Inventory Management

Introduction

An inventory is a stock or store of goods. A typical firm has about 30 percent of its current assets and as much as 90 percent of its working capital invested in inventory. Because inventories may represent a significant portion of total assets, a reduction of inventories can result in a significant increase in return on investment (ROI), a ratio of profit after taxes to total assets.

A typical firm carries the following different kinds of inventories.

- Raw materials and purchased parts.
- Partially completed goods, called work-in-process (WIP).
- Finished-goods inventories (manufacturing firms) or merchandise (retail stores).
- Replacement parts, tools, and supplies.
- Goods-in-transit to warehouses or customers (pipeline inventory).

Cost Information

Three basic costs are associated with inventory management: holding, transaction (ordering), and shortage costs.

Holding or **carrying costs** relate to physically having items in storage. Costs include interest, insurance, taxes, depreciation, obsolescence, deterioration, spoilage, pilferage, breakage, and warehousing costs (heat, light, rent, security). They also include opportunity costs associated with having funds which could be used elsewhere tied up in inventory. Note that it is the *variable* portion of these costs that is pertinent.

The significance of the various components of holding cost depends on the type of item involved, although taxes, interest, and insurance are generally based on the dollar value of the inventory. Items that are easily concealed (e.g., pocket cameras, transistor radios, calculators) or fairly expensive (cars, TVs) are prone to theft. Fresh seafood, meats and poultry, produce, and baked goods are subject to rapid deterioration and spoilage. Daily products, salad dressings, medications, batteries, and film also have limited shelf lives.

Holding costs are stated in either of two ways: as a percentage of unit price or as a dollar amount per unit. In any case, typical annual holding costs range from 20 percent to 40 percent of the value of an item. In other words, to hold a $100 item for one year could cost from $20 to $40.

Ordering costs are the costs of ordering and receiving inventory. They are the costs that vary with the actual placement of an order. These include determining how much is needed, preparing invoices, shipping costs, inspecting goods upon arrival for quality and quantity, and moving the goods to temporary storage. Ordering costs are generally expressed as a fixed dollar amount per order, regardless of order size.

When a firm produces its own inventory instead of ordering it from a supplier, the costs of machine setup (e.g., preparing equipment for the job by adjusting the machine, changing cutting tools) are analogous to ordering costs; that is, they are expressed as a fixed charge per production run, regardless of the size of the run.

Shortage costs result when demand exceeds the supply of inventory on hand. These costs can include the opportunity cost of not making a sale, loss of customer good will, late charges, and similar costs. Furthermore, if the shortage occurs in an item carried for internal use (e.g., to supply an assembly line), the cost of lost production or downtime is considered a shortage cost. Such cost can easily run into hundreds of dollars a minute or more. Shortage costs are sometimes difficult to measure, and they may be subjectively estimated.

Classification System

The **A-B-C approach** classifies inventory items according to some measure of importance, usually annual dollar usage (i.e., dollar value per unit multiplied by annual usage rate), and then allocates control efforts accordingly. Typically, three classes of items are used: A (very important), B (moderately important), and C (least important).

The actual number of categories varies from organization to organization, depending on the extent to which a firm wants to differentiate control efforts. With three classes of items, A items generally account for about 15 to 20 percent of the number of items in inventory but about 60 to 70 percent of the dollar usage. At the other end of the scale, C items might account for about 60 percent of the number of items but only about 10 percent of the dollar usage of an inventory.

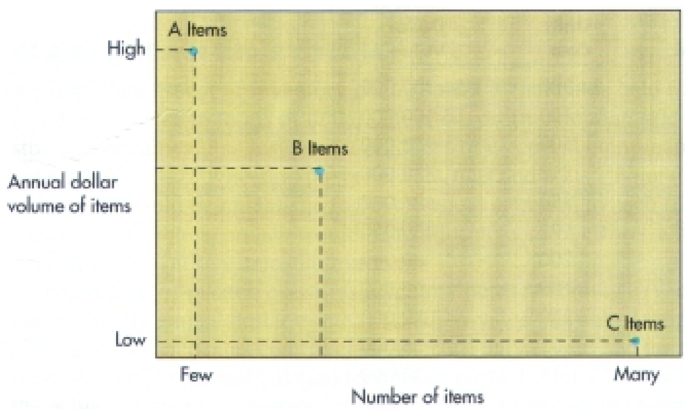

A items should receive close attention through frequent reviews of amounts on hand and control over withdraws, where possible, to make sure that customer service levels are attained. The C items should receive only loose control (two-bin systems, bulk orders), and the B items should have controls that lie between the two extremes.

Note that C items are not necessarily unimportant; incurring a stockout of C items such as the nuts and bolts used to assemble manufactured goods can result in a costly shutdown of an assembly line. However, due to the low annual dollar volume of C items, there may not be much additional cost incurred by ordering larger quantities of some items, or ordering them a bit earlier.

Example: Classify the inventory items as A, B, or C based on annual dollar value, given the following information:

Solution:

Item	Annual Demand	Unit Cost	Annual Dollar Value
1	1,000	$4,300	$4,300,000
2	5,000	720	3,600,000
3	1,900	500	950,000
4	1,000	710	710,000
5	2,500	250	625,000
6	2,500	192	480,000
7	400	200	80,000
8	500	100	50,000
9	200	210	42,000
10	1,000	35	35,000
11	3,000	10	30,000
12	9,000	3	27,000

The first two items have a relatively high annual dollar value, so it seems reasonable to classify them as A items. The next four items appear to have moderate annual dollar value and should be classified as B items. The reminder are C items, based on their relatively low dollar value.

Although annual dollar volume may be the primary factor in classifying inventory items, other useful factors may include the risk of obsolescence, the risk of a stockout, the distance of a supplier, and so on.

In addition to customer service, another application of the A-B-C approach is as a guide to **cycle counting**, which is a physical count of items in inventory. The purpose of cycle counting is to reduce discrepancies between the amounts indicated by inventory records and the actual quantities of inventory on hand.

Key questions concerning cycle counting are:

1. How much accuracy is needed?
2. When should cycle counting be performed?
3. Who should do it?

APICS, formerly known as the American Production and Inventory Control Society, recommends the following guidelines for inventory record accuracy: ± 0.2 percent for A items, ± 1 percent for B items, and ± 5 percent for C items.

Some companies use certain events to trigger cycle counting, whereas others do it on periodic (scheduled) basis. Events that can trigger a physical count of inventory include a out-of-stock report written on an item indicated by inventory records to be in stock, an inventory report that indicates a low or zero balance of an item, and a specified level of activity (e.g., every 2,000 units sold).

Some companies use regular stockroom personnel to do cycle counting during periods of slow activity while others contract with outside firms to do it on a periodic basis. Use of an outside firm provides an independent check on inventory and may reduce the risk of problems created by dishonest employees. Still other firms maintain full time personnel to do cycle counting.

How Much To Order: Economic Order Quantity Model

Economic order quantity (**EOQ**) **models** identify the optimal order quantity by minimizing the sum of certain annual costs that vary with order size. These models are:

1. Economic order quantity model.
2. Economic production quantity model.
3. Quantity discount model.

Economic Order Quantity (EOQ) Model

The basic economic order quantity (EOQ) model identifies the order size that will minimize the sum of the annual costs of holding inventory and ordering inventory. The unit cost of inventory items is not generally included, because the unit cost is unaffected by the order size. If holding costs are specified as a percentage of unit cost, then unit cost is indirectly included in the total cost as a part of holding costs.

Assumptions of the basic EOQ model are:

1. Only one product is involved.
2. Annual demand requirements are known.
3. Demand is spread evenly throughout the year so that the demand rate is reasonable constant.
4. Lead time does not vary.
5. Each order is received in a single delivery.
6. There are no quantity discounts.

Inventory ordering and usage occur in cycles. A cycle begins with receipt of an order of Q units, which are withdrawn at a constant rate over time. When the quantity on hand is just sufficient to satisfy demand during lead time, an order for Q units is submitted to the supplier. Because it is assumed that both the usage rate and the lead time do not vary, the order will be received at the precise instant that the inventory on hand falls to zero. Thus, orders are timed to avoid both excess stock and stockout (i.e., running out of stock).

Annual carrying cost is computed by multiplying the average amount of inventory on hand by the cost to carry one unit for one year, even though any given unit would not necessarily be held for a year. The average inventory is simply half of the order quantity: the amount on hand decreases steadily from Q units to 0, for an average of (Q+0)/2, or Q/2. Using the symbol H to represent the average annual carrying cost per unit, the total annual carrying cost is

$$\text{annual carrying cost} = \frac{Q}{2}H$$

where

Q = order quantity in units
H = holding (carrying) cost per unit

Carrying cost is thus a linear function of Q: carrying costs increase or decrease in direct proportion to changes in the order quantity Q.

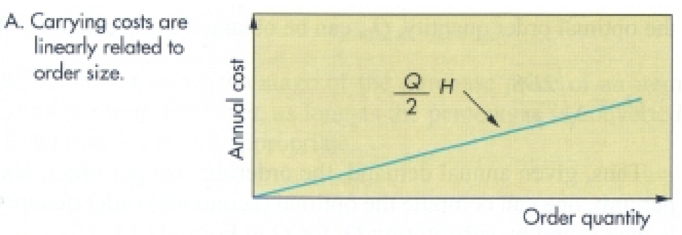

On the other hand, annual ordering cost will decrease as order size increases, because for a given annual demand, the larger the order size, the fewer the number of orders needed. In general, the number of orders per year will be D/Q, where D = annual demand and Q = order size. Unlike carrying costs, ordering costs are relatively insensitive to order size; regardless of the amount of an order, certain amount of activities must be done, such as determine how much is needed, periodically evaluate sources of supply, and prepare the invoice. Even inspection of the shipment to verify quality and quantity characteristics is not strongly influenced by order size, since large shipments are sampled rather than completely inspected. Hence, there is a fixed ordering cost. *Annual ordering cost* is a function of the number of orders per year and the ordering cost per order:

$$\text{annual ordering cost} = \frac{D}{Q}S$$

where

D = demand, usually in units per year
S = ordering cost

Because the numbers of orders per year, D/Q, decreases as Q increases, annual ordering cost is inversely related to order size.

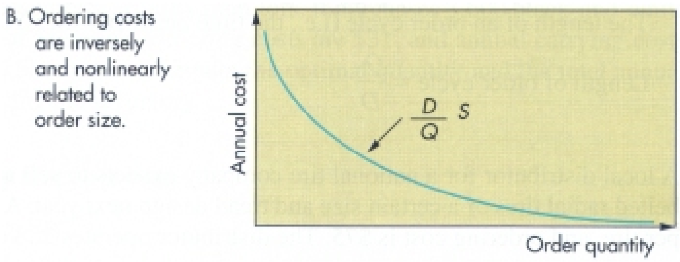

The total annual cost associated with carrying and ordering inventory when Q units are ordered each time is

$$TC = \begin{matrix}\text{annual}\\ \text{carrying}\\ \text{cost}\end{matrix} + \begin{matrix}\text{annual}\\ \text{ordering}\\ \text{cost}\end{matrix} = \frac{Q}{2}H + \frac{D}{Q}S$$

Note that D and H must be in the same units, e.g., months, years. The figure below reveals that the total cost curve is U-shape (i.e., convex, with one minimum) and that *it reaches its minimum at the quantity where carrying and ordering costs are equal.*

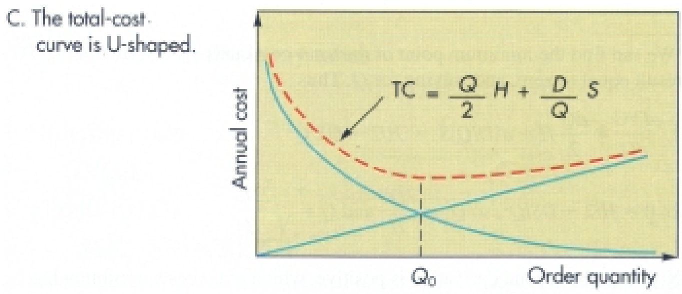

An expression for the optimal order quantity, Q_0, can be obtained using calculus. The result of the formula is

$$Q_0 = \sqrt{\frac{2DS}{H}}$$

Thus, given annual demand, the ordering cost per order, and the annual carrying cost per unit, one can compute the optimal (economic) order quantity. The minimum total cost is then found by substituting Q_0 for Q in the total cost formula above.

The length of an order cycle (i.e., the time between orders) is

$$\text{length of order cycle} = \frac{Q_0}{D}$$

Example: A local distributor for a national tire company expects to sell approximately 9,600 steel belted radial tires of a certain size and tread design next year. Annual carrying cost is $16 per tire, and ordering cost is $75. The distributor operates 288 days a year.

a. What is the EOQ?
b. How many times per year does the store reorder?
c. What is the length of an order cycle?
d. What is the total annual cost if the EOQ quantity is ordered?

Solution:

D = 9,600 tires per year
H = $16 per unit per year
S = $75

a. $$Q_0 = \sqrt{\frac{2DS}{H}} = \sqrt{\frac{2(9{,}600)75}{16}} = 300 \text{ tires}$$

b. $$\text{Number of orders per year}: \frac{D}{Q_0} = \frac{9{,}600 \text{ tires}}{300 \text{ tires}} = 32$$

c. $$\text{Length of order cycle}: \frac{Q_0}{D} = \frac{300 \text{ tires}}{9{,}600 \text{ tires}} = \frac{1}{32} \text{ of a year} \Rightarrow \frac{1}{32} \times 288 = 9 \text{ workdays}$$

d. $$\begin{aligned} TC &= \text{carrying cost} + \text{ordering cost} \\ &= \frac{Q_0}{2}H + \frac{D}{Q_0}S \\ &= \frac{300}{2}\$16 + \frac{9{,}600}{300}\$75 \\ &= \$2{,}400 + \$2{,}400 \\ &= \$4{,}800 \end{aligned}$$

Note that the ordering and carrying costs are equal at the EOQ, as illustrated previously.

Carrying cost is sometimes stated as a percentage of the purchase price of an item rather than as a dollar amount per unit. However, as long as the percentage is converted into a dollar amount, the EOQ formula is still appropriate.

Example: Piddling Manufacturing assembles security monitors. It purchases 3,600 black-and-white cathode ray tubes a year at $65 each. Ordering costs are $31, and annual carrying costs are 20 percent of the purchase price. Compute the optimal quantity and the total annual cost of ordering and carrying the inventory.

Solution:

$D = 3{,}600$ cathode ray tubes per year

$S = \$31$

$H = .20(\$65) = \$13$

$$Q_0 = \sqrt{\frac{2DS}{H}} = \sqrt{\frac{2(3{,}600)(31)}{13}} = 131 \text{ cathode ray tubes}$$

$$TC = \text{carrying costs} + \text{ordering costs}$$
$$= \frac{Q_0}{2}H + \frac{D}{Q_0}S$$
$$= \frac{131}{2}13 + \frac{3{,}600}{131}31$$
$$= \$852 + \$852 = \$1{,}704$$

Holding and ordering costs, and annual demand, are typically estimated values rather than values that can be precisely determined, say, from accounting records. Holding costs are sometimes *designated* rather than computed by managers. Consequently, the EOQ should be regarded as an *approximate* quantity rather than an exact quantity. Thus, rounding the calculated value is perfectly acceptable; stating a value to several decimal places would tend to give an unrealistic impression of the impression involved. An obvious question is: How good is this "approximate" EOQ in terms of minimizing cost? The answer is that the EOQ is fairly robust; the total cost curve is relatively flat near the EOQ, especially to the right of the EOQ. In other words, even if the resulting EOQ differs from the actual EOQ, total costs will not increase much as all. This is particularly true for quantities larger than the real EOQ, because the total cost curve rises very slowly to the right of the EOQ.

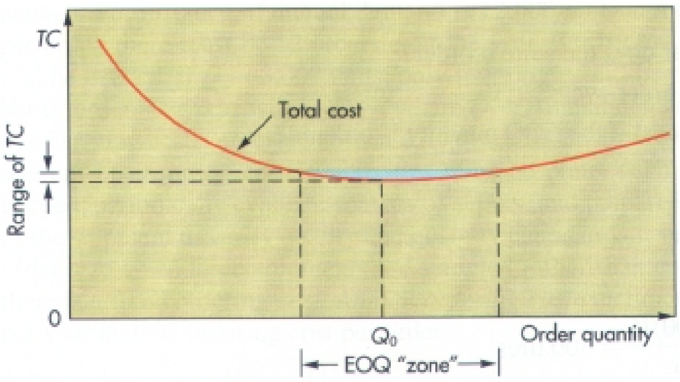

Economic Production Quantity (EPQ) Model

In certain instances, the capacity to produce a single part exceeds the part's usage or demand rate. As long as production continues, inventory will continue to grow. In such instances, it makes sense to periodically produce such items in batches, or *lots*.

The assumptions of the EPQ model are similar to those of the EOQ model, except that instead of orders received in a single delivery, units are received incrementally during production. The assumptions are:

1. Only one item is involved.
2. Annual demand is known.
3. The usage rate is constant.
4. Usage occurs continuously, but production occurs periodically.
5. The production rate is constant.
6. Lead time does not vary.
7. There are no quantity discounts.

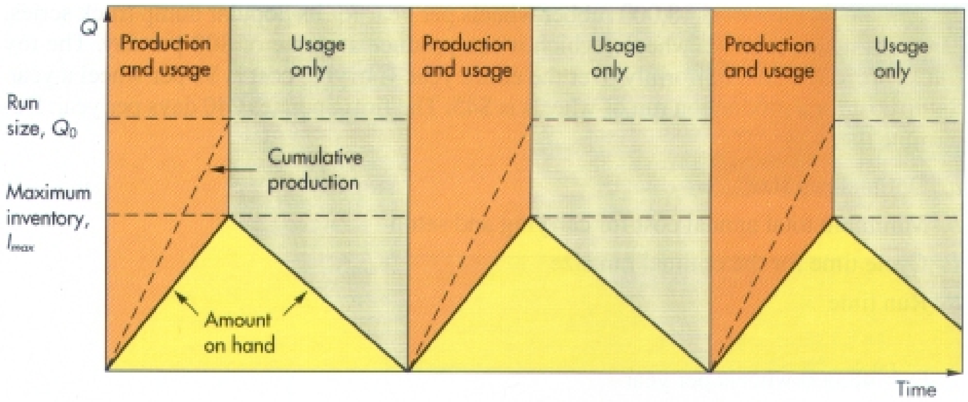

During the production phase of the cycle, inventory builds up at a rate equal to the difference between production and usage rates. As long as production occurs, the inventory will continue to build; when production ceases, the inventory level will begin to decrease. Hence, the inventory level will be maximum at the point where production ceases. When the amount of inventory on hand is exhausted, production is resumed, and the cycle repeats itself.

Because the company makes the product itself, there are no ordering costs as such. Nonetheless, with every production run (batch), there are setup costs --- the costs required to prepare the equipment for the job, such as cleaning, adjusting, and changing tools and fixtures. Setup costs are analogous to ordering costs because they are independent of the lot (run) size. They are treated in the formula in exactly the same way. The larger the run size, the fewer the number of runs needed and, hence, the lower the annual setup cost.

$$TC_{min} = \text{carrying cost} + \text{setup cost} = \frac{I_{max}}{2}H + \frac{D}{Q_0}S$$

where

I_{max} = maximum inventory.

The economic run quantity is

$$Q_0 = \sqrt{\frac{2DS}{H}}\sqrt{\frac{p}{p-u}}$$

where

p = production or delivery rate
u = usage rate

The cycle time (the time between orders or between the beginning of runs) for the economic run size model is a function of the run size and usage (demand) rate:

$$\text{cycle time} = \frac{Q_0}{u}$$

Similarly, the run time (the production phase of the cycle) is a function of the run size and the production rate:

$$\text{production rate} = \frac{Q_0}{p}$$

The maximum and average inventory levels are

$$I_{max} = \frac{Q_0}{p}(p-u) \quad \text{and} \quad I_{average} = \frac{I_{max}}{2}$$

Example: A toy manufacturer uses 48,000 rubber wheels per year for its popular dump truck series. The firm makes its own wheels, which it can produce at a rate of 800 per day. The toy trucks are assembled uniformly over the entire year. Carrying cost is $1 per wheel a year. Setup cost for a production run of wheels is $45. The firm operates 240 days per year. Determine the:

a. Optimal run size
b. Minimum total annual cost for carrying and setup
c. Cycle time for the optimal run size
d. Run time

Solution:

D = 48,000 wheels per year
S = $45
H = $1 per wheel per year
p = 800 wheels per day
u = 48,000 wheels per 240 days, or 200 wheels per day

a.
$$Q_0 = \sqrt{\frac{2DS}{H}}\sqrt{\frac{p}{p-u}} = \sqrt{\frac{2(48{,}000)45}{1}}\sqrt{\frac{800}{800-200}} = 2{,}400 \text{ wheels}$$

$$TC_{\min} = \text{carrying cost} + \text{setup cost} = \frac{I_{\max}}{2}H + \frac{D}{Q_0}S$$

$$I_{\max} = \frac{Q_0}{p}(p-u) = \frac{2{,}400}{800}(800-200) = 1{,}800 \text{ wheels}$$

b.
$$TC = \frac{1{,}800}{2}\times \$1 + \frac{48{,}000}{2{,}400}\times \$45 = \$900 + \$900 = \$1{,}800$$

$$\text{Cycle time} = \frac{Q_0}{u} = \frac{2{,}400 \text{ wheels}}{200 \text{ wheels per day}} = 12 \text{ days}$$

c. Thus, a run of wheels will be made every 12 days.

$$\text{Run time} = \frac{Q_0}{p} = \frac{2{,}400 \text{ wheels}}{800 \text{ wheels per day}} = 3 \text{ days}$$

d. Thus, each run will require 3 days to complete.

Quantity Discount Model

Quantity discounts are price reduction for large orders offered to customers to include them to buy in large quantities.

If quantity discounts are offered, the buyer must weight the potential benefits of reduced purchase price and fewer orders that will result from buying in large quantities against the increase in carrying costs caused by higher average inventories. The buyer's goal with quantity discounts is to select the order quantity that will minimize total cost, where total cost is the sum of carrying cost, ordering cost, and purchasing cost:

$$TC = \text{carrying cost} + \text{ordering cost} + \text{purchasing cost}$$
$$= \frac{Q}{2}H + \frac{D}{Q}S + PD$$

where P = unit price.

Recall that in the basic EOQ model, determination of order size does not involve the purchasing cost. The rationale for not including unit price is that the assumption of no quantity discounts, price per unit is the same for all order sizes (when the demand, D, is known and a constant). Inclusion of unit price in the total cost computation is that case would merely increase the total cost by the amount PD, P times D. A graph of total annual purchase cost versus quantity would be a horizontal line. Hence, including purchasing costs would merely raise the total-cost curve by the same amount (PD) at every point. That would not change the EOQ curve.

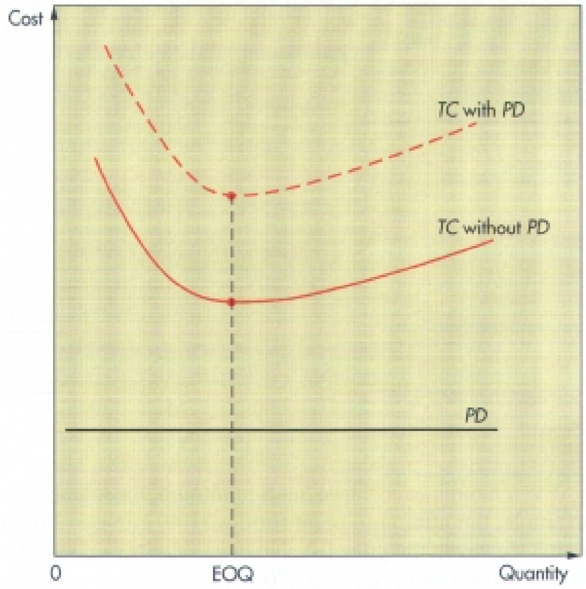

When quantity discounts are offered, there is a separate U-shape total-cost curve for each unit price. Again, including unit price merely raises each curve by a constant amount. However, because the unit prices are all different, each curve is raised by a different amount: smaller unit price will raise a total-cost curve less than larger unit price. Note that no one curve applies to the entire range of quantities; each curve applies to only a *portion* of the range.

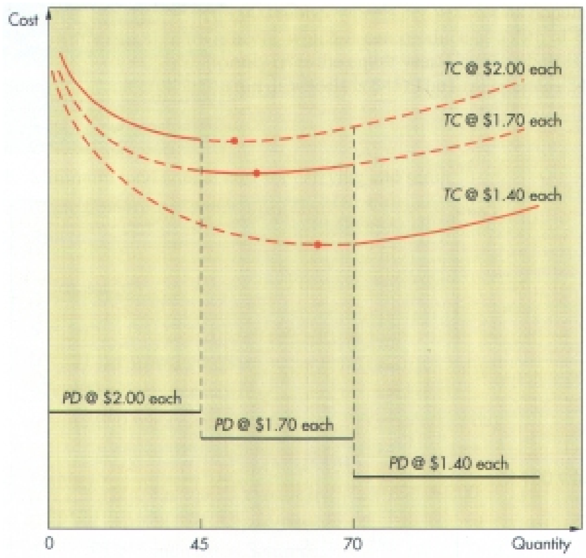

Hence, the applicable or feasible total cost is initially on the curve with the highest unit price and then drops down, curve by curve, at the *price breaks*, which are the minimum quantities needed to obtain the discounts. The result is a total-cost curve with steps at price breaks.

Even though each curve has a minimum, those curves are not necessarily feasible (if we consider the quantity discount rages). The actual total-cost curve is denoted by the solid lines; only those price-quantity combinations are possible. The objective of the quantity discount model is to identify an order quantity that will represent the lowest cost for the entire set of curves.

There are two general cases of the model. In one, carrying costs are constant; in the other, carrying costs are stated as a percentage of purchase price.

When carrying costs are constant, there will be a single minimum point: all curves will have their minimum point at the same quantity. Consequently, the total-cost curves line up vertically, differing only in that the lower unit prices are reflected by lower total-cost curves.

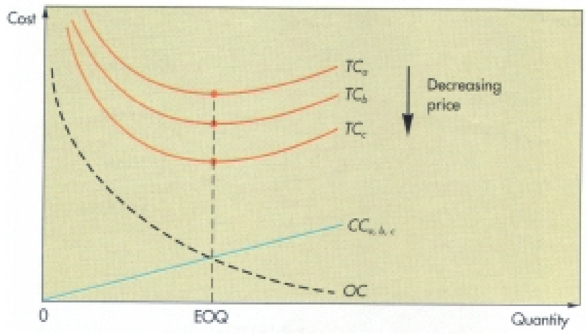

When carrying costs are specified as a percentage of unit price, each curve will have a different minimum point. Because carrying costs are a percentage of price, lower prices will mean lower carrying costs and larger minimum points. Thus, as price decreases, each curve's minimum point will be to the right of the next higher curve's minimum point.

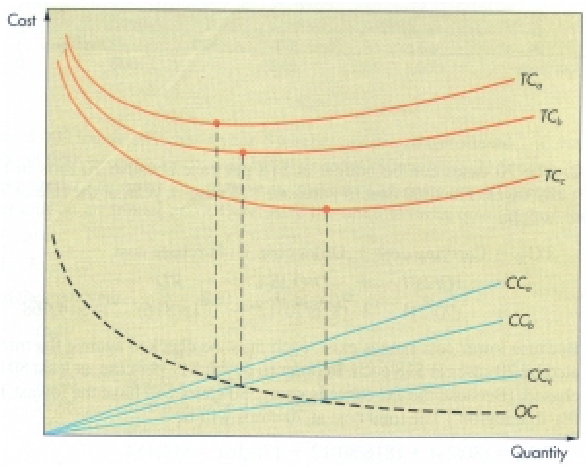

The procedure for determining the overall EOQ differs slightly, depending on which of these two cases is relevant. For carrying costs that are constant, the procedure is as follows:

1. Compute the common minimum point.
2. Only one of the unit prices will have the minimum point in its feasible rage since the ranges do not overlap. Identify that range.
    a. If the feasible minimum point is on the lowest price range, that is the optimal order quantity.
    b. If the feasible minimum point is in any other range, compute the total cost for the minimum point and for the price breaks of all lower unit costs. Compare the total costs; the quantity (minimum point or price break) that yields the lowest total cost is the optimal order quantity.

Example: The maintenance department of a large hospital uses about 816 cases of liquid cleanser annually. Ordering costs are $12, carrying costs are $4 per case per year, and the new price schedule indicates that orders of less than 50 cases will cost $20 per case, 50 to 79 cases will cost $18 per case, 80 to 99 cases will cost $17 per case, and larger orders will cost $16 per case. Determine the optimal order quantity and the total cost.

Solution: D = 816 cases per year, S = $12, H = $4 per case per year.

Range	Price
1 to 49	$20
50 to 79	18
80 to 99	17
100 or more	16

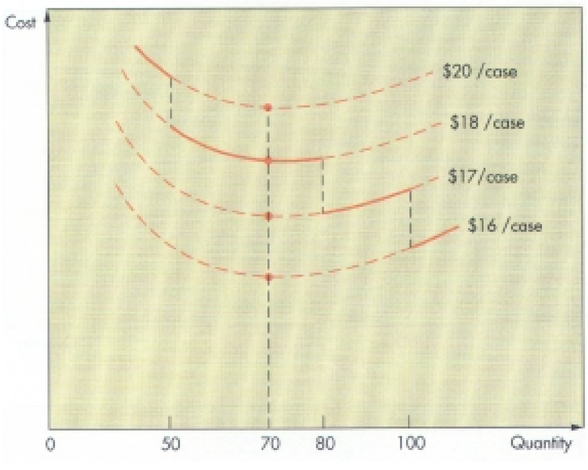

1. Compute the common EOQ: $\sqrt{\frac{2DS}{H}} = \sqrt{\frac{2(816)12}{4}} = 70$ cases
2. The 70 cases can be bought at $18 per case because 70 falls in the range of 50 to 79 cases. The total cost to purchase 816 cases a year, at the rate of 70 cases per order, will be

$$TC_{70} = \text{carrying cost} + \text{ordering cost} + \text{purchase cost}$$
$$= \frac{Q}{2}H + \frac{D}{Q_0}S + PD$$
$$= \frac{70}{2}4 + \frac{816}{70}12 + 18(816) = \$14{,}968$$

Because lower cost ranges exist, each must be checked against the minimum cost generated by 70 cases at $18 each. In order to buy at $17 per case, at least 80 cases must be purchased. (Because the TC curve is rising, 80 cases will have the lowest TC for that curve's feasible region.) The total cost at 80 cases will be

$$TC_{80} = (80/2)4 + (816/80)12 + 17(816) = \$14{,}154$$

To obtain a cost of $16 per case, at least 100 cases per order are required, and the total cost will be

$$TC_{100} = (100/2)4 + (816/100)12 + 16(816) = \$13{,}354$$

Therefore, because 100 cases per order yields the lowest total cost, 100 cases is the overall optimal order quantity.

When carrying costs are expressed as a percentage of price, determine the best purchase quantity with the following procedure:

1. Beginning with the lowest unit price, compute the minimum points for each price range until you find a feasible minimum point (i.e., until a minimum point falls in the quantity range for its price).
2. If the minimum point for the lowest unit price is feasible, it is the optimal order quantity. If the minimum point is not feasible in the lowest price range, compare the total cost at the price break for all *lower* prices with the total cost of the largest feasible minimum point. The quantity that yields the lowest total cost is the optimum.

Example: Surge Electric uses 4,000 toggle switches a year. Switches are priced as follows: 1 to 499, 90 cents each; 500 to 999, 85 cents each; and 1,000 or more, 80 cents each It costs approximately $30 to prepare an order and receive it, and carrying costs are 40 percent of purchase price per unit on an annual basis. Determine the optimal order quantity and the total annual cost.

Solution: D = 4,000 switches per year, S = $30, H = .40P.

Range	Unit Price	H
1 to 499	$0.90	$0.36
500 to 999	$0.85	$0.34
1,000 or more	$0.80	$0.32

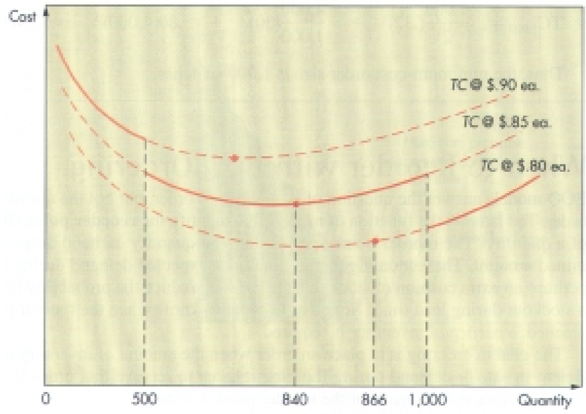

Find the minimum point for each price, starting with the lowest price, until you locate a feasible minimum point.

$$\text{minimum point}_{0.80} = \sqrt{\frac{2DS}{H}} = \sqrt{\frac{2(4{,}000)30}{0.32}} = 866 \text{ switches}$$

Because an order size of 866 switches will cost $0.85 each rather than $0.80 each, 866 is not a feasible minimum point for $0.80 per switch. Next, try $0.85 per unit.

$$\text{minimum point}_{0.85} = \sqrt{\frac{2DS}{H}} = \sqrt{\frac{2(4{,}000)30}{0.34}} = 840 \text{ switches}$$

This is feasible; it falls in the $0.85 per switch range of 500 to 999.

Now compute the total cost for 840, and compare it to the total cost of the minimum quantity necessary to obtain a price of $0.80 per switch.

$$\begin{aligned} TC_{840} &= \text{carrying cost} + \text{ordering cost} + \text{purchasing cost} \\ &= \frac{Q}{2}H + \frac{D}{Q}S + PD \\ &= \frac{840}{2}(.34) + \frac{4{,}000}{840}(30) + 0.85(4{,}000) = \$3{,}686 \end{aligned}$$

$$TC_{1{,}000} = \frac{1{,}000}{2}(.32) + \frac{4{,}000}{1{,}000}(30) + 0.80(4{,}000) = \$3{,}480$$

Thus, the minimum-cost order size is 1,000 switches.

When to Reorder with EOQ Ordering

The **reorder point (ROP)** occurs when the quantity on hand drops to a predetermined amount. That amount generally includes expected demand during lead time and perhaps an extra cushion of stock, which serves to reduce the probability of experiencing a stockout during lead time. Note that in order to know when the reorder point has been reached, a *perpetual* inventory is required.

There are four determinants of the reorder point quantity:

1. The rate of demand (usually based on a forecast).
2. The lead time.
3. The extent of demand and/or lead time variability.
4. The degree of stockout risk acceptable to manager.

If demand and lead time are both constant, the reorder point is simply

$$ROP = d \times LT$$

where

d = demand rate (units per day or week)
LT = lead time in days or weeks

Note: Demand and lead time must have the same time units.

Example: Tingly takes Two-A-Day vitamins, which are delivered to his home by a routeman seven days after an order is called in. At what point should Tingly reorder?

Solution:

Usage	=	2 vitamins a day
Lead time	=	7 days
ROP	=	Usage x Lead time
	=	2 vitamins per day x 7 days
	=	14 vitamins

Thus, Tingly should reorder when 14 vitamin tablets are left.

When variability is present in demand or lead time, it creates the possibility that actual demand will exceed expected demand. Consequently, it becomes necessary to carry additional inventory, called **safety stock**, to reduce the risk of running out of inventory (a stockout) during lead time. The reorder point then increases by the amount of the safety stock:

ROP = expected demand during lead time + safety stock

The following figure illustrates how safety stock can reduce the risk of a stockout during lead time (LT).

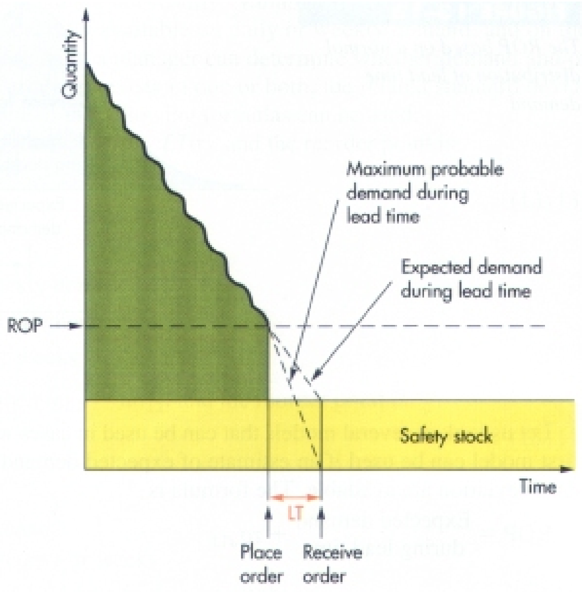

Note that stockout protection is needed only during lead time. If there is a sudden surge at any point during the cycle, that will trigger another order. Once that order is received, the danger of an imminent stockout is negligible.

Because it costs money to hold safety stock, a manager must carefully weight the cost of carrying safety stock against the reduction in stockout risk it provides. The customer *service level* increases as the risk of stockout decreases. Order cycle **service level** can be defined as the probability that demand will not exceed supply during lead time (i.e., that the amount of stock on hand will be sufficient to meet demand). Hence, a service level of 95 percent implies a probability of 95 percent that demand will not exceed supply during lead time. An equivalent statement that demand will be satisfied in 95 percent of such instances does not mean that 95 percent of demand will be satisfied. The risk of a stockout is the complement of service level; a customer service level of 95 percent implies a stockout risk of 5 percent. That is,

service level = 100 percent - stockout risk

The amount of safety stock that is appropriate for a given situation depends on the following factors:

1. The average demand rate and average lead time.
2. Demand and lead time variability.
3. The desired service level.

For a given order cycle service level, the greater the variability in either demand rate or lead time, the greater the amount of safety stock that will be needed to achieve that service level. Similarly, for a given amount of variation in demand rate or lead time, achieving an increase in the service level will require increasing the amount of safety stock. Selection of a service level may reflect stockout costs (e.g., lost sales, customer dissatisfaction) or it might simply be a policy variable (e.g., the manager wants to achieve a specified service level for a certain item).

The following model can be used if an estimate of expected demand during lead time and its standard deviation are available.

$$ROP = \text{expected demand during lead time} + z\sigma_{dLT}$$

where

z = number of standard deviations
σ_{dLT} = the standard deviation of lead time demand

The model generally assumes that any variability in demand rate or lead time can be adequately described by a normal distribution. However, this is not a strict requirement; the models provide approximate reorder points even where actual distributions depart from normal.

The value of z used in a particular instance depends on the stockout risk that the manager is willing to accept. Generally, the smaller the risk the manager is willing to accept, the greater the value of z.

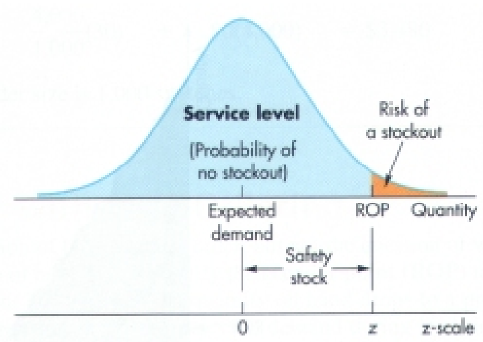

Example: Suppose that the manager of a construction supply house determined from historical records that demand for sand during lead time averages 50 tons. In addition, suppose the manager determined that demand during lead time could be described by a normal distribution that has a mean of 50 tons and a standard deviation of 5 tons. Answer these questions, assuming that the manager is willing to accept a stockout risk of no more than 3 percent:

a. What value of z is appropriate?
b. How much safety stock should be held?
c. What reorder point should be used?

Solution: Expected lead time demand = 50 tons, σ_{dLT} = 5 tons, risk = 3 percent.

a. From the probability table of a standard normal distribution, using a service level of 1 - 0.03 = 0.97, you obtain a value of z = +1.88.
b. Safety stock = $z\sigma_{dLT}$ = 1.88(5) = 9.40 tons.
c. ROP = expected lead time demand + safety stock = 50 + 9.40 = 59.40 tons.

When data on lead time demand are not readily available, the previous formula cannot be used. Nevertheless, data are generally available on daily or weekly demand, and on the length of lead time. Using those data, a manager can determine whether demand and/or lead time is variable, and if variability exists in one or both, the related standard deviation(s). For those situations, one of the following formulas can be used:

1. If only demand is variable, then $\sigma_{dLT} = \sqrt{LT}\sigma_d$, and the reorder point is
   $ROP = \bar{d} \times LT + z\sqrt{LT}\sigma_d$
   where
   $\bar{d}$ = average daily or weekly demand
   σ_d = standard deviation of demand per day or week
   LT = lead time in days or weeks
2. If only lead time is variable, then $\sigma_{dLT} = d\sigma_{LT}$, and the reorder point is
   $ROP = d \times \overline{LT} + zd\sigma_{LT}$
   where
   d = daily or weekly demand
   $\overline{LT}$ = average lead time in days or weeks
   σ_{LT} = standard deviation of lead time in days or weeks
3. If both demand and lead time are variable, then
   $\sigma_{dLT} = \sqrt{\overline{LT}\sigma_d^2 + \bar{d}^2\sigma_{LT}^2}$
   and the reorder point is
   $ROP = \bar{d} \times \overline{LT} + z\sqrt{\overline{LT}\sigma_d^2 + \bar{d}^2\sigma_{LT}^2}$

Note: Each of these models assumes that demand and lead time are independent.

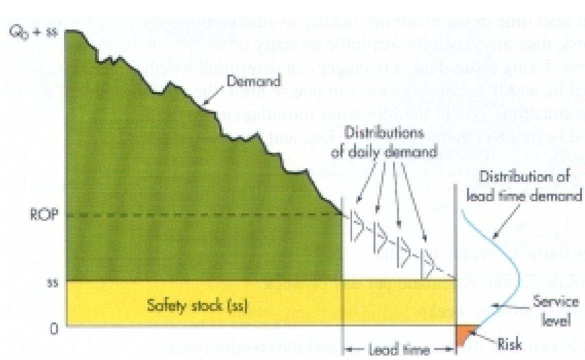

Example: A restaurant uses an average of 50 jars of a special sauce each week. Weekly usage of sauce has a standard deviation of 3 jars. The manager is willing to accept no more than a 10 percent risk of stockout during lead time, which is two weeks. Assume the distribution of usage is normal.

a. Which of the above formula is appropriate for this situation? Why?
b. Determine the value of z.
c. Determine the ROP.

Solution:

$\bar{d} = 50$ jars per week

$\sigma_d = 3$ jars per week

LT = 2 weeks

acceptable risk = 10 percent, so service level is 0.90

a. Because only demand is variable (i.e., has a standard deviation), formula $ROP = \bar{d} \times LT + z\sqrt{LT}\sigma_d$ is appropriate.
b. From the probability of the standard normal distribution, using a service level of 0.90, you obtain z = +1.28.
c. $ROP = \bar{d} \times LT + z\sqrt{LT}\sigma_d = 50 \times 2 + 1.28\sqrt{2}(3) = 100 + 5.43 = 105.43$ jars.

Shortage and Service Levels

The ROP computation does not reveal the expected amount of shortage for a given lead time service level. The expected number of units short can, however, be vary useful to a manager: This quantity can easily be determined from the same information used to compute the ROP, with one additional piece of information (see Table 13-3 on page 569). Use of the table assumes that the distribution of lead time demand can be adequately represented by a normal distribution. If it can, the expected number of units short in each order cycle is given by this formula:

$$E(n) = E(z)\sigma_{dLT}$$

where

$E(n)$ = expected number of units short per order cycle

$E(z)$ = standardized number of units short obtained from Table 13-3

σ_{dLT} = standard deviation of lead time demand

(p.569, Table 13-3 not included here.)

Example: Suppose the standard deviation of lead time demand is known to be 20 units. Lead time demand is approximately normal.

a. For a lead time service level of 90 percent, determine the expected number of units short for an order cycle.
b. What lead time service level would an expected shortage of 2 units imply?

Solution: σ_{dLT} = 20 units.

a. Lead time (cycle) service level = 0.90. From Table 13-3, E(z)=0.048. E(n)=0.048(20 units)=0.96, or about 1 unit.
b. For E(n)=2, we have E(z) = E(n)/σ_{dLT} = 2/20 = 0.10. From Table 13-3, this implies a service level of approximately 81.5 percent.

Having determined the expected number of units short for an order cycle, you can determine the expected number of units short per year. It is simply the expected number of units short per cycle multiplied by the number of cycles (orders) per year. Thus,

$$E(N) = E(n)\frac{D}{Q}$$

where

E(N) = expected number of units short per year.

Example: Given the following information, determine the expected number of units short per year.

D = 1,000   Q = 250   E(n) = 2.5

Solution: $E(N) = E(n)\frac{D}{Q} = 2.5\left(\frac{1,000}{250}\right) = 10.0 \text{ units per year}$

It is sometimes convenient to think of service level in annual terms. One definition of annual service level is the percentage of demand filled directly from inventory. This is also known as the *fill rate*. Thus, if D = 1,000 and 990 units were filled directly from inventory (shortage totaling 10 units over the year were recorded), the annual service level (fill rate) would be 990/1,000 = 99 percent. The annual service level and the lead time service level can be related using the following formula:

$$SL_{annual} = 1 - \frac{E(N)}{D}$$

Using $E(N) = E(n)\frac{D}{Q} = E(z)\sigma_{dLT}\frac{D}{Q}$, we have

$$SL_{annual} = 1 - \frac{E(z)\sigma_{dLT}}{Q}$$

Example: Given a lead time service level of 90, D = 1,000, Q = 250, and σ_{dLT} = 16, determine the annual service level, and the amount of cycle safety stock that would provide an annual service level of 0.98. From Table 13-3, E(z) = 0.048 for a 90 percent lead time service level.

Solution:

a.
$$SL_{annual} = 1 - \frac{E(z)\,\sigma_{dLT}}{Q} = 1 - 0.048\frac{16}{250} = 0.997$$
$$0.98 = 1 - E(z)\frac{16}{250} \Rightarrow E(z) = 0.312 \Rightarrow z \approx 0.19$$

b.
$$\text{Annual safety stock level} = z\sigma_{dLT} = 0.19 \times 16 = 3.04 \approx 3 \text{ units}$$

Note that in the preceding example, a lead time service level of 90 percent provided an annual service level of 99.7 percent. Naturally, different values of D, Q, and σ_{dLT} will tend to produce different results for a cycle service level of 90 percent. Nonetheless, the annual service level will usually be greater than the cycle service level. In addition, since the annual service level as defined relates to the percentage of units short per year. It makes sense to base cycle service levels on a specified annual service level. This means setting annual level and using that value to obtain the service level for the order cycles.

How Much To Order: Fixed-Order-Interval Model

The **fixed-order-interval (FOI) model** is used when orders must be placed at fixed time intervals (weekly, twice a month, etc.) Fixed-interval ordering is widely used in retail businesses (e.g., drugstores, small grocery stores).

Reasons of using the model include: a supplier's policy might encourage orders at fixed intervals; grouping orders for items from the same supplier can produce savings in shipping costs; some situations do not require continuous monitoring of inventory levels, only periodic check is needed.

In the fixed-quantity arrangement, orders are triggered by a quantity (ROP), while in the fixed-interval arrangement orders are triggered by *time*. Therefore, the fixed-interval system must have stockout protection for lead time plus the next order cycle, but the fixed-quantity system needs protection only during lead time. Note, for example, the larger dip into safety stock during the second order cycle with the fixed-interval model below.

Fixed quantity

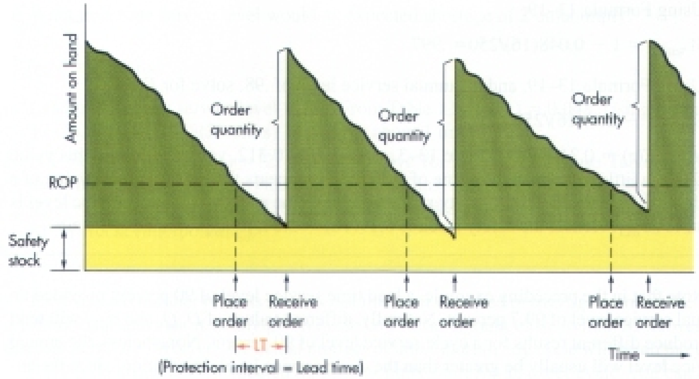

Fixed interval

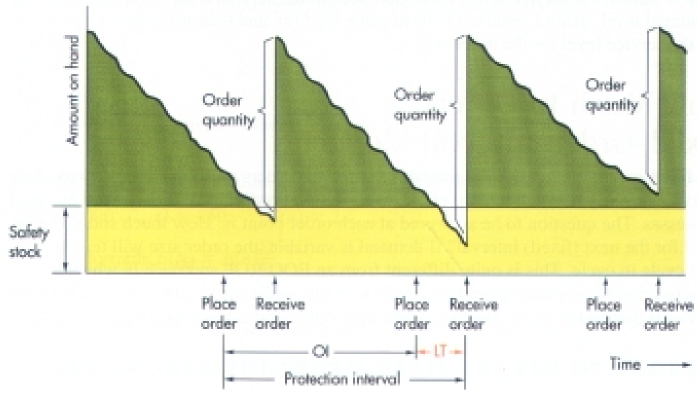

Both models are sensitive to demand experience just priori to ordering, but in somewhat different ways. In the fixed-quantity model, a higher-than-normal demand causes a *shorter time* between orders, whereas in the fixed-interval model, the result is a *larger order size*. Another difference is that the fixed-quantity model requires close monitoring of inventory levels in order to know *when* the amount on hand has reached the reorder point. The fixed-interval model requires only a periodic review (i.e., physical inspection) of inventory levels just priori to placing an order to determine how much is needed.

If both the demand rate and lead time are constant, the fixed-interval model and the fixed-quantity model function identically. Like the ROP model, the fixed-interval model can have variations in demand only, in lead time only, or in both demand and lead time. For the sake of simplicity, only variable demand and constant lead time case is discussed here. Order size of the case is determined by the following formula:

$$\text{amount to order} = \text{expected demand during interval} + \text{safety stock} - \text{amount on hand}$$
$$= \bar{d}(OI + LT) + z\sigma_d\sqrt{OI + LT} - A$$

where

OI = order interval (length of time between orders)

A = amount on hand at reorder time

As in previous models, we assume that demand during the interval is normally distributed.

Example: Given the following information, determine the amount to order.

$$\bar{d} = 30 \text{ units per day}$$
$$\sigma_d = 3 \text{ units per day}$$
$$LT = 2 \text{ days}$$
$$OI = 7 \text{ days}$$
$$\text{desired service level} = 99 \text{ percent}$$
$$\text{amount on hand at reorder time} = 71 \text{ units}$$

Solution: z = 2.33 for 99 percent service level.

$$\text{amount to order} = \bar{d}(OI + LT) + z\sigma_d\sqrt{OI + LT} - A$$
$$= 30(7+2) + 2.33(3)\sqrt{7+2} - 71 = 220 \text{ units}$$

An issue related to fixed-interval ordering is the risk of a stockout. From the perspective (i.e., the point in time) of placing an order, there are two points in the order cycle at which a stockout could occur. One is shortly after the order is placed, while waiting to receive the current order. The second point is near the end of the cycle, while waiting to receive the next order.

To find out the initial risk of a stockout, use $ROP = \bar{d} \times LT + z\sqrt{LT}\sigma_d$, setting ROP equal to the quantity on hand when the order is placed, and solve for z, then obtain the service level for that value of z and subtract it from 1.00 to get the risk of a stockout.

To find out the risk of a stockout at the end of the order cycle, use the formula $\text{amount to order} = \bar{d}(OI + LT) + z\sigma_d\sqrt{OI + LT} - A$ and solve for z. Then, obtain the service level for that value of z and subtract it from 1.00 to get the risk of a stockout.

Example: Given the following information:

$LT = 4$ days	$A = 43$ units
$OI = 12$ days	$Q = 171$ units
$\bar{d} = 10$ units/day	$\sigma_d = 2$ units/day

Determine the risk of stockout at

a. The end of the initial lead time.
b. The end of the second lead time.

Solution:

a. Using formula $ROP = \bar{d} \times LT + z\sqrt{LT}\sigma_d$, we get 43 = 10(4) + z(2)(2). z = 0.75. The service level is 0.7734. The risk is 1-0.7734 = 0.2266, which is fairly high.

b. Using $\text{amount to order} = \overline{d}(OI + LT) + z\sigma_d \sqrt{OI + LT} - A$, we get 171 = 10(4+12) + z(4)(2) - 43. z = 6.75. The service level is virtually 100 percent. The risk is essentially equal to zero.

Benefits and Disadvantages

The fixed-interval system results in the tight control needed for A items in an A-B-C classification due to its periodic reviews. In addition, when two or more items come from the same supplier, grouping orders can yield savings in ordering, packing, and shipping costs. Moreover, it may be the only practical approach if inventory withdraws cannot be closely monitored.

On the negative side, the fixed-interval system necessitates a larger amount of safety stock for a given risk of stockout because of the need to protect against shortages during an entire order interval plus lead time (instead of lead time only), and this increases the carrying cost. Also, there are the costs of periodic reviews.

Single-Period Model

The **single-period model** (sometimes referred to as the *newsboy problem*) is used to handle ordering of perishables (fresh fruits, vegetables, seafood, cut flowers) and items that have a limited useful life (newspapers, magazines, spare parts for specialized equipment). The *period* for spare parts is the life of the equipment, assuming that the parts can not be used for other equipment. What sets unsold or unused goods apart is that they are not typically carried over from one period to the next, at least not without penalty. Day-old baked goods, for instance, are often sold at reduced prices, leftover seafood may be discarded, and out-of-date magazines may be offered to used book stores at bargain rates. There may even be some cost associated with disposal of leftover goods.

Analysis of single-period situations generally focuses on two costs: shortage and excess.

Shortage cost may include a charge for loss of customer goodwill as well as the opportunity cost of lost sales. Generally, **shortage cost** is simply unrealized profit per unit. That is,

$C_{shortage} = C_s$ = revenue per unit - cost per unit

If a shortage or a stockout relates to an item used in production or to a spare part for a machine, then shortage cost refers to the actual cost of lost production.

Excess cost pertains to items left over at the end of the period. In effect, excess cost is the difference between purchase cost and salvage value. That is,

$C_{excess} = C_e$ = original cost per unit - salvage value per unit

If there is cost associated with disposing of excess items, the salvage will be negative and will therefore *increase* the excess cost per unit.

The goal of the single-period model is to identify the order quantity, or stocking level, that will minimize the long-run excess and shortage costs.

There are two general categories of problems that we will consider: those for which demand can be approximated using a continuous distribution (perhaps a theoretical one such as a uniform or normal distribution) and those for which demand can be approximated using a discrete distribution (say, historical frequencies or a theoretical distribution such as the Poisson). The kind of inventory can indicate which type of model might be appropriate. For example, demand for petroleum, liquids, and gases tends to vary over some *continuous scale*, thus lending itself to description by a continuous distribution. Demand for tractors, cars, and computers is expressed in terms of the *number of units* demanded and lends itself to description by a discrete distribution.

Continuous Stocking Levels

When the demand is *uniform*, choosing the stocking level is similar to balancing a seesaw. We have excess cost per unit (C_e) on one end of the distribution and shortage cost per unit (C_s) on the other. The optimal stocking level is analogous to the fulcrum of the seesaw; the stocking level equalizes the cost weights.

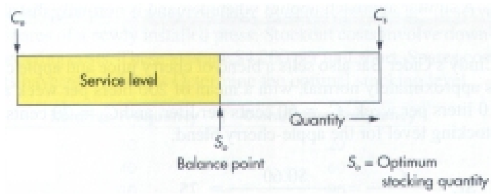

The *service level* is the *probability* that demand will not exceed the stocking level, and computation of the service level is the key to determining the optimal stocking level, S_o.

$$\text{service level} = \frac{C_s}{C_s + C_e}$$

where

C_s = shortage cost per unit

C_e = excess cost per unit

Example: Sweet cider is delivered weekly to Cindy's Cider Bar. Demand varies uniformly between 300 liters and 500 liters per week. Cindy pays 20 cents per liter for the cider and charges 80 cents per liter for it. Unsold cider has no salvage value and cannot be carried over into the next week due to spoilage. Find the optimal stocking level and its stockout risk for that quantity.

Solution:

$C_e = 0.20 - 0 = \$0.20$ per unit

$C_s = 0.80 - 0.20 = \$0.60$ per unit

$$SL = \frac{C_s}{C_s + C_e} = \frac{0.60}{0.60 + 0.20} = 0.75$$

$S_0 = 300 + 0.75(500 - 300) = 450$ liters

The stockout risk is $1.00 - 0.75 = 0.25$

A similar approach applies when demand is normally distributed.

Example: Cindy's Cider Bar also sells a blend of cherry juice and apple cider. Demand for the blend is approximately normal, with a mean of 200 liters per week and a standard deviation of 10 liters per week. C_s = 60 cents per liter, and C_e = 20 cents per liter. Find the optimal stocking level for the apple-cherry blend.

Solution:

$$SL = \frac{C_s}{C_s + C_e} = \frac{0.60}{0.60 + 0.20} = 0.75 \Rightarrow z = 0.675$$

$S_0 = 200 + 0.675(10) = 206.75$ liters

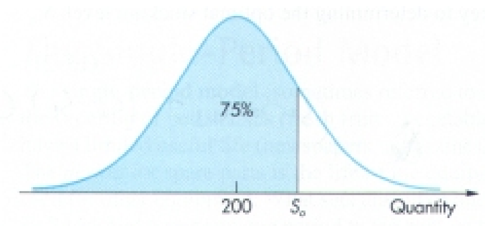

Discrete Stocking Levels

When stocking levels are discrete rather than continuous, the service level computed using the ratio $C_s/(C_s+C_e)$ usually does not coincide with a feasible stocking level (e.g., the optimal amount may be between five and six units). The solution is to stock at the *next higher level* (e.g., six units). In other words, choose the stocking level so that the desired service level is equaled or exceeded.

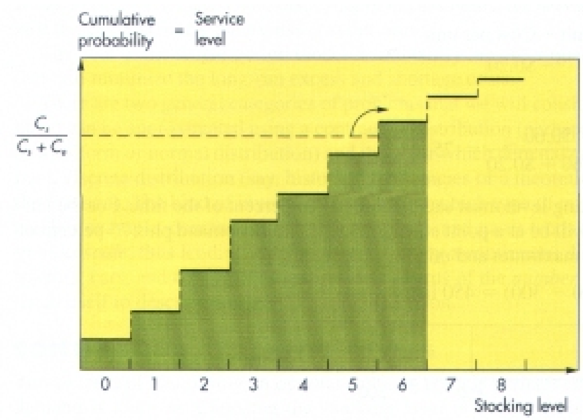

If the computed service level is exactly equal to the cumulative probability associated with one of the stocking levels, there are *two* equivalent stocking levels in terms of minimizing long-run cost --- the one with equal probability and the next higher one.

Example: Historical records on the use of spare parts for several large hydraulic presses are to serve as an estimate of usage for spares of a newly installed press. Stockout costs involve downtown expenses and special ordering costs. These average $4,200 per unit short. Spares cost $800 each, and unused parts have zero salvage. Determine the optimal stocking level.

Number of spares used	Relative frequency	Cumulative frequency
0	0.20	0.20
1	0.40	0.60
2	0.30	0.90
3	0.10	1.00
4 or more	0.00	
	1.00	

Solution:

$C_s = \$4{,}200 \quad C_e = \$800$

$$SL = \frac{C_s}{C_s + C_e} = \frac{4200}{4200+800} = 0.84 \Rightarrow S_0 = 2$$

Example: Demand for long-stemmed red roses at a small flower shop can be approximated using a Poisson distribution that has a mean of four dozen per day. Profit on the roses is $3 per dozen. Leftover flowers are marked down and sold the next day at a loss of $2 per dozen. Assume that all marked-down flowers are sold. What is the optimal stocking level?

Solution:

$C_s = \$3 \quad C_e = \$2$

$$SL = \frac{C_s}{C_s + C_e} = \frac{3}{3+2} = 0.60 \Rightarrow S_0 = 4 \text{ dozens}$$

Demand (dzs per day)	Cumulative frequency
0	0.018
1	0.092
2	0.238
3	0.434
4	0.629
5	0.785
…	…