lot size equals the average demand during the review period T and is given as

$$\text{Average lot size, } Q = D_T = D \times T \tag{12.20}$$

In Figure 12-7, we show the inventory profile for a periodic review policy with lead time $L = 4$ and reorder interval $T = 7$. On day 7, the company places an order that determines available inventory until day 18 (as illustrated in Figure 12-7 by the dashed line from point 1 to point 2). As a result, the safety inventory must be sufficient to buffer demand variability over $T + L = 7 + 4 = 11$ days.

We illustrate the periodic review policy for Walmart in Example 12-13 (see worksheet *Example 12-13*).

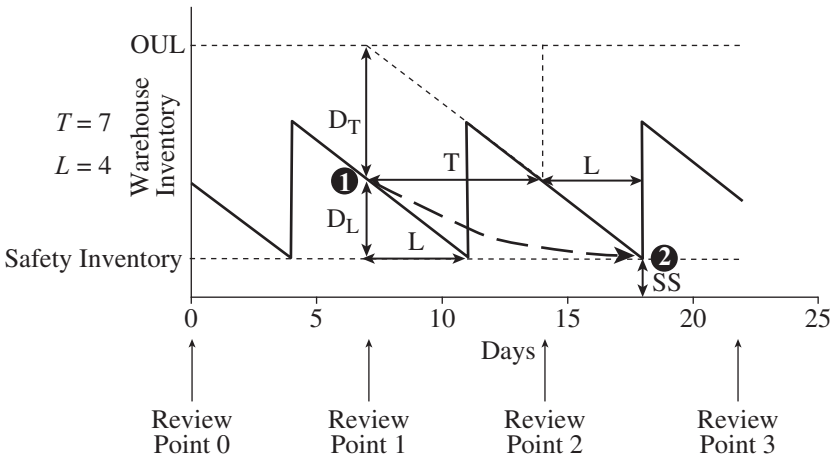

FIGURE 12-7 Inventory Profile for Periodic Review Policy with $L = 4$, $T = 7$

EXAMPLE 12-13 Evaluation Safety Inventory for a Periodic Review Policy

Weekly demand for Legos at a Walmart store is normally distributed, with a mean of 2,500 boxes and a standard deviation of 500. The replenishment lead time is two weeks, and the store manager has decided to review inventory every four weeks. Assuming a periodic-review replenishment policy, evaluate the safety inventory that the store should carry to provide a CSL of 90 percent. Evaluate the OUL for such a policy.

Analysis:

In this case, we have

Average demand per period, $D = 2{,}500$

Standard deviation of demand per period, $\sigma_D = 500$

Average lead time for replenishment, $L = 2$ weeks

Review interval, $T = 4$ weeks

We first obtain the distribution of demand during the time interval $T + L$. Using Equation 12.2, demand during the time interval $T + L$ is normally distributed, with

$$\text{Mean demand during } T + L \text{ periods, } D_{T+L} = (T + L)D = (2 + 4)2{,}500 = 15{,}000$$

$$\text{Standard deviation of demand during } T + L \text{ periods, } \sigma_{T+L} = \sqrt{T + L}\sigma_D = (\sqrt{4 + 2})500 = 1{,}225$$

From Equation 12.19, the required safety inventory for a $CSL = 0.90$ is given as

$$ss = F_S^{-1}(CSL) \times \sigma_{T+L} = NORMSINV(CSL) \times \sigma_{T+L}$$
$$= NORMSINV(0.90) \times 1{,}225 = 1{,}570 \text{ boxes}$$

Using Equation 12.18, the OUL is given by

$$OUL = D_{T+L} + ss = 15{,}000 + 1{,}570 = 16{,}570$$

The store manager thus orders the difference between 16,570 and current inventory every four weeks.

We can now compare the safety inventory required when using continuous and periodic review policies. With a continuous review policy, the safety inventory is used to cover for demand uncertainty over the lead time L. With a periodic review policy, the safety inventory is used to cover for demand uncertainty over the lead time and the review interval $L + T$. Given that higher uncertainty must be accounted for, periodic review policies require a higher level of safety inventory. This argument can be confirmed by comparing the results in Examples 12-4 and 12-13. For a 90 percent CSL, the store manager requires a safety inventory of 906 boxes when using a continuous review and a safety inventory of 1,570 boxes when using a periodic review.

Key Point

Periodic review replenishment policies require more safety inventory than continuous review policies for the same lead time and level of product availability.

Of course, periodic review policies are somewhat simpler to implement because they do not require continuous tracking of inventory. Given the broad use of bar codes and POS systems as well as the emergence of RFID technology, continuous tracking of all inventories is much

more commonplace today than it was a decade ago. In some instances, companies partition their products based on their value. High-value products are managed using continuous review policies, and low-value products are managed using periodic review policies. This makes sense if the cost of perpetual tracking of inventory is more than the savings in safety inventory that result from switching all products to a continuous review policy.

12.7 MANAGING SAFETY INVENTORY IN A MULTIECHELON SUPPLY CHAIN

In our discussion so far, we have assumed that each stage of the supply chain has a well-defined demand and supply distribution that it uses to set its safety inventory levels. In practice, this is not true for multiechelon supply chains. Consider a simple multiechelon supply chain with a supplier feeding a retailer that sells to the final customer. The retailer needs to know demand as well as supply uncertainty to set safety inventory levels. Supply uncertainty, however, is influenced by the level of safety inventory the supplier chooses to carry. If a retailer order arrives when the supplier has enough inventory, the supply lead time is short. In contrast, if the retailer order arrives when the supplier is out of stock, the replenishment lead time for the retailer increases. Thus, if the supplier increases its level of safety inventory, the retailer can reduce the safety inventory it holds. This implies that the level of safety inventory at all stages in a multiechelon supply chain should be related.

All inventory between a stage and the final customer is called the *echelon inventory.* Echelon inventory at a retailer is just the inventory at the retailer or in the pipeline coming to the retailer. Echelon inventory at a distributor, however, includes inventory at the distributor and all retailers served by the distributor. In a multiechelon setting, ROPs and OULs at any stage should be based on echelon inventory and not local inventory. Thus, a distributor should decide its safety inventory levels based on the level of safety inventory carried by all retailers supplied by it. The more safety inventory retailers carry, the less safety inventory the distributor needs to carry. As retailers decrease the level of safety inventory they carry, the distributor must increase its safety inventory to ensure regular replenishment at the retailers.

If all stages in a supply chain attempt to jointly manage their echelon inventory, the issue of how the inventory is divided among various stages becomes important. Carrying inventory upstream in a supply chain allows for more aggregation and thus reduces the amount of inventory required. Carrying inventory upstream, however, increases the probability that the final customer will have to wait because product is not available at a stage close to him or her. Thus, in a multiechelon supply chain, a decision must be made with regard to the level of safety inventory carried at different stages. If inventory is expensive to hold and customers are willing to tolerate a delay, it is better to increase the amount of safety inventory carried upstream, far from the final customer, to exploit the benefits of aggregation. If inventory is inexpensive to hold and customers are time sensitive, it is better to carry more safety inventory downstream, closer to the final customer.

12.8 THE ROLE OF IT IN INVENTORY MANAGEMENT

Besides the basics of formalizing inventory replenishment procedures for thousands of SKUs, the two most significant contributions of IT systems can be improved inventory visibility and better coordination in the supply chain.

An excellent example of the benefits of improved inventory visibility is Nordstrom, a department store chain in the United States. The company was always very good at managing its inventories (IT systems played an important role here) but had historically separated its online inventories and its store inventories. In September 2009, the company started integrating store inventories onto its website. Customers are now able to access inventory no matter where it is available. If they prefer home delivery, Nordstrom can now use store inventory to serve them. If,

however, they prefer to pick up the item themselves, Nordstrom allows them to reserve it for pickup. The increased inventory visibility allows Nordstrom to serve its online customers better while also drawing more traffic to stores. In 2010, Walmart also added a similar feature, called "Pick Up Today," which allows customers to place orders online and pick them up a few hours later at a retail store. Customers are alerted (typically through a text message) when the order is ready. Redbox uses inventory visibility at each of its vending machines to guide customers to the closest kiosk that has the desired DVD in stock. In each example, the increased visibility provided by IT systems allows the firm to improve product availability to the customer without increasing inventories.

Another area in which improved visibility could play a significant role is locating in-store or in-warehouse inventory. Often, a store or warehouse has inventory available but in the wrong place. The net result is a loss in product availability despite carrying inventory. Good RFID systems have the potential to address this issue. Although there has been limited success using RFID systems at the item level in stores (there has been some success with high-value apparel), there has been success in areas such as warehousing of aircraft spare parts.

IT systems have also played a significant role in better integrating different stages of the supply chain. A classic example is the continuous replenishment program (CRP) set up between Procter and Gamble (P&G) and Walmart that allowed P&G to replenish diaper inventory at Walmart based on the visibility of available inventories and sales at Walmart. This coordination allowed the two firms to improve service levels while reducing inventories. Over time, the program evolved into collaborative planning, forecasting, and replenishment (discussed in greater detail in Chapter 10), which allows better coordination of planning and replenishment across multiple supply chain partners through improved visibility of inventories and sales. Even though each of these programs uses IT as a foundation, it is important to acknowledge that success requires important organizational changes and leadership commitment as discussed in Chapter 10. Good IT systems are a necessary but not a sufficient condition for success.

It is important to recognize that the value of the IT system in each of the cases discussed here is tightly linked to the accuracy of the inventory information. Inaccurate inventory information leads to flawed decisions and could in the worst case create mistrust among supply chain partners attempting to coordinate the decisions and actions. A study by DeHoratius and Raman (2008) found that about 65 percent of the inventory records checked for a retailer were inaccurate. That is, for 65 percent of the records checked, the inventory on hand did not match the inventory showing in the IT system. Without reasonably accurate inventory records, the value provided by an IT system will be limited.

12.9 ESTIMATING AND MANAGING SAFETY INVENTORY IN PRACTICE

1. ***Account for the fact that supply chain demand is lumpy.*** In practice, a manufacturer or distributor does not order one unit at a time but instead often orders in a large lot. Thus, demand observed by different stages of the supply chain tends to be lumpy. Lumpiness adds to the variability of demand. For example, when using a continuous review policy, lumpiness may lead to inventory dropping far below the ROP before a replenishment order is placed. On average, inventory will drop below the ROP by half the average size of an order. The lumpiness can be accounted for in practice by raising the safety inventory suggested by the models discussed earlier by half the average size of an order.

2. ***Adjust inventory policies if demand is seasonal.*** In practice, demand is often seasonal, with the mean and the standard deviation of demand varying by the time of year. Thus, a given ROP or OUL may correspond to ten days of demand during the low-demand season and only two days of demand during the peak demand season. If the lead time is one week, stockouts are certain to occur during the peak season. In the presence of seasonality, it is not appropriate to select an average demand and standard deviation over the year to evaluate fixed ROPs and OULs. Both the mean and the standard deviation of demand must be adjusted by the time of year to

reflect changing demand. Corresponding adjustments in the ROPs, OULs, and safety inventories must be made over the year. It is best to evaluate all inventory parameters such as ROPs and OULs in terms of days of demand. A simple heuristic that keeps days of demand constant over time helps account for seasonality by automatically adjusting the ROP and OUL.

3. ***Use simulation to test inventory policies.*** Given that demand is most likely not normally distributed and may be seasonal, it is a good idea to use a computer simulation to test and adjust inventory policies before they are implemented. The simulation should use a demand pattern that truly reflects actual demand, including any lumpiness as well as seasonality. The inventory policies obtained using the models discussed in the chapter can then be tested and adjusted if needed to obtain the desired service levels. Surprisingly powerful simulations can be built using Excel, as we discuss in Chapter 13. Identifying problems in a simulation can save a lot of time and money compared to facing these problems once the inventory policy is in place.

4. ***Start with a pilot.*** Even a simulation cannot identify all problems that may arise when using an inventory policy. Once an inventory policy has been selected and tested using simulation, it is often a good idea to start implementation with a pilot program of products that are representative of the entire set of products in inventory. By starting with a pilot, many of the problems (both in the inventory policies themselves and in the process of applying the policies) can be solved. Getting these problems solved before the policy is rolled out to all the products can save a lot of time and money.

5. ***Monitor service levels.*** Once an inventory policy has been implemented, it is important that its performance be tracked and monitored. Monitoring is crucial because it allows a supply chain to identify when a policy is not working well and make adjustments before supply chain performance is affected significantly. Monitoring requires not just tracking the inventory levels but also tracking any stockouts that may result. Historically, firms have not tracked stockouts very well, partly because they are difficult to track and partly because of the perception that stockouts affect the customer but not the firm itself. Stockouts can be difficult to measure in a situation such as a supermarket, where the customer simply does not buy the product when it is not on the shelf. However, there are simple ways to estimate stockouts. At a supermarket, the fraction of time that a shelf does not contain a product may be used to estimate the fill rate. Stockouts are in fact easier to estimate online, where the number of clicks on an out-of-stock product can be measured. Given the fraction of clicks that turn into orders and the average size of an order, demand during a stockout can be estimated.

6. ***Focus on reducing safety inventories.*** Given that safety inventory is often a large fraction of the total inventory in a supply chain, the ability to reduce safety inventory without hurting product availability can significantly increase supply chain profitability. This is particularly important in the high-tech industry, in which product life cycles are short. In this chapter, we discussed a variety of managerial levers that can help reduce safety inventories without hurting availability. Supply chain managers must focus continuously on using these levers to reduce safety inventories.

12.10 SUMMARY OF LEARNING OBJECTIVES

1. ***Describe different measures of product availability.*** The three basic measures of product availability are product fill rate, order fill rate, and cycle service level. Product fill rate is the fraction of demand for a product that is filled from inventory. Order fill rate is the fraction of orders that are completely filled. Cycle service level is the fraction of replenishment cycles in which no stockouts occur.

2. ***Understand the role of safety inventory in a supply chain.*** Safety inventory helps a supply chain provide customers with a high level of product availability in spite of supply and demand uncertainty. It is carried just in case demand exceeds the amount forecasted or supply arrives later than expected.