Source: Corbis

Chapter 12
Inventory planning and control

Introduction

Operations managers often have an ambivalent attitude towards inventories. On the one hand, they are costly, sometimes tying up considerable amounts of working capital. They are also risky because items held in stock could deteriorate, become obsolete or just get lost and, furthermore, they take up valuable space in the operation. On the other hand, they provide some security in an uncertain environment that one can deliver items in stock should customers demand them. This is the dilemma of inventory management: in spite of the cost and the other disadvantages associated with holding stocks, they do facilitate the smoothing of supply and demand. In fact, they exist only because supply and demand are not exactly in harmony with each other (see Figure 12.1).

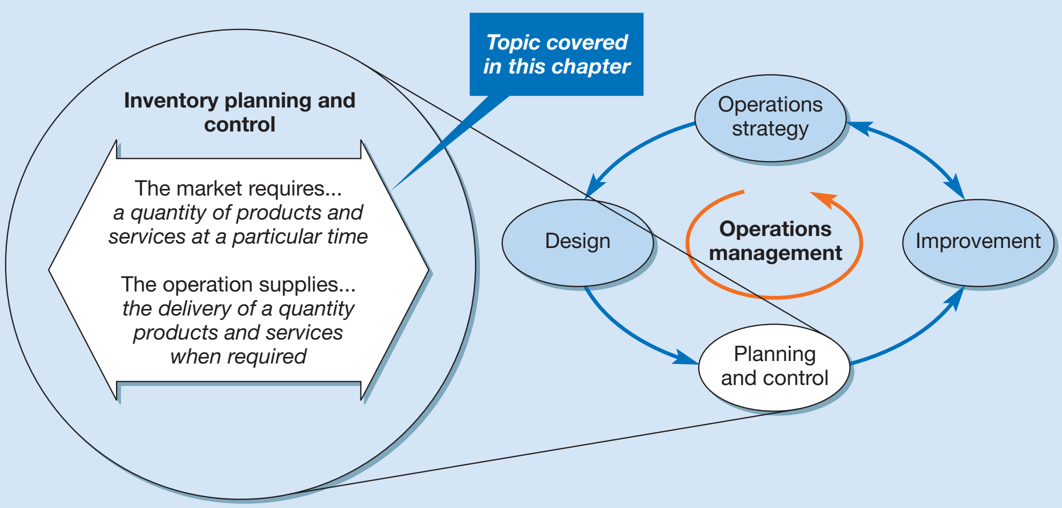

Figure 12.1 This chapter covers inventory planning and control

Key questions ???

- What is inventory?
- Why is inventory necessary?
- What are the disadvantages of holding inventory?
- How much inventory should an operation hold?
- When should an operation replenish its inventory?
- How can inventory be controlled?

Operations in practice

The UK's National Blood Service[1]

No inventory manager likes to run out of stock. But for blood services, such as the UK's National Blood Service (NBS), the consequences of running out of stock can be particularly serious. Many people owe their lives to transfusions that were made possible by the efficient management of blood, stocked in a supply network that stretches from donation centres through to hospital blood banks. The NBS supply chain has three main stages:

- collection, which involves recruiting and retaining blood donors, encouraging them to attend donor sessions (at mobile or fixed locations) and transporting the donated blood to their local blood centre;
- processing, which breaks blood down into its constituent parts (red cells, platelets and plasma) as well as over 20 other blood-based 'products';
- distribution, which transports blood from blood centres to hospitals in response to both routine and emergency requests. Of the Service's 200,000 deliveries a year, about 2500 are emergency deliveries.

GO TO WEB! 12A

Inventory accumulates at all three stages and in individual hospitals' blood banks. Within the supply chain, around 11.5 per cent of donated red blood cells are lost. Much of this is due to losses in processing, but around 5 per cent is not used because it has 'become unavailable', mainly because it has been stored for too long. Part of the Service's inventory control task is to keep this 'time expired' loss to a minimum. In fact, only small losses occur within the NBS; most blood is lost when it is stored in hospital blood banks that are outside its direct control. However, it does attempt to provide advice and support to hospitals to enable them to use blood efficiently.

Blood components and products need to be stored under a variety of conditions, but will deteriorate over time. This varies depending on the component; platelets

Source: Alamy/Van Hilversum

have a shelf life of only five days and demand can fluctuate significantly. This makes stock control particularly difficult. Even red blood cells that have a shelf life of 35 days may not be acceptable to hospitals if they are close to their 'use-by date'. Stock accuracy is crucial. Giving a patient the wrong type of blood can be fatal.

At a local level demand can be affected significantly by accidents. One serious accident involving a cyclist used 750 units of blood, which completely exhausted the available supply (miraculously, he survived). Large-scale accidents usually generate a surge of offers from donors wishing to make immediate donations. There is also a more predictable seasonality to the donating of blood, with a low period during the summer vacation. Yet there is always an unavoidable tension between maintaining sufficient stocks to provide a very high level of supply dependability to hospitals and minimizing wastage. Unless blood stocks are controlled carefully, they can easily go past the 'use-by date' and be wasted. But avoiding outdated blood products is not the only inventory objective at NBS. It also measures the percentage of requests that it was able to meet in full, the percentage emergency requests delivered within two hours, the percentage of units banked to donors bled, the number of new donors enrolled and the number of donors waiting longer than 30 minutes before they were able to donate.

The traceability of donated blood is also increasingly important. Should any problems with a blood product arise, its source can be traced back to the original donor.

What is inventory?

Inventory
Also known as stock, the stored accumulation of transformed resources in a process; usually applies to material resources but may also be used for inventories of information; inventories of customers or customers of customers are usually queues.

Inventory, or 'stock' as it is more commonly called in some countries, is defined here as the *stored accumulation of material resources in a transformation system*. Sometimes the term 'inventory' is also used to describe any capital-transforming resource, such as rooms in a hotel or cars in a vehicle-hire firm, but we will not use that definition here. Usually the term refers only to *transformed resources*. So a manufacturing company will hold stocks of materials, a tax office will hold stocks of information and a theme park will hold stocks of customers. Note that when it is customers who are being processed we normally refer to the 'stocks' of them as 'queues'. This chapter will deal particularly with inventories of materials. Inventories of customers are referred to in Chapter 11. However, this does not imply that this chapter is relevant only when examining predominantly materials-processing operations such as manufacturing operations. All operations keep physical stocks of materials of some sort.

Revisiting operations objectives – the roles of inventory

Most of us are accustomed to keeping inventory for use in our personal lives, but often we don't think about it. For example, most families have some stocks of food and drinks so that they don't have to go out to the shops before every meal. Holding a variety of food ingredients in stock in the kitchen cupboard or freezer gives us the ability to respond quickly (with *speed*) in preparing a meal whenever unexpected guests arrive. It also allows us the *flexibility* to choose a range of menu options without having to go to the time and trouble of purchasing further ingredients. We may purchase some items because we have found something of exceptional *quality* but intend to save it for a special occasion. Many people buy multiple packs to achieve lower *costs* for a wide range of goods from toilet paper to beer, and large packages of shampoo or toothpaste are usually cheaper than small ones. We don't usually intend to use it all up on the day we bought it, but we believe that it saves us time and money to buy larger amounts less frequently. Before going shopping next time, we check our inventory and if a regularly used item is below a certain amount, or occasionally even at zero level, we list it for repurchase. In general, our inventory planning protects us from critical stock-outs, so this approach gives a level of *dependability* of supplies.

It is, however, entirely possible to manage our inventory planning differently. For example, some people (students?) are short of available cash and/or space and so cannot 'invest' in large inventories of goods. They may shop locally for much smaller quantities. They forfeit the cost benefits of bulk-buying but do not have to transport heavy or bulky supplies. They also reduce the risk of forgetting an item in the cupboard and letting it go out of date. Essentially, they purchase against specific known requirements (the next meal). However, they may find that the local shop is temporarily out of stock of a particular item, forcing them, for example, to drink coffee without their usual milk. How we control our own supplies is therefore a matter of choice which can affect their quality (e.g. freshness), availability or speed of response, dependability of supply, flexibility of choice and cost. It is the same for most organizations. Significant levels of inventory can be held for a range of sensible and pragmatic reasons but it must also be tightly controlled for other equally good reasons.

All operations keep inventories

If you walk around any operation you will see several types of stored material. Table 12.1 gives some examples for several operations. However, there are differences between the examples of inventory given in Table 12.1. Some are relatively trivial to the operation in question: for example, the cleaning materials which are stored in the television factory are of much lower value than the stocks of steel, plastic and electronic components which it also holds. More importantly, the television plant would probably not immediately stop if it ran out of cleaning materials, whereas if it ran out of any of its component parts its activities would be severely disrupted. However, cleaning materials would be a far more important item of inventory for an industrial cleaning company, not only because it uses far more of this input but also because its main operation would stop if it ever ran out of them.

The value of inventories

Perhaps the most obvious difference between the operations in Table 12.1 is in the value of the inventories which they hold. In some, it is relatively small compared with the costs of the total inputs to the operation. In others, it will be far higher, especially where storage is the prime purpose of the operation. For example, the value of the goods held in the warehouse is likely to be very high compared with its day-to-day expenditure on such things as labour, rent and running costs. Sometimes the value of the inventories can be so high that it is not even included in the organization's general financial accounts; this would be true, for example, of the precious metals refiner.[2]

Why inventories exist

No matter what is being stored as inventory, or where it is positioned in the operation, it will be there because there is a difference in the timing or rate of supply and demand. If the supply of any item occurred exactly when it was demanded, the item would never be stored. A common analogy is the water tank shown in Figure 12.2. If, over time, the rate of supply of water to the tank differs from the rate at which it is demanded, a tank of water (inventory) will be needed if supply is to be maintained. When the rate of supply exceeds the rate of demand, inventory increases; when the rate of demand exceeds the rate of supply, inventory decreases. So if an operation can match supply and demand rates, it will also succeed in reducing its inventory levels.

Types of inventory

The various reasons for an imbalance between the rates of supply and demand at different points in any operation lead to the different types of inventory. There are five of

Table 12.1 Examples of inventory held in operations

Operation	*Examples of inventory held in operations*
Hotel	Food items, drinks, toilet items, cleaning materials
Hospital	Wound dressings, disposable instruments, whole blood, food, drugs, cleaning materials
Retail store	Goods to be sold, wrapping materials
Warehouse	Goods being stored, packaging materials
Automotive parts distributor	Automotive parts in main depot, automotive parts at local distribution points
Television manufacturer	Components, raw materials, part-finished sub-assemblies, finished televisions, cleaning materials
Precious metals refiner	Material (gold, platinum, etc.) waiting to be processed, material partly processed, fully refined material

GO TO WEB! 12B

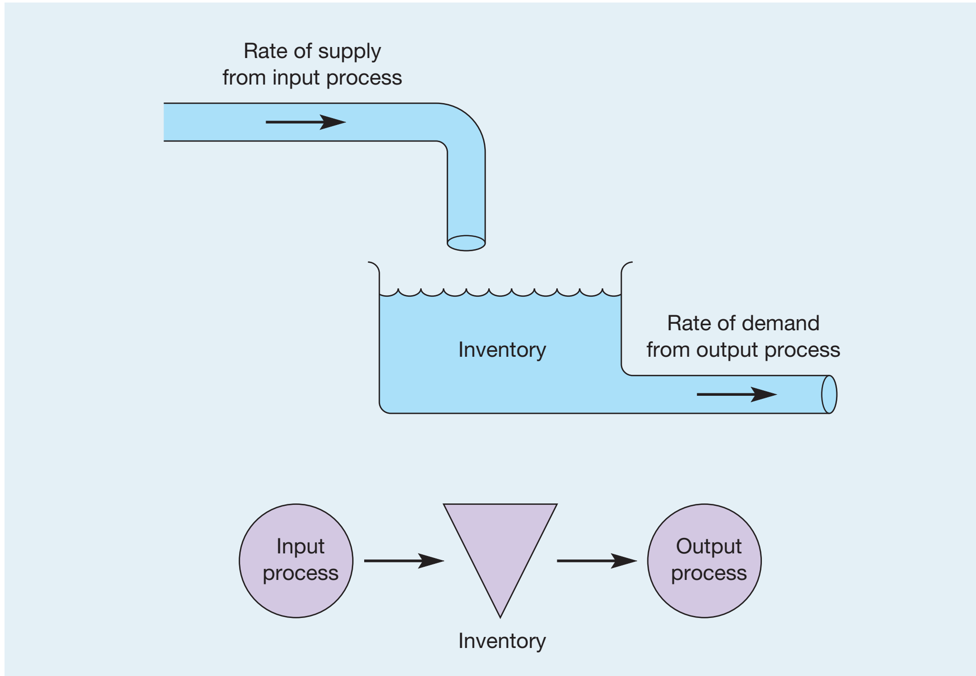

Figure 12.2 Inventory is created to compensate for the differences in timing between supply and demand

these: buffer inventory, cycle inventory, de-coupling inventory, anticipation inventory and pipeline inventory.

Buffer inventory

Buffer inventory
An inventory that compensates for unexpected fluctuations in supply and demand, can also be called safety inventory.

Safety inventory

Buffer inventory is also called safety inventory. Its purpose is to compensate for the unexpected fluctuations in supply and demand. For example, a retail operation can never forecast demand perfectly, even when it has a good idea of the most likely demand level. It will order goods from its suppliers such that there is always a certain amount of most items in stock. This minimum level of inventory is there to cover against the possibility that demand will be greater than expected during the time taken to deliver the goods. This is *buffer* or *safety inventory*. It can also compensate for the uncertainties in the process of the supply of goods into the store, perhaps because of the unreliability of certain suppliers or transport firms.

Cycle inventory

Cycle inventory
Inventory that occurs when one stage in a process cannot supply all the items it produces simultaneously and so has to build up inventory of one item while it processes the others.

Cycle inventory occurs because one or more stages in the process cannot supply all the items it produces simultaneously. For example, suppose a baker makes three types of bread, each of which is equally popular with its customers. Because of the nature of the mixing and baking process, only one kind of bread can be produced at any time. The baker would have to produce each type of bread in batches (batch processes were described in Chapter 4) as shown in Figure 12.3. The batches must be large enough to satisfy the demand for each kind of bread between the times when each batch is ready for sale. So even when demand is steady and predictable, there will always be some inventory to compensate for the intermittent supply of each type of bread. Cycle inventory only results from the need to produce products in batches and the amount of it depends on volume decisions, which are described in a later section of this chapter.

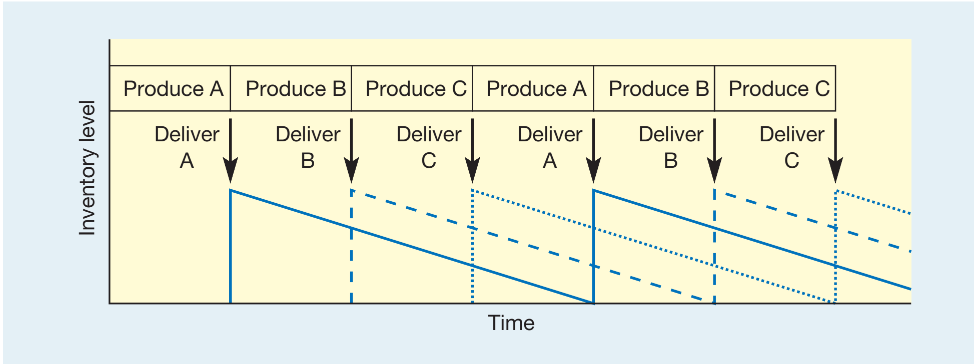

Figure 12.3 **Cycle inventory in a bakery**

De-coupling inventory

Wherever an operation is designed to use a process layout (introduced in Chapter 7), the transformed resources move intermittently between specialized areas or departments that comprise similar operations. Each of these areas can be scheduled to work relatively independently in order to maximize the local utilization and efficiency of the equipment and staff. As a result, each batch of work-in-progress inventory joins a queue, awaiting its turn in the schedule for the next processing stage. This also allows each operation to be set to the optimum processing speed (cycle time), regardless of the speed of the steps before and after. Thus **de-coupling inventory** creates the opportunity for independent scheduling and processing speeds between process stages.

De-coupling inventory
The inventory that is used to allow work centres or processes to operate relatively independently.

Anticipation inventory

In Chapter 11 we saw how anticipation inventory can be used to cope with seasonal demand. Again, it was used to compensate for differences in the timing of supply and demand. Rather than trying to make the product (such as chocolate) only when it was needed, it was produced throughout the year ahead of demand and put into inventory until it was needed. **Anticipation inventory** is most commonly used when demand fluctuations are large but relatively predictable. It might also be used when supply variations are significant, such as in the canning or freezing of seasonal foods.

Anticipation inventory
Inventory that is accumulated to cope with expected future demand or interruptions in supply.

Pipeline inventory

Pipeline inventory
The inventory that exists because material cannot be transported instantaneously.

Pipeline inventory exists because material cannot be transported instantaneously between the point of supply and the point of demand. If a retail store orders a consignment of items from one of its suppliers, the supplier will allocate the stock to the retail store in its own warehouse, pack it, load it onto its truck, transport it to its destination and unload it into the retailer's inventory. From the time that stock is allocated (and therefore it is unavailable to any other customer) to the time it becomes available for the retail store, it is pipeline inventory.

Pipeline inventory also exists within processes where the layout is geographically spread out. For example, a large European manufacturer of specialized steel regularly moves cargoes of part-finished materials between its two mills in the UK and Scandinavia using a dedicated vessel that shuttles between the two countries every week. All the thousands of tonnes of material in transit are pipeline inventory.

Some disadvantages of holding inventory

Although inventory plays an important role in many operations' performance, there are a number of negative aspects of inventory:

- Inventory ties up money, in the form of working capital, which is therefore unavailable for other uses, such as reducing borrowings or making investment in productive fixed assets (we shall expand on the idea of working capital later).
- Inventory incurs storage costs (leasing space, maintaining appropriate conditions, etc.).
- Inventory may become obsolete as alternatives become available.
- Inventory can be damaged or deteriorate.
- Inventory could be lost, or be expensive to retrieve, as it gets hidden among other inventory.
- Inventory might be hazardous to store (for example flammable solvents, explosives, chemicals and drugs), requiring special facilities and systems for safe handling.
- Inventory uses space that could be used to add value.
- Inventory involves administrative and insurance costs.

The position of inventory

Not only are there several reasons for supply–demand imbalance, there could also be several points where such imbalance exists between different stages in the operation. Figure 12.4 illustrates different levels of complexity of inventory relationships within an operation. Perhaps the simplest level is the single-stage inventory system, such as a retail store, which will have only one stock of goods to manage. An automotive parts distribution operation will have a central depot and various local distribution points which contain inventories. In many manufacturers of standard items, there are three types of inventory. The **raw material** and **components inventories** (sometimes called input inventories) receive goods from the operation's suppliers; the raw materials and components work their way through the various stages of the production process but spend considerable amounts of time as **work-in-progress** (or work-in-process) (WIP) before finally reaching the **finished goods inventory**.

Raw material

Components inventories

Work-in-progress (WIP) The number of units within a process waiting to be processed further (also called work-in-process).

Finished goods inventory

Multi-echelon inventory

A development of this last system is the **multi-echelon inventory** system. This maps the relationship of inventories between the various operations within a supply network (see Chapter 6). In Figure 12.4(d) there are five interconnected sets of inventory systems. The second-tier supplier's (yarn producer's) inventories will feed the first-tier supplier's (cloth producer's) inventories, which will in turn supply the main operation. The products are distributed to local warehouses from where they are shipped to the final customers. We will discuss the behaviour and management of such multi-echelon systems in the next chapter.

Day-to-day inventory decisions

At each point in the inventory system, operations managers need to manage the day-to-day tasks of running the system. Orders will be received from internal or external customers; these will be despatched and demand will gradually deplete the inventory. Orders will need to be placed for replenishment of the stocks; deliveries will arrive and require storing. In managing the system, operations managers are involved in three major types of decision:

- *How much to order*. Every time a replenishment order is placed, how big should it be (sometimes called the *volume decision*)?
- *When to order*. At what point in time, or at what level of stock, should the replenishment order be placed (sometimes called the *timing decision*)?
- *How to control the system*. What procedures and routines should be installed to help make these decisions? Should different priorities be allocated to different stock items? How should stock information be stored?

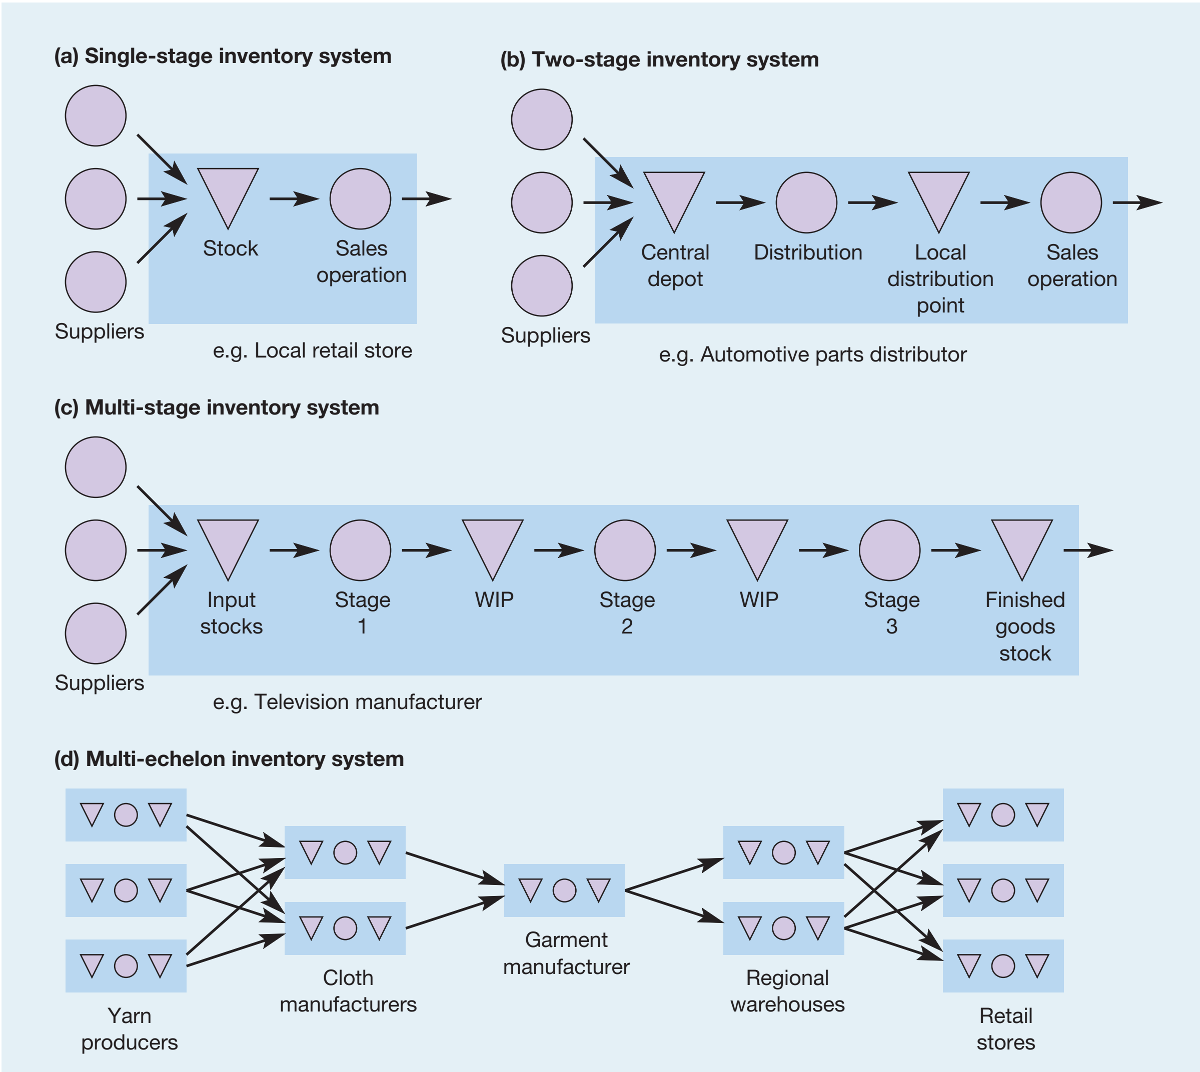

Figure 12.4 **(a) Single-stage, (b) two-stage, (c) multi-stage and (d) multi-echelon inventory systems**

The volume decision – how much to order

GO TO WEB! 12C

To illustrate this decision, consider again the example of the food and drinks we keep in our home. In managing this inventory we implicitly make decisions on *order quantity*, which is how much to purchase at one time. In making this decision we are balancing two sets of costs: the costs associated with going out to purchase the food items and the costs associated with holding the stocks. The option of holding very little or no inventory of food and purchasing each item only when it is needed has the advantage that it requires little money since purchases are made only when needed. However, it would involve purchasing provisions several times a day, which is inconvenient. At the very opposite extreme, making one journey to the local superstore every few months and purchasing all the provisions we would need until our next visit reduces the time and costs incurred in making the purchase but requires a very large amount of money each time the trip is made – money which could otherwise be in the bank and earning interest. We might also have to invest in extra cupboard units and a very large freezer. Somewhere between these extremes there will lie an ordering strategy which will minimize the total costs and effort involved in the purchase of food.

Inventory costs

The same principles apply in commercial order-quantity decisions as in the domestic situation. In making a decision on how much to purchase, operations managers must try to identify the costs which will be affected by their decision. Some costs are directly associated with order size:

1. *Cost of placing the order*. Every time that an order is placed to replenish stock, a number of transactions is needed which incurs costs to the company. These include the clerical tasks of preparing the order and all the documentation associated with it, arranging for the delivery to be made, arranging to pay the supplier for the delivery and the general costs of keeping all the information which allows us to do this. Also, if we are placing an 'internal order' on part of our own operation, there are still likely to be the same types of transaction concerned with internal administration. In addition, there could be a 'changeover' cost incurred by the part of the operation which is to supply the items, caused by the need to change from producing one type of item to another.
2. *Price discount costs*. In many industries suppliers offer discounts on the normal purchase price for large quantities; alternatively they might impose extra costs for small orders.
3. *Stock-out costs*. If we misjudge the order-quantity decision and our inventory runs out of stock, there will be costs to us incurred by failing to supply our customers. If the customers are external, they may take their business elsewhere; if internal, stock-outs could lead to idle time at the next process, inefficiencies and eventually, again, dissatisfied external customers.
4. *Working capital costs*. Soon after we receive a replenishment order, the supplier will demand payment for their goods. Eventually, when (or after) we supply our own customers, we in turn will receive payment. However, there will probably be a lag between paying our suppliers and receiving payment from our customers. During this time we will have to fund the costs of inventory. This is called the *working capital* of inventory. The costs associated with it are the interest we pay the bank for borrowing it or the opportunity costs of not investing it elsewhere.
5. *Storage costs*. These are the costs associated with physically storing the goods. Renting, heating and lighting the warehouse, as well as insuring the inventory, can be expensive, especially when special conditions are required such as low temperature or high security.
6. *Obsolescence costs*. When we order large quantities, this usually results in stocked items spending a long time stored in inventory. Then there is a risk that the items might either become obsolete (in the case of a change in fashion, for example) or deteriorate with age (in the case of most foodstuffs).
7. *Operating inefficiency costs*. According to just-in-time philosophies, high inventory levels prevent us seeing the full extent of problems within the operation. This argument is fully explored in Chapter 15.

Consignment stock

There are two points to be made about this list of costs. The first is that some of the costs will decrease as order size is increased; the first three costs are like this, whereas the other costs generally increase as order size is increased. The second point is that it may not be the same organization that incurs the costs. For example, sometimes suppliers agree to hold **consignment stock**. This means that they deliver large quantities of inventory to their customers to store but will charge for the goods only as and when they are used. In the meantime they remain the supplier's property so do not have to be financed by the customer, who does however provide storage facilities.

Inventory profiles

An inventory profile is a visual representation of the inventory level over time. Figure 12.5 shows a simplified inventory profile for one particular stock item in a retail operation. Every time an order is placed, Q items are ordered. The replenishment order arrives in one batch instantaneously. Demand for the item is then steady and perfectly predictable at a rate of D units per month. When demand has depleted the stock of the items entirely, another order of Q items instantaneously arrives and so on. Under these circumstances:

$$\text{The average inventory} = \frac{Q}{2} \text{ (because the two shaded areas in Figure 12.5 are equal)}$$

$$\text{The time interval between deliveries} = \frac{Q}{D}$$

$$\text{The frequency of deliveries} = \text{the reciprocal of the time interval} = \frac{D}{Q}$$

The economic order quantity (EOQ) formula

Economic order quantity (EOQ)
The quantity of items to order that supposedly minimizes the total cost of inventory management, derived from various EOQ formulae.

The most common approach to deciding how much of any particular item to order when stock needs replenishing is called the **economic order quantity** (EOQ) approach. This attempts to find the best balance between the advantages and disadvantages of holding stock. For example, Figure 12.6 shows two alternative order-quantity policies for an item. Plan A, represented by the unbroken line, involves ordering in quantities of 400 at a time. Demand in this case is running at 1000 units per year. Plan B, represented by the dotted line, uses smaller but more frequent replenishment orders. This time only 100 are ordered at a time, with orders being placed four times as often. However, the average inventory for plan B is one-quarter of that for plan A.

To find out whether either of these plans, or some other plan, minimizes the total cost of stocking the item, we need some further information, namely the total cost of holding one unit in stock for a period of time (C_h) and the total costs of placing an order (C_o). Generally, holding costs are taken into account by including:

- working capital costs;
- storage costs;
- obsolescence risk costs.

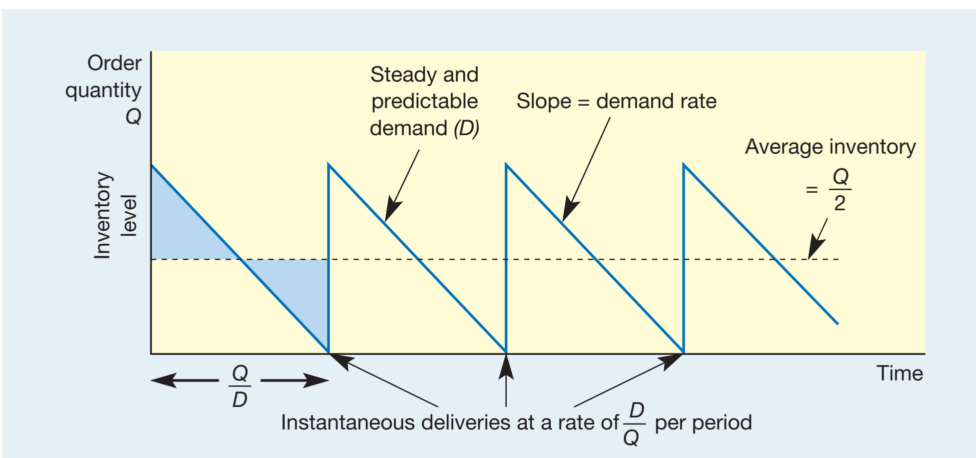

Figure 12.5 **Inventory profiles chart the variation in inventory level**

GO TO WEB! 12D

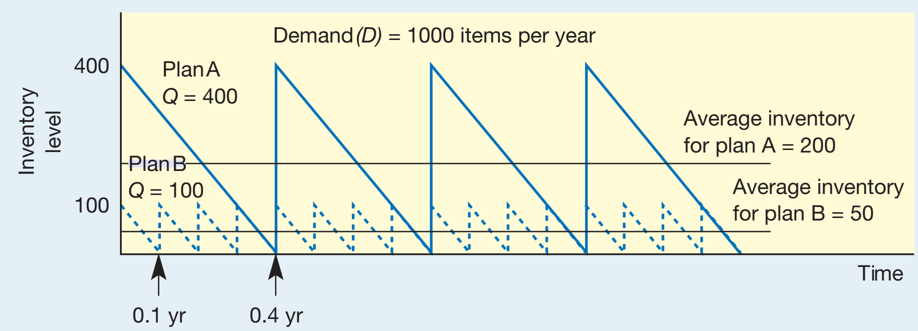

Figure 12.6 Two alternative inventory plans with different order quantities (Q)

Order costs are calculated by taking into account:

- cost of placing the order (including transportation of items from suppliers if relevant);
- price discount costs.

In this case the cost of holding stocks is calculated at £1 per item per year and the cost of placing an order is calculated at £20 per order.

We can now calculate total holding costs and ordering costs for any particular ordering plan as follows:

$$\text{Holding costs} = \text{holding cost/unit} \times \text{average inventory}$$

$$= C_h \times \frac{Q}{2}$$

$$\text{Ordering costs} = \text{ordering cost} \times \text{number of orders per period}$$

$$= C_o \times \frac{D}{Q}$$

$$\text{So, total cost, } C_t = \frac{C_h Q}{2} + \frac{C_o D}{Q}$$

We can now calculate the costs of adopting plans with different order quantities. These are illustrated in Table 12.2. As we would expect with low values of Q, holding costs are low but the costs of placing orders are high because orders have to be placed very frequently. As Q

Table 12.2 Costs of adoption of plans with different order quantities

Demand (D) = 1000 units per year Order costs (C_o) = £20 per order			*Holding costs (C_h) = £1 per item per year*		
Order quantity (Q)	*Holding costs (0.5Q × C_h)*	+	*Order costs ((D/Q) × C_o)*	=	*Total costs*
50	25		20 × 20 = 400		425
100	50		10 × 20 = 200		250
150	75		6.7 × 20 = 134		209
200	100		5 × 20 = 100		200*
250	125		4 × 20 = 80		205
300	150		3.3 × 20 = 66		216
350	175		2.9 × 20 = 58		233
400	200		2.5 × 20 = 50		250

* Minimum total cost

increases, the holding costs increase but the costs of placing orders decrease. Initially the decrease in ordering costs is greater than the increase in holding costs and the total cost falls. After a point, however, the decrease in ordering costs slows, whereas the increase in holding costs remains constant and the total cost starts to increase. In this case the order quantity, Q, which minimizes the sum of holding and order costs, is 200. This 'optimum' order quantity is called the *economic order quantity*. This is illustrated graphically in Figure 12.7.

A more elegant method of finding the EOQ is to derive its general expression. This can be done using simple differential calculus as follows. From before:

Total cost = holding cost + order cost

$$C_t = \frac{C_h Q}{2} + \frac{C_o D}{Q}$$

The rate of change of total cost is given by the first differential of C_t with respect to Q:

GO TO WEB! 12E

$$\frac{dC_t}{dQ} = \frac{C_h}{2} - \frac{C_o D}{Q^2}$$

The lowest cost will occur when $\frac{dC_t}{dQ} = 0$, that is:

$$0 = \frac{C_h}{2} - \frac{C_o D}{Q_o^2}$$

where Q_o = the *EOQ*. Rearranging this expression gives:

$$Q_o = EOQ = \sqrt{\frac{2C_o D}{C_h}}$$

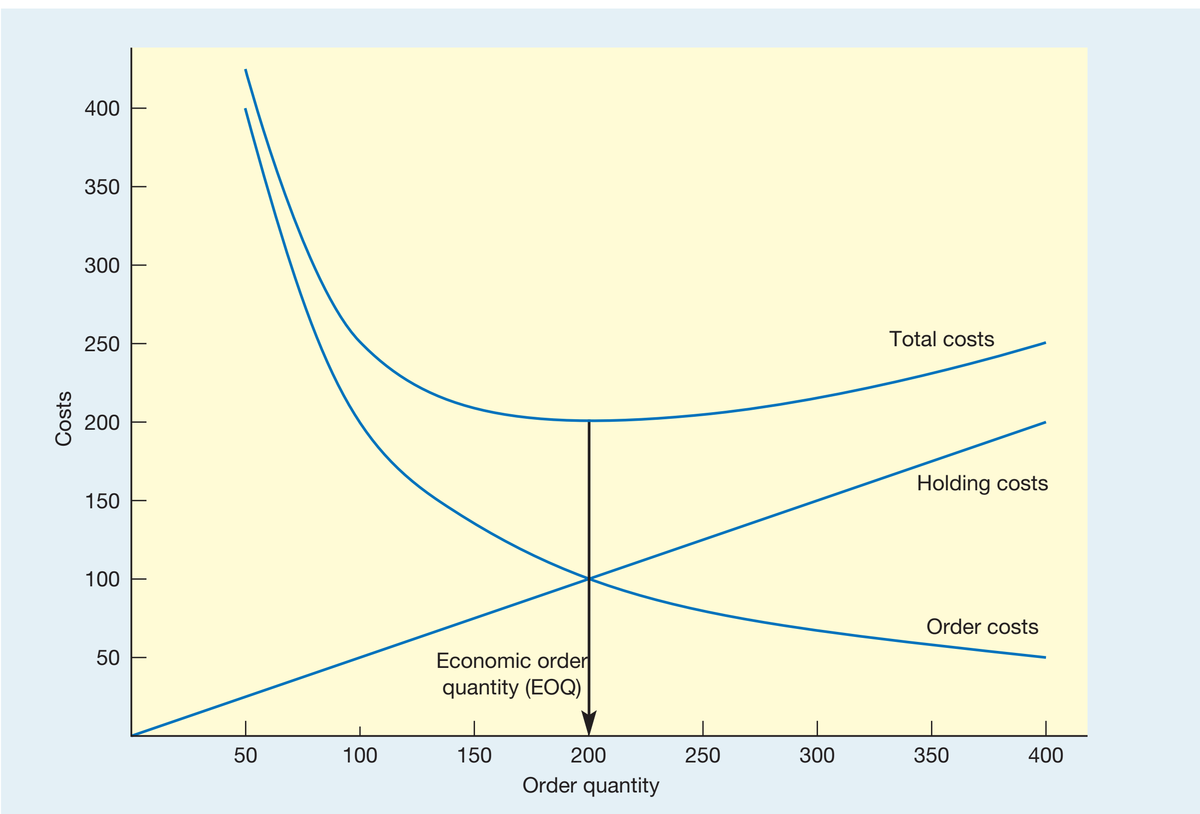

Figure 12.7 Graphical representation of the economic order quantity

When using the *EOQ*:

$$\text{Time between orders} = \frac{EOQ}{D}$$

$$\text{Order frequency} = \frac{D}{EOQ} \text{ per period}$$

Sensitivity of the EOQ

Examination of the graphical representation of the total cost curve in Figure 12.7 shows that, although there is a single value of *Q* which minimizes total costs, any relatively small deviation from the EOQ will not increase total costs significantly. In other words, costs will be near-optimum provided a value of *Q* which is reasonably close to the EOQ is chosen. Put another way, small errors in estimating either holding costs or order costs will not result in a significant deviation from the EOQ. This is a particularly convenient phenomenon because, in practice, both holding and order costs are not easy to estimate accurately.

Worked example

A building materials supplier obtains its bagged cement from a single supplier. Demand is reasonably constant throughout the year and last year the company sold 2000 tonnes of this product. It estimates the costs of placing an order at around £25 each time an order is placed and calculates that the annual cost of holding inventory is 20 per cent of purchase cost. The company purchases the cement at £60 per tonne. How much should the company order at a time?

$$EOQ \text{ for cement} = \sqrt{\frac{2C_oD}{C_h}}$$

$$= \sqrt{\frac{2 \times 25 \times 2000}{0.2 \times 60}}$$

$$= \sqrt{\frac{100\,000}{12}}$$

$$= 91.287 \text{ tonnes}$$

After calculating the EOQ the operations manager feels that placing an order for 91.287 tonnes *exactly* seems somewhat over-precise. Why not order a convenient 100 tonnes?

Total cost of ordering plan for Q = 91.287:

$$= \frac{C_hQ}{2} + \frac{C_oD}{Q}$$

$$= \frac{(0.2 \times 60) \times 91.287}{2} + \frac{25 \times 2000}{91.287}$$

$$= £1095.454$$

Total cost of ordering plan for Q = 100:

$$= \frac{(0.2 \times 60) \times 100}{2} + \frac{25 \times 2000}{100}$$

$$= £1100$$

The extra cost of ordering 100 tonnes at a time is £1100 – £1095.45 = £4.55. The operations manager therefore should feel confident in using the more convenient order quantity.

Gradual replacement – the economic batch quantity (EBQ) model

Although the simple inventory profile shown in Figure 12.5 made some simplifying assumptions, it is broadly applicable in most situations where each complete replacement order arrives at one point in time. In many cases, however, replenishment occurs over a time period rather than in one lot. A typical example of this is where an internal order is placed for a batch of parts to be produced on a machine. The machine will start to produce the parts and ship them in a more or less continuous stream into inventory, but at the same time demand is continuing to remove parts from the inventory. Provided the rate at which parts are being made and put into the inventory (P) is higher than the rate at which demand is depleting the inventory (D), the size of the inventory will increase. After the batch has been completed the machine will be reset (to produce some other part) and demand will continue to deplete the inventory level until production of the next batch begins. The resulting profile is shown in Figure 12.8. Such a profile is typical for cycle inventories supplied by batch processes, where items are produced internally and intermittently. For this reason the minimum-cost batch quantity for this profile is called the **economic batch quantity (EBQ)**. It is also sometimes known as the economic manufacturing quantity (EMQ) or the production order quantity (POQ). It is derived as follows:

Economic batch quantity (EBQ)
The amount of items to be produced by a machine or process that supposedly minimizes the costs associated with production and inventory holding.

$$\text{Maximum stock level} = M$$

$$\text{Slope of inventory build-up} = P - D$$

Also, as is clear from Figure 12.8:

$$\text{Slope of inventory build-up} = M \div \frac{Q}{P}$$

$$= \frac{MP}{Q}$$

So,

$$\frac{MP}{Q} = P - D$$

$$M = \frac{Q(P - D)}{P}$$

$$\text{Average inventory level} = \frac{M}{2}$$

$$= \frac{Q(P - D)}{2P}$$

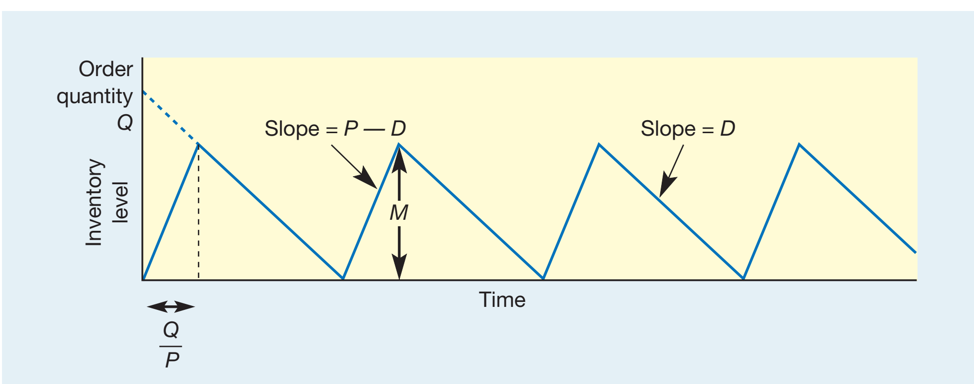

Figure 12.8 **Inventory profile for gradual replacement of inventory**

As before:

Total cost = holding cost + order cost

$$C_t = \frac{C_h Q(P - D)}{2P} + \frac{C_o D}{Q}$$

$$\frac{dC_t}{dQ} = \frac{C_h(P - D)}{2P} - \frac{C_o D}{Q^2}$$

Again, equating to zero and solving Q gives the minimum-cost order quantity EBQ:

$$EBQ = \sqrt{\frac{2C_o D}{C_h(1 - (D/P)}}$$

Worked example

The manager of a bottle-filling plant which bottles soft drinks needs to decide how long a 'run' of each type of drink to process. Demand for each type of drink is reasonably constant at 80,000 per month (a month has 160 production hours). The bottling lines fill at a rate of 3000 bottles per hour, but take an hour to clean and reset between different drinks. The cost (of labour and lost production capacity) of each of these changeovers has been calculated at £100 per hour. Stock-holding costs are counted at £0.1 per bottle per month.

$$D = 80{,}000 \text{ per month}$$
$$= 500 \text{ per hour}$$

$$EBQ = \sqrt{\frac{2C_o D}{C_h(1 - (D/P)}}$$

$$= \sqrt{\frac{2 \times 100 \times 80{,}000}{0.1(1 - (500/3000))}}$$

$$EBQ = 13{,}856$$

The staff who operate the lines have devised a method of reducing the changeover time from 1 hour to 30 minutes. How would that change the EBQ?

New C_o = £50

$$\text{New } EBQ = \sqrt{\frac{2 \times 50 \times 80{,}000}{0.1(1 - (500/3000))}}$$

$$= 9798$$

Critical commentary

The approach to determining order quantity which involves optimizing costs of holding stock against costs of ordering stock, typified by the EOQ and EBQ models, has always been subject to criticisms. Originally these concerned the validity of some of the assumptions of the model; more recently they have involved the underlying rationale of the approach itself. The criticisms fall into four broad categories, all of which we shall examine further:

→

- **the assumptions included in the EOQ models are simplistic;**
- **the real costs of stock in operations are not as assumed in EOQ models;**
- **the models are really descriptive and should not be used as prescriptive devices;**
- **cost minimization is not an appropriate objective for inventory management.**

Responding to the criticisms of EOQ

In order to keep EOQ-type models relatively straightforward, it was necessary to make assumptions. These concerned such things as the stability of demand, the existence of a fixed and identifiable ordering cost, that the cost of stock-holding can be expressed by a linear function, shortage costs which were identifiable and so on. While these assumptions are rarely strictly true, most of them can approximate to reality. Furthermore, the shape of the total cost curve has a relatively flat optimum point which means that small errors will not significantly affect the total cost of a near-optimum order quantity. However, at times the assumptions do pose severe limitations to the models. For example, the assumption of steady demand (or even demand which conforms to some known probability distribution) is untrue for a wide range of the operation's inventory problems. For example, a bookseller might be very happy to adopt an EOQ-type ordering policy for some of its most regular and stable products such as dictionaries and popular reference books. However, the demand patterns for many other books could be highly erratic, dependent on critics' reviews and word-of-mouth recommendations. In such circumstances it is simply inappropriate to use EOQ models.

Cost of stock

Other questions surround some of the assumptions made concerning the nature of stock-related costs. For example, placing an order with a supplier as part of a regular and multi-item order might be relatively inexpensive, whereas asking for a special one-off delivery of an item could prove far more costly. Similarly with stock-holding costs – although many companies make a standard percentage charge on the purchase price of stock items, this might not be appropriate over a wide range of stock-holding levels. The marginal costs of increasing stock-holding levels might be merely the cost of the working capital involved. Or it might necessitate the construction or lease of a whole new stock-holding facility such as a warehouse. Operations managers using an EOQ-type approach must check that the decisions implied by the use of the formulae do not exceed the boundaries within which the cost assumptions apply. In Chapter 15 we explore the just-in-time approach that sees inventory as being largely negative. However, it is useful at this stage to examine the effect on an EOQ approach of regarding inventory as being more costly than previously believed. Increasing the slope of the holding cost line increases the level of total costs of *any* order quantity, but more significantly shifts the minimum cost point substantially to the left, in favour of a lower economic order quantity. In other words, the less willing an operation is to hold stock on the grounds of cost, the more it should move towards smaller, more frequent ordering.

Using EOQ models as prescriptions

Perhaps the most fundamental criticism of the EOQ approach again comes from the Japanese-inspired 'lean' and JIT philosophies. The EOQ tries to optimize order decisions. Implicitly the costs involved are taken as fixed, in the sense that the task of operations managers is to find out what are the true costs rather than to change them in any way. EOQ is essentially a reactive approach. Some critics would argue that it fails to ask the right question. Rather than asking the EOQ question of 'What is the optimum order quantity?', operations managers should really be asking, 'How can I change the operation in some way so as to reduce the overall level of inventory I need to hold?' The EOQ approach may be a reasonable description of stock-holding costs but should not necessarily be taken as a strict prescription over what decisions to take. For example, many organizations have made considerable efforts to reduce the effective cost of placing an order. Often they have done this by working to reduce changeover times on machines. This means that less time is taken changing over from one product to the other and therefore less operating capacity is lost, which in turn reduces the cost of the changeover. Under these circumstances, the order cost curve in the EOQ formula reduces and, in turn, reduces the effective economic order quantity. Figure 12.9 shows the EOQ formula represented graphically with increased holding costs (see the previous discussion) and reduced order costs. The net effect of this is to significantly reduce the value of the EOQ.

Should the cost of inventory be minimized?

Many organizations (such as supermarkets and wholesalers) make most of their revenue and profits simply by holding and supplying inventory. Because their main investment is in the inventory it is critical that they make a good return on this capital by ensuring that it has the highest possible 'stock turn' (defined later in this chapter) and/or gross profit margin. Alternatively, they may also be concerned to maximize the use of space by seeking to maximize the profit earned per square metre. The EOQ model does not address these objectives. Similarly for products that deteriorate or go out of fashion, the EOQ model can result in excess inventory of slower-moving items. In fact, the EOQ model is rarely used in such organizations and there is more likely to be a system of periodic review (described later) for regular ordering of replenishment inventory. A typical builders' supply merchant might carry around 50,000 different items of stock (stock-keeping units or SKUs). However, most of these cluster into larger families of items such as paints, sanitary ware or metal fixings. Single orders are placed at regular intervals for all the required replenishments in the supplier's range and these are then delivered together at one time. If such deliveries were made weekly, then on average the individual item order quantities will be for only one week's usage. Less popular items, or ones with erratic demand patterns, can be individually ordered at the same time or (when urgent) can be delivered the next day by carrier.

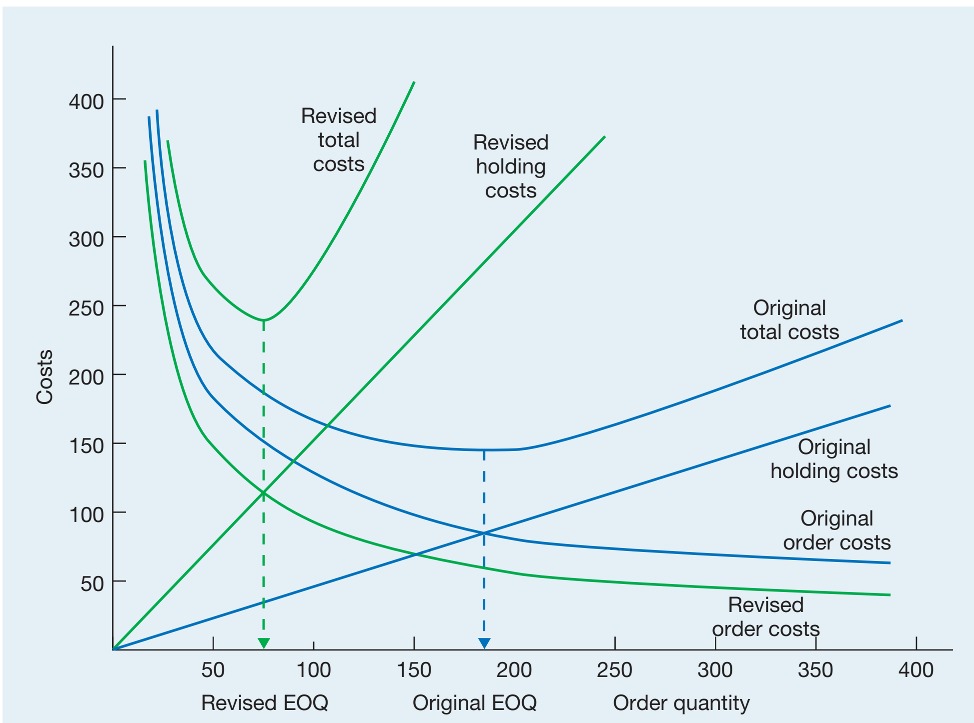

Figure 12.9 **If the true costs of stock holding are taken into account, and if the cost of ordering (or changeover) is reduced, the economic order quantity (EOQ) is much smaller**

Short case **Howard Smith Paper Group**[3]

The Howard Smith Paper Group operates the most advanced warehousing operation within the European paper merchanting sector, delivering over 120,000 tonnes of paper annually. The function of a paper merchant is to provide the link between the paper mills and the printers or converters. This is illustrated in Figure 12.10. It is a sales- and service-driven business, so the role of the operation function is to deliver whatever the salesperson has promised to the customer. Usually, this means precisely the right product at the right time at the right place and in the right quantity. The company's operations are divided into two areas, 'logistics' which combines all warehousing and logistics tasks, and 'supply side' which includes inventory planning, purchasing and merchandising decisions. Its main stocks are held at the national distribution centre, located in Northampton in the middle part of the UK. This location was chosen because it is at the centre of the company's main customer location and also because it has good access to motorways.

The key to any efficient merchanting operation lies in its ability to do three things well. First, it must efficiently store the desired volume of required inventory. Second, it must have a 'goods inward' programme that sources the required volume of desired inventory. Third, it must be able to fulfil customer orders by 'picking' the desired goods fast and accurately from its warehouse. The warehouse is operational 24 hours per day, 5 days per week. A total of 52 staff are employed in the warehouse, including maintenance and cleaning staff. Skill sets are not an issue, since all pickers are trained for all tasks. This facilitates easier capacity management, since pickers can be deployed where most urgently needed. Contract labour is used on occasions, although this is less effective

Source: Howard Smith Paper Group

Dispatch activity at Howard Smith Paper Group

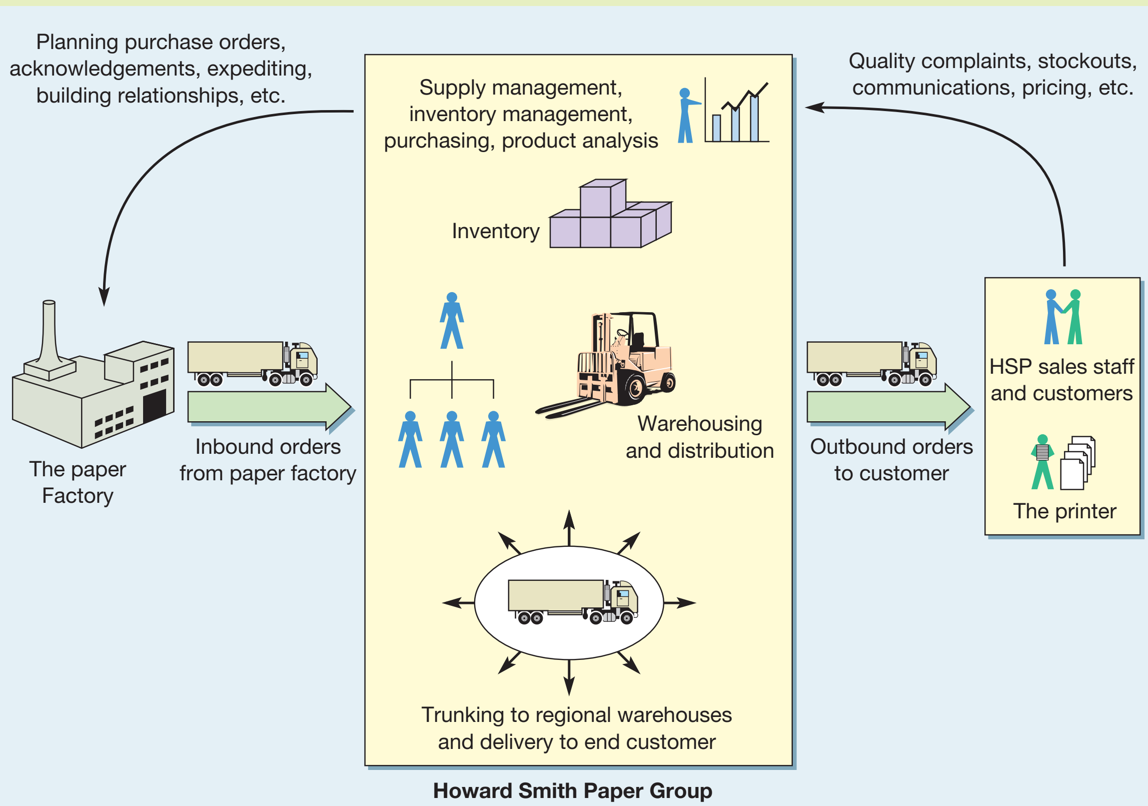

Figure 12.10 The role of the paper merchant

because the staff tend to be less motivated and have to learn the job.

At the heart of the company's operations is a warehouse known as a 'dark warehouse'. All picking and movement within the dark warehouse is fully automatic and there is no need for any person to enter the high bay stores and picking area. The important difference with this warehouse operation is that pallets are brought to the pickers. Conventional paper merchants send pickers with handling equipment into the warehouse aisles for stock. A warehouse computer system (WCS) controls the whole operation without the need for human input. It manages pallet location and retrieval, robotic crane missions, automatic conveyors, bar code label production and scanning, and all picking routines and priorities. It also calculates operator activity and productivity measures, as well as issuing documentation and planning transportation schedules. The fact that all product is identified by a unique bar code means that accuracy is guaranteed. The unique user log-on ensures that any picking errors can be traced back to the name of the picker, to ensure further errors do not occur. The WCS is linked to the company's ERP system (we will deal with ERP in Chapter 14), such that once the order has been placed by a customer, computers manage the whole process from order placement to order despatch.

Question

1 Why has the Howard Smith Paper Group invested so much capital in automating its inventory storage and control capabilities?

The timing decision – when to place an order

When we assumed that orders arrived instantaneously and demand was steady and predictable, the decision on when to place a replenishment order was self-evident. An order would be placed as soon as the stock level reached zero. This would arrive instantaneously and prevent any stock-out occurring. If replenishment orders do not arrive instantaneously but have a lag between the order being placed and it arriving in the inventory, we can calculate the timing of a replacement order as shown in Figure 12.11. The lead time for an order to arrive is two weeks in this case, so the **re-order point (ROP)** is the point at which stock will fall to zero minus the order lead time. Alternatively, we can define the point in terms of the level which the inventory will have reached when a replenishment order needs to be placed. In this case this occurs at a **re-order level (ROL)** of 200 items.

Re-order point
The point in time at which more items are ordered, usually calculated to ensure that inventory does not run out before the next batch of inventory arrives.

Re-order level
The level of inventory at which more items are ordered, usually calculated to ensure that inventory does not run out before the next batch of inventory arrives.

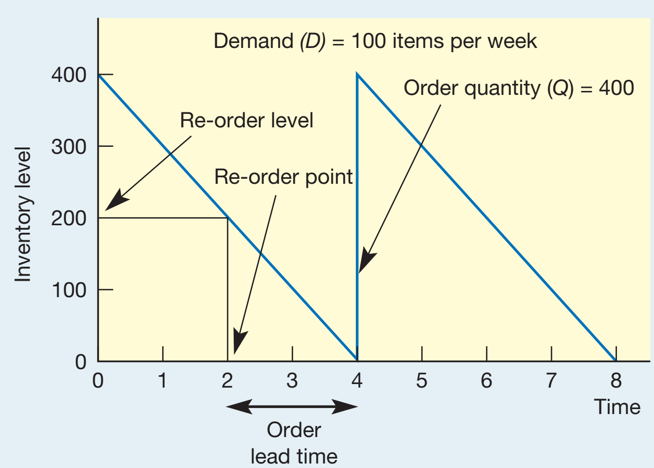

Figure 12.11 Re-order level (ROL) and re-order point (ROP) are derived from the order lead time and demand rate

However, this assumes that both the demand and the order lead time are perfectly predictable. In most cases, of course, this is not so. Both demand and the order lead time are likely to vary to produce a profile which looks something like that in Figure 12.12. In these circumstances it is necessary to make the replenishment order somewhat earlier than would be the case in a purely deterministic situation. This will result in, on average, some stock still being in the inventory when the replenishment order arrives. This is buffer (safety) stock. The earlier the replenishment order is placed, the higher will be the expected level of safety stocks (s) when the replenishment order arrives. But because of the variability of both lead time (t) and demand rate (d), there will sometimes be a higher-than-average level of safety stock and sometimes lower. The main consideration in setting safety stock is not so much the average level of stock when a replenishment order arrives but rather the probability that the stock will not have run out before the replenishment order arrives.

Lead-time usage
The amount of inventory that will be used between ordering replenishment and the inventory arriving, usually described by a probability distribution to account for uncertainty in demand and lead time.

The key statistic in calculating how much safety stock to allow is the probability distribution which shows the **lead-time usage**. The lead-time usage distribution is a combination of the distributions which describe lead-time variation and the demand rate during the lead time. If safety stock is set below the lower limit of this distribution, there will be shortages every single replenishment cycle. If safety stock is set above the upper limit of the distribution, there is no chance of stock-outs occurring. Usually, safety stock is set to give a predetermined likelihood that stock-outs will not occur. Figure 12.12 shows that, in this case, the first replenishment order arrived after t_1, resulting in a lead-time usage of d_1. The second replenishment order took longer, t_2, and demand rate was also higher, resulting in a lead-time usage of d_2. The third order cycle shows several possible inventory profiles for different conditions of lead-time usage and demand rate.

GO TO WEB! 12F

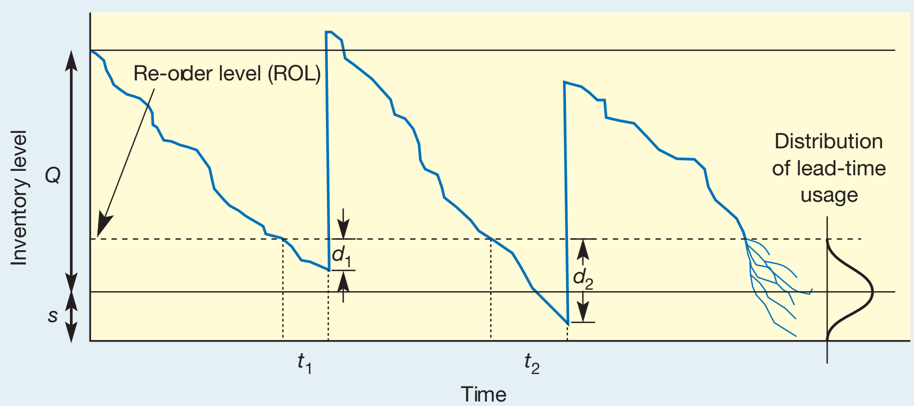

Figure 12.12 Safety stock (*s*) helps to avoid stock-outs when demand and/or order lead time are uncertain

Worked example

A company which imports running shoes for sale in its sports shops can never be certain of how long, after placing an order, the delivery will take. Examination of previous orders reveals that out of ten orders one took one week, two took two weeks, four took three weeks, two took four weeks and one took five weeks. The rate of demand for the shoes also varies between 110 pairs per week and 140 pairs per week. There is a 0.2 probability of the demand rate being either 110 or 140 pairs per week and a 0.3 chance of demand being either 120 or 130 pairs per week. The company needs to decide when it should place replenishment orders if the probability of a stock-out is to be less than 10 per cent.

Both lead time and the demand rate during the lead time will contribute to the lead-time usage. So the distributions which describe each will need to be combined. Figure 12.13 and Table 12.3 show how this can be done. Taking lead time to be either one, two, three, four or five weeks, and demand rate to be either 110, 120, 130 or 140 pairs per week, and also assuming the two variables to be independent, the distributions can be combined as shown in Table 12.4. Each element in the matrix shows a possible lead-time usage with the probability of its occurrence. So if the lead time is one week and the demand rate is 110 pairs per week, the actual lead-time usage will be 1 × 110 = 110 pairs. Since there is a 0.1 chance of the lead time being one week, and a 0.2 chance of demand rate being 110 pairs per week, the probability of both these events occurring is 0.1 × 0.2 = 0.02.

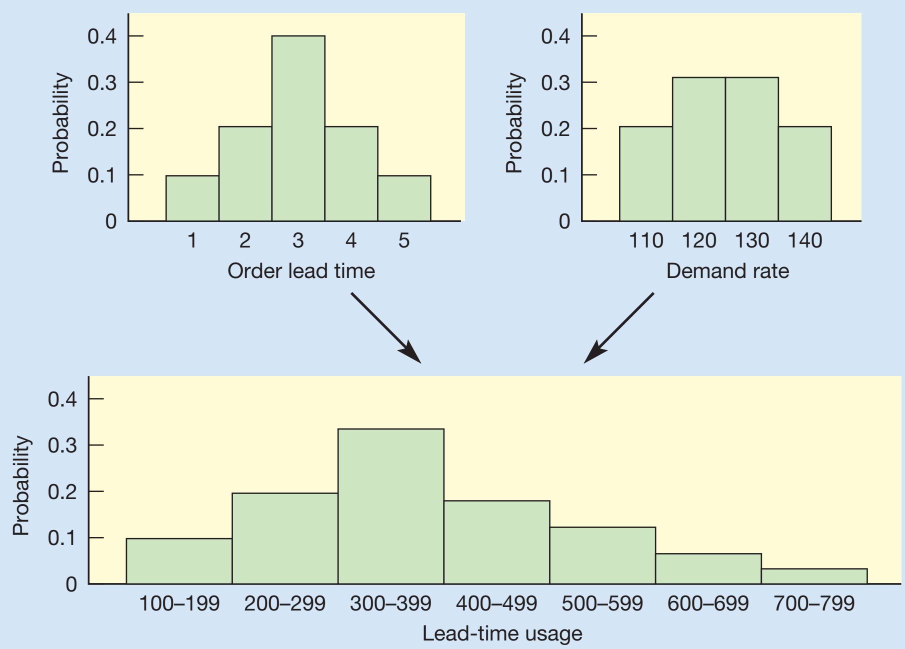

Figure 12.13 The probability distributions for order lead time and demand rate combine to give the lead-time usage distribution

Table 12.3 Matrix of lead-time and demand-rate probabilities

			Lead-time probabilities				
			1 *0.1*	*2* *0.2*	*3* *0.4*	*4* *0.2*	*5* *0.1*
Demand-rate probabilities	110	0.2	110 (0.02)	220 (0.04)	330 (0.08)	440 (0.04)	550 (0.02)
	120	0.3	120 (0.03)	240 (0.06)	360 (0.12)	480 (0.06)	600 (0.03)
	130	0.3	130 (0.03)	260 (0.06)	390 (0.12)	520 (0.06)	650 (0.03)
	140	0.2	140 (0.02)	280 (0.04)	420 (0.08)	560 (0.04)	700 (0.02)

We can now classify the possible lead-time usages into histogram form. For example, summing the probabilities of all the lead-time usages which fall within the range 100–199 (all the first column) gives a combined probability of 0.1. Repeating this for subsequent intervals results in Table 12.4.

Table 12.4 Combined probabilities

Lead-time usage	100–199	200–299	300–399	400–499	500–599	600–699	700–799
Probability	0.1	0.2	0.32	0.18	0.12	0.06	0.02

This shows the probability of each possible range of lead-time usage occurring, but it is the cumulative probabilities that are needed to predict the likelihood of stock-out (see Table 12.5).

Table 12.5 Combined probabilities

Lead-time usage, *X*	100	200	300	400	500	600	700	800
Probability of usage being greater than *X*	1.0	0.9	0.7	0.38	0.2	0.08	0.02	0

Setting the re-order level at 600 would mean that there is only a 0.08 chance of usage being greater than available inventory during the lead time, i.e. there is a less than 10 per cent chance of a stock-out occurring.

Continuous and periodic review

Continuous review
An approach to managing inventory that makes inventory-related decisions when inventory reaches a particular level, as opposed to period review.

Periodic review
An approach to making inventory decisions that defines points in time for examining inventory levels and then makes decisions accordingly, as opposed to continuous review.

The approach we have described to making the replenishment timing decision is often called the **continuous review** approach. This is because, to make the decision in this way, there must be a process to review the stock level of each item continuously and then place an order when the stock level reaches its re-order level. The virtue of this approach is that, although the timing of orders may be irregular (depending on the variation in demand rate), the order size (Q) is constant and can be set at the optimum economic order quantity. Such continual checking on inventory levels can be time-consuming, especially when there are many stock withdrawals compared with the average level of stock, but in an environment where all inventory records are computerized, this should not be a problem unless the records are inaccurate.

An alternative and far simpler approach, but one which sacrifices the use of a fixed (and therefore possibly optimum) order quantity, is called the **periodic review** approach. Here, rather than ordering at a predetermined re-order level, the periodic approach orders at a fixed and regular time interval. So the stock level of an item could be found, for example, at the end of every month and a replenishment order placed to bring the stock up to a predetermined level. This level is calculated to cover demand between the replenishment order being placed and the following replenishment order arriving. Figure 12.14 illustrates the parameters for the periodic review approach.

At time T_1 in Figure 12.14 the inventory manager would examine the stock level and order sufficient to bring it up to some maximum, Q_m. However, that order of Q_1 items will not arrive until a further time of t_1 has passed, during which demand continues to deplete the stocks. Again, both demand and lead time are uncertain. The Q_1 items will arrive and bring the stock up to some level lower than Q_m (unless there has been no demand during t_1). Demand then continues until T_2, when again an order Q_2 is placed which is the difference between the current stock at T_2 and Q_m. This order arrives after t_2, by which time demand has depleted the stocks further. Thus the replenishment order placed at T_1 must be able to cover for the demand which occurs until T_2 and t_2. Safety stocks will need to be calculated, in a similar manner to before, based on the distribution of usage over this period.

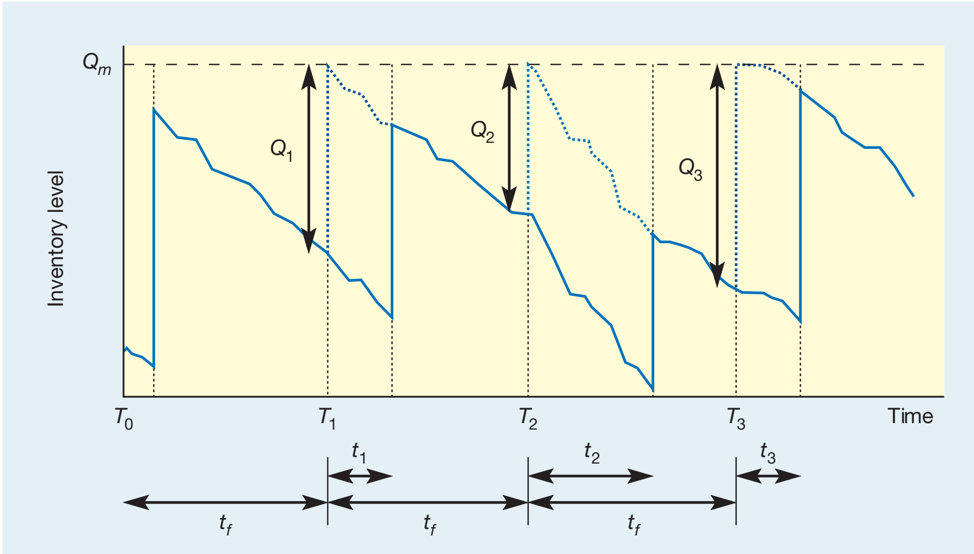

Figure 12.14 A periodic review approach to order timing with probabilistic demand and lead time

The time interval

The interval between placing orders, t_1, is usually calculated on a deterministic basis and derived from the EOQ. So, for example, if the demand for an item is 2,000 per year, the cost of placing an order £25 and the cost of holding stock £0.5 per item per year:

$$\text{EOQ} = \sqrt{\frac{2C_oD}{C_h}} = \sqrt{\frac{2 \times 2000 \times 25}{0.5}} = 447$$

The optimum time interval between orders, t_f, is therefore:

$$t_f = \frac{EOQ}{D} = \frac{447}{2000} \text{ years}$$

$$= 2.68 \text{ months}$$

It may seem paradoxical to calculate the time interval assuming constant demand when demand is, in fact, uncertain. However, uncertainties in both demand and lead time can be allowed for by setting Q_m to allow for the desired probability of stock-out based on usage during the period t_f + lead time.

Two-bin and three-bin systems

Keeping track of inventory levels is especially important in continuous review approaches to re-ordering. A simple and obvious method of indicating when the re-order point has been reached is necessary, especially if there is a large number of items to be monitored. The two- and three-bin systems illustrated in Figure 12.15 are such methods. The simple **two-bin system** involves storing the re-order point quantity plus the safety inventory quantity in the second bin and using parts from the first bin. When the first bin empties, that is the signal to order the next re-order quantity. Sometimes the safety inventory is stored in a third bin (the **three-bin system**), so it is clear when demand is exceeding that which was expected. Different 'bins' are not always necessary to operate this type of system. For example, a common practice in retail operations is to store the second 'bin' quantity upside-down behind or under the first 'bin' quantity. Orders are then placed when the upside-down items are reached.

Two-bin system

Three-bin system

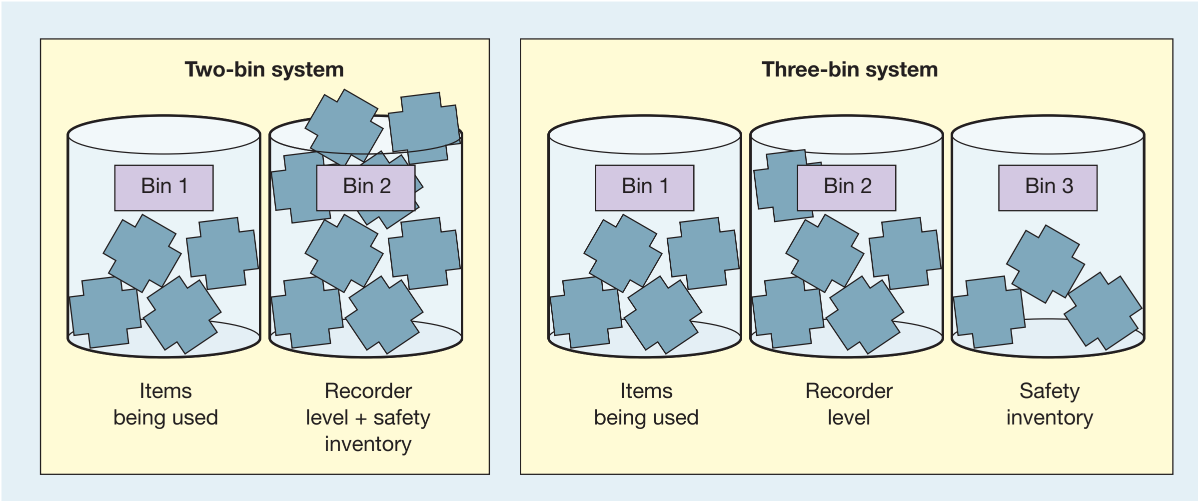

Figure 12.15 The two-bin and three-bin systems of re-ordering

Inventory analysis and control systems

The models we have described, even the ones which take a probabilistic view of demand and lead time, are still simplified compared with the complexity of real stock management. Coping with many thousands of stocked items, supplied by many hundreds of different suppliers, with possibly tens of thousands of individual customers, makes for a complex and dynamic operations task. In order to control such complexity, operations managers have to do two things. First, they have to discriminate between different stocked items, so that they can apply a degree of control to each item which is appropriate to its importance. Second, they need to invest in an information-processing system which can cope with their particular set of inventory control circumstances.

Usage value
A term used in inventory control to indicate the quantity of items used or sold multiplied by their value or price.

Pareto law
A general law found to operate in many situations that indicates that 20 per cent of something causes 80 per cent of something else, often used in inventory management (20 per cent of products produce 80 per cent of sales value) and improvement activities (20 per cent of types of problems produce 80 per cent of disruption).

ABC inventory control
An approach to inventory control that classes inventory by its usage value and varies the approach to managing it accordingly.

Inventory priorities – the ABC system

In any inventory which contains more than one stocked item, some items will be more important to the organization than others. Some, for example, might have a very high usage rate, so if they ran out many customers would be disappointed. Other items might be of particularly high value, so excessively high inventory levels would be particularly expensive. One common way of discriminating between different stock items is to rank them by the usage value (their usage rate multiplied by their individual value). Items with a particularly high usage value are deemed to warrant the most careful control, whereas those with low usage values need not be controlled quite so rigorously. Generally, a relatively small proportion of the total range of items contained in an inventory will account for a large proportion of the total usage value. This phenomenon is known as the Pareto law (after the person who described it), sometimes referred to as the 80/20 rule. It is called this because, typically, 80 per cent of an operation's sales are accounted for by only 20 per cent of all stocked item types. The Pareto law is also used elsewhere in operations management (see, for example, Chapter 18). Here the relationship can be used to classify the different types of items kept in an inventory by their usage value. ABC inventory control allows inventory managers to concentrate their efforts on controlling the more significant items of stock.

- *Class A* items are those 20 per cent or so of high-usage value items which account for around 80 per cent of the total usage value.

- *Class B* items are those of medium-usage value, usually the next 30 per cent of items which often account for around 10 per cent of the total usage value.
- *Class C* items are those low-usage value items which, although comprising around 50 per cent of the total types of items stocked, probably account for only around 10 per cent of the total usage value of the operation.

Worked example

Table 12.6 shows all the parts stored by an electrical wholesaler. The 20 different items stored vary in terms of both their usage per year and cost per item as shown. However, the wholesaler has ranked the stock items by their usage value per year. The total usage value per year is £5,569,000. From this it is possible to calculate the usage value per year of each item as a percentage of the total usage value and from that a running cumulative total of the usage value as shown. The wholesaler can then plot the cumulative percentage of all stocked items against the cumulative percentage of their value. So, for example, the part with stock number A/703 is the highest-value part and accounts for 25.14 per cent of the total inventory value. As a part, however, it is only one-twentieth or 5 per cent of the total number of items stocked. This item, together with the next highest value item (D/012), accounts for only 10 per cent of the total number of items stocked, yet accounts for 47.37 per cent of the value of the stock and so on.

Table 12.6 Warehouse items ranked by usage value

Stock no.	*Usage (items/year)*	*Cost (£/item)*	*Usage value (£000/year)*	*% of total value*	*Cumulative % of total value*
A/703	700	20.00	1400	25.14	25.14
D/012	450	2.75	1238	22.23	47.37
A/135	1000	0.90	900	16.16	63.53
C/732	95	8.50	808	14.51	78.04
C/375	520	0.54	281	5.05	83.09
A/500	73	2.30	168	3.02	86.11
D/111	520	0.22	114	2.05	88.16
D/231	170	0.65	111	1.99	90.15
E/781	250	0.34	85	1.53	91.68
A/138	250	0.30	75	1.34	93.02
D/175	400	0.14	56	1.01	94.03
E/001	80	0.63	50	0.89	94.92
C/150	230	0.21	48	0.86	95.78
F/030	400	0.12	48	0.86	96.64
D/703	500	0.09	45	0.81	97.45
D/535	50	0.88	44	0.79	98.24
C/541	70	0.57	40	0.71	98.95
A/260	50	0.64	32	0.57	99.52
B/141	50	0.32	16	0.28	99.80
D/021	20	0.50	10	0.20	100.00
Total			5 569	100.00	

→

This is shown graphically in Figure 12.16. Here the wholesaler has classified the first four part numbers (20 per cent of the range) as Class A items and will monitor the usage and ordering of these items closely and frequently. A few improvements in order quantities or safety stocks for these items could bring significant savings. The six next part numbers, C/375 through to A/138 (30 per cent of the range), are to be treated as Class B items with slightly less effort devoted to their control. All other items are Class C items whose stocking policy is reviewed only occasionally.

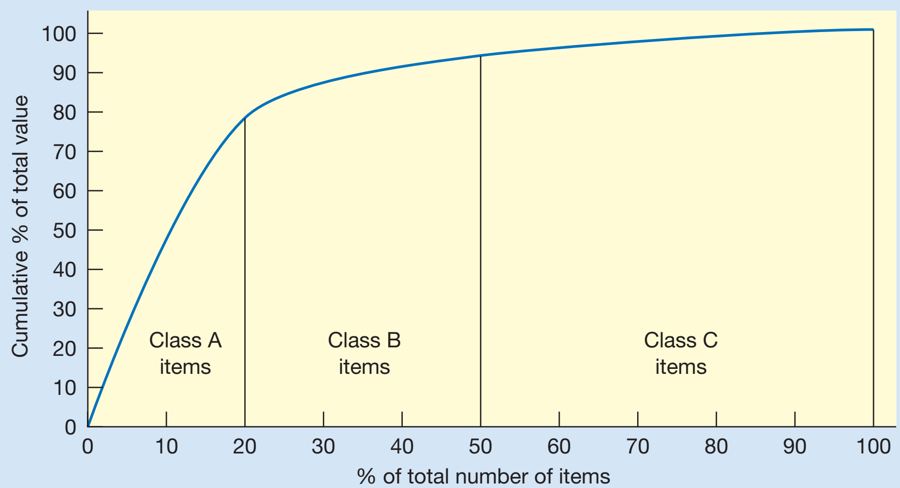

Figure 12.16 Pareto curve for items in a warehouse

Although annual usage and value are the two criteria most commonly used to determine a stock classification system, other criteria might also contribute towards the (higher) classification of an item:

- *Consequence of stock-out*. High priority might be given to those items which would seriously delay or disrupt other operations, or the customers, if they were not in stock.
- *Uncertainty of supply*. Some items, although of low value, might warrant more attention if their supply is erratic or uncertain.
- *High obsolescence or deterioration risk*. Items which could lose their value through obsolescence or deterioration might need extra attention and monitoring.

Some more complex stock classification systems might include these criteria by classifying on an A, B, C basis for each. For example, a part might be classed as A/B/A meaning it is an A category item by value, a class B item by consequence of stock-out and a class A item by obsolescence risk.

Critical commentary

This approach to inventory classification can sometimes be misleading. Many professional inventory managers point out that the Pareto law is often misquoted. It does not say that 80 per cent of the SKUs account for only 20 per cent inventory value. It accounts for 80 per cent of inventory 'usage' or throughput value; in other words sales value. In fact, it is the slow-moving items (the C category items) that often pose the greatest challenge in inven-

tory management. Often these slow-moving items, although accounting for only 20 per cent of sales, require a large part (typically between one-half and two-thirds) of the total investment in stock. This is why slow-moving items are a real problem. Moreover, if errors in forecasting or ordering result in excess stock in 'A class' fast-moving items, it is relatively unimportant in the sense that excess stock can be sold quickly. However, excess stock in slow-moving C items will be there a long time. According to some inventory managers, it is the A items that can be left to look after themselves, it is the B and even more the C items that need controlling.

Measuring inventory

In our example of ABC classifications we used the monetary value of the annual usage of each item as a measure of inventory usage. Monetary value can also be used to measure the absolute level of inventory at any point in time. This would involve taking the number of each item in stock, multiplying it by its value (usually the cost of purchasing the item) and summing the value of all the individual items stored. This is a useful measure of the investment that an operation has in its inventories but gives no indication of how large that investment is relative to the total throughput of the operation. To do this we must compare the total number of items in stock against their rate of usage. There are two ways of doing this. The first is to calculate the amount of time the inventory would last, subject to normal demand, if it were not replenished. This is sometimes called the number of weeks' (or days', months', years' etc.) cover of the stock. The second method is to calculate how often the stock is used up in a period, usually one year. This is called the **stock turn** or turnover of stock and is the reciprocal of the stock-cover figure mentioned earlier.

Stock turn

Worked example

A small specialist wine importer holds stocks of three types of wine, Chateau A, Chateau B and Chateau C. Current stock levels are 500 cases of Chateau A, 300 cases of Chateau B and 200 cases of Chateau C. Table 12.7 shows the number of each held in stock, their cost per item and the demand per year for each.

Table 12.7 Stock, cost and demand for three stocked items

Item	*Average number in stock*	*Cost per item (£)*	*Annual demand*
Chateau A	500	3.00	2000
Chateau B	300	4.00	1500
Chateau C	200	5.00	1000

$$\begin{aligned}\text{The total value of stock} &= \Sigma\ (\text{average stock level} \times \text{cost per item}) \\ &= (500 \times 3) + (300 \times 4) + (200 \times 5) \\ &= 3700\end{aligned}$$

→

The amount of *stock cover* provided by each item stocked is as follows (assuming 50 sales weeks per year):

$$\text{Chateau A, stock cover} = \frac{\text{stock}}{\text{demand}} = \frac{500}{2000} \times 50 = 12.5 \text{ weeks}$$

$$\text{Chateau B, stock cover} = \frac{\text{stock}}{\text{demand}} = \frac{300}{1500} \times 50 = 10 \text{ weeks}$$

$$\text{Chateau C, stock cover} = \frac{\text{stock}}{\text{demand}} = \frac{200}{1000} \times 50 = 10 \text{ weeks}$$

The *stock turn* for each item is calculated as follows:

$$\text{Chateau A, stock turn} = \frac{\text{demand}}{\text{stock}} = \frac{2000}{500} = 4 \text{ times/year}$$

$$\text{Chateau B, stock turn} = \frac{\text{demand}}{\text{stock}} = \frac{1500}{300} = 5 \text{ times/year}$$

$$\text{Chateau C, stock turn} = \frac{\text{demand}}{\text{stock}} = \frac{1000}{200} = 5 \text{ times/year}$$

To find the average stock cover or stock turn for the total items in the inventory, the individual item measures can be weighted by their demand levels as a proportion of total demand (4500). Thus:

$$\text{Average stock cover} = \left(12.5 \times \frac{2000}{4500}\right) + \left(10 \times \frac{1500}{4500}\right) + \left(10 \times \frac{1000}{4500}\right)$$

$$= 11.11$$

$$\text{Average stock turn} = \left(4 \times \frac{2000}{4500}\right) + \left(5 \times \frac{1500}{4500}\right) + \left(5 \times \frac{1000}{4500}\right)$$

$$= 4.56$$

Inventory information systems

Most inventories of any significant size are managed by computerized systems. The many relatively routine calculations involved in stock control lend themselves to computerized support. This is especially so since data capture has been made more convenient through the use of barcode readers and the point-of-sale recording of sales transactions. Many commercial systems of stock control are available, although they tend to share certain common functions.

Updating stock records

Every time a transaction takes place (such as the sale of an item, the movement of an item from a warehouse into a truck or the delivery of an item into a warehouse) the position, status and possibly value of the stock will have changed. This information must be recorded so that operations managers can determine their current inventory status at any time.

Generating orders

The two major decisions we have described previously, namely how much to order and when to order, can both be made by a computerized stock control system. The first decision, setting the value of how much to order (Q), is likely to be taken only at relatively infrequent intervals. Originally almost all computer systems automatically calculated order quantities by using the EOQ formulae covered earlier. Now more sophisticated algorithms are used, often using probabilistic data and based on examining the marginal return on investing in stock. The system will hold all the information which goes into the ordering algorithm but might periodically check to see whether demand or order lead times, or any of the other parameters, have changed significantly and recalculate *Q* accordingly. The decision on when to order is a far more routine affair which computer systems make according to whatever decision rules operations managers have chosen to adopt: either continuous review or periodic review. Furthermore, the systems can automatically generate whatever documentation is required, or even transmit the re-ordering information electronically through an EDI system.

Generating inventory reports

Inventory control systems can generate regular reports of stock value for the different items stored, which can help management monitor its inventory control performance. Similarly, customer service performance, such as the number of stock-outs or the number of incomplete orders, can be regularly monitored. Some reports may be generated on an exception basis. That is, the report is generated only if some performance measure deviates from acceptable limits.

Forecasting

Inventory replenishment decisions should ideally be made with a clear understanding of forecast future demand. The inventory control system can compare actual demand against forecast and adjust the forecast in the light of actual levels of demand. Control systems of this type are treated in more detail in Chapter 14.

Short case **Manor Bakeries**[4]

Inventory management is one of the most important operations management activities at Manor Bakeries, Europe's largest manufacturer of 'ambient' packaged cakes and pies. (Ambient means that they can be stored at room temperature.) Its brands include Lyons and Mr Kipling. Its fleet of vans routinely restocks the shelves of thousands of small retailers and also distributes to major supermarkets, but here the re-ordering process is usually managed by the supermarket's own inventory management systems. Cakes are produced at four factories, on production lines, some of which are operated continuously. Although considerable effort is made to forecast sales accurately, there is always uncertainty. Yet there are limits to how much inventory can be used to compensate for demand fluctuations because supermarkets require products to be on their shelves for most of their shelf-life, allowing only a few days for Manor to transport, store and deliver the products.

Input stocks of raw materials must also be carefully managed at each factory. Bulk ingredients such as flour and sugar are delivered to giant storage silos, but managing the hundreds of other ingredients (butter, nuts, dried fruits, pasteurized egg, etc.) is more complex. Some of these are not expensive but are used in huge volumes, while others are very expensive but usage is small. Some ingredients have a short shelf-life and have to be stored in special conditions. Some are easily available, others are specially imported and are on long lead times and fresh annual crops such as fruit can vary in quality and availability. Packaging is frequently changed to reflect new promotions and price changes. Yet running out of stock is serious. It can disrupt production schedules and lead to stock-outs of finished products, affecting both sales and customer relations. Inventory also occurs because of the way products are produced on the production lines. Although some products sell enough to warrant their own production lines, most lines have been designed to make a range of similar products. So products are made in batches, sufficient to last until the next production run.

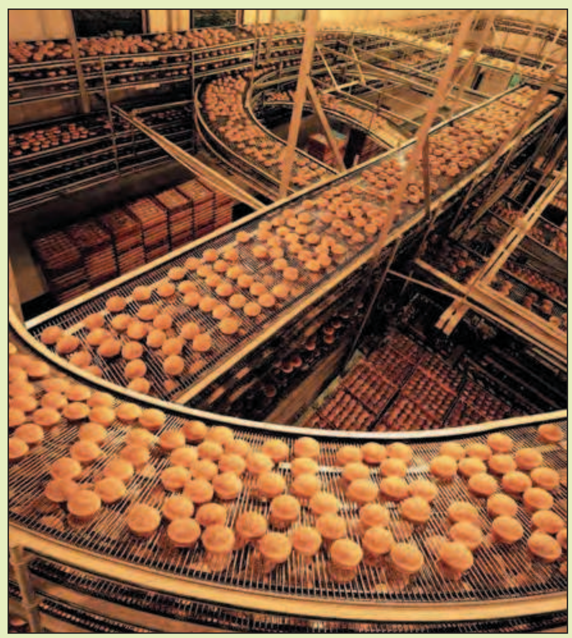

Source: RHM Ltd.

Questions

1. What are the factors which constitute inventory holding costs, order costs and stock-out costs at Manor Bakeries?
2. What makes its inventory planning and control so complex?

Common problems with inventory systems

Our description of inventory systems has been based on the assumption that operations (a) have a reasonably accurate idea of costs such as holding cost or order cost, and (b) have accurate information that really does indicate the actual level of stock and sales. But data inaccuracy often poses one of the most significant problems for inventory managers. This is because most computer-based inventory management systems are based on what is called the **perpetual inventory principle**. This is the simple idea that stock records are (or should be) automatically updated every time that items are recorded as having been received into an inventory or taken out of the inventory. So,

Perpetual inventory principle
A principle used in inventory control that inventory records should be automatically updated every time items are received or taken out of stock.

opening stock level + receipts in – despatches out = new stock level

Any errors in recording these transactions and/or in handling the physical inventory can lead to discrepancies between the recorded and actual inventory and these errors are perpetuated until physical stock checks are made (usually quite infrequently). In practice there are many opportunities for errors to occur, if only because inventory transactions are numerous. This means that it is surprisingly common for the majority of inventory records to be inaccurate. The underlying causes of errors include:

- keying errors; entering the wrong product code;
- quantity errors; a mis-count of items put into or taken from stock;
- damaged or deteriorated inventory not recorded as such, or not correctly deleted from the records when it is destroyed;
- the wrong items being taken out of stock, but the records not being corrected when they are returned to stock;
- delays between the transactions being made and the records being updated;
- items stolen from inventory (common in retail environments, but also not unusual in industrial and commercial inventories).

Summary answers to key questions

The Companion Website to the book – www.pearsoned.co.uk/slack – also has a brief 'Study Guide' to each chapter.

What is inventory?

- Inventory, or stock, is the stored accumulation of the transformed resources in an operation. Sometimes the words 'stock' and 'inventory' are also used to describe transforming resources, but the terms *stock control* and *inventory control* are nearly always used in connection with transformed resources.
- Almost all operations keep some kind of inventory, most usually of materials but also of information and customers (customer inventories are normally called queues).

Why is inventory necessary?

- Inventory occurs in operations because the timing of supply and the timing of demand do not always match. Inventories are needed, therefore, to smooth the differences between supply and demand.
- There are five main reasons for keeping inventory:
  - to cope with random or unexpected interruptions in supply or demand (buffer inventory);
  - to cope with an operation's inability to make all products simultaneously (cycle inventory);

- to allow different stages of processing to operate at different speeds and with different schedules (de-coupling inventory);
- to cope with planned fluctuations in supply or demand (anticipation inventory);
- to cope with transportation delays in the supply network (pipeline inventory).

What are the disadvantages of holding inventory?

- Inventory is often a major part of working capital, tying up money which could be used more productively elsewhere.
- If inventory is not used quickly, there is an increasing risk of damage, loss, deterioration or obsolescence.
- Inventory invariably takes up space (for example, in a warehouse) and has to be managed, stored in appropriate conditions, insured and physically handled when transactions occur. It therefore contributes to overhead costs.

How much inventory should an operation hold?

- This depends on balancing the costs associated with holding stocks against the costs associated with placing an order. The main stock-holding costs are usually related to working capital, whereas the main order costs are usually associated with the transactions necessary to generate the information to place an order.
- The best known approach to determining the amount of inventory to order is the economic order quantity (EOQ) formula. The EOQ formula can be adapted to different types of inventory profile using different stock behaviour assumptions.
- The EOQ approach, however, has been subject to a number of criticisms regarding the true cost of holding stock, the real cost of placing an order and the use of EOQ models as prescriptive devices.

When should an operation replenish its inventory?

- Partly this depends on the uncertainty of demand. Orders are usually timed to leave a certain level of average safety stock when the order arrives. The level of safety stock is influenced by the variability of both demand and the lead time of supply. These two variables are usually combined into a lead-time usage distribution.
- Using re-order level as a trigger for placing replenishment orders necessitates the continual review of inventory levels. This can be time-consuming and expensive. An alternative approach is to make replenishment orders of varying size but at fixed time periods.

How can inventory be controlled?

- The key issue here is how managers discriminate between the levels of control they apply to different stock items. The most common way of doing this is by what is known as the ABC classification of stock. This uses the Pareto principle to distinguish between the different values of, or significance placed on, types of stock.
- Inventory is usually managed through sophisticated computer-based information systems which have a number of functions: the updating of stock records, the generation of orders, the generation of inventory status reports and demand forecasts. These systems critically depend on maintaining accurate inventory records.

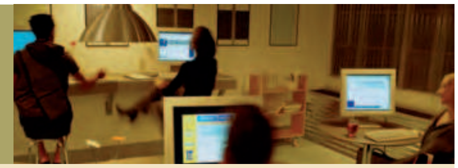

Case study
Trans-European Plastics

Trans-European Plastics (TEP) is one of Europe's largest manufacturers of plastic household items. Its French factory makes a range of over 500 products that are sold to wholesalers and large retailers throughout Europe. The company despatches orders within 24 hours of receipt using an international carrier. All customers would expect to receive their requirements in full within one week. The manufacturing operation is based on batch production, employing 24 large injection-moulding machines. Weekly production schedules are prepared by the Planning and Control office, detailing the sequence of products (moulds and colours) to be used, the quantity required for each batch and the anticipated timing of each production run. Mould changes ('set-ups') take on average three hours, at an estimated cost of €500 per set-up.

Source: Alamy/ArchivBerlin Fotoagentur GmbH

Concerned about the declining delivery reliability, increased levels of finished goods inventory and falling productivity (apparently resulting from 'split-batches' where only part of a planned production batch is produced to overcome immediate shortages), the CEO, Francis Lamouche, employed consultants to undertake a complete review of operations. On 2 January, a full physical inventory check was taken. A representative sample of 20 products from the range is shown in Table 12.8.

Table 12.8 Details of a representative sample of 20 TEP products

*Product reference number**	*Description*	*Unit manuf'g variable cost (Euro)*	*Last 12 mths' sales (000s)*	*Physical inventory 2 Jan (000s)*	*Re-order quantity (000s)*	*Standard moulding rate** (items/hour))*
016GH	Storage bin large	2.40	10	0	5	240
033KN	Storage jar + lid	3.60	60	6	4	200
041GH	10 litre bucket	0.75	2200	360	600	300
062GD	Grecian-style pot	4.50	40	15	20	180
080BR	Bathroom mirror	7.50	5	6	5	250
101KN	1 litre jug	0.90	100	22	20	600
126KN	Pack (10) bag clips	0.45	200	80	50	2000
143BB	Baby bath	3.75	50	1	2	90
169BB	Baby potty	2.25	60	0	4	180
188BQ	Barbecue table	16.20	10	8	5	120
232GD	Garden bird bath	3.00	2	6	4	200
261GH	Broom head	1.20	60	22	20	400
288KN	Pack (10) clothes pegs	1.50	10	17	50	1000
302BQ	Barbecue salad fork	0.30	5	12	8	400
351GH	Storage bin small	1.50	25	1	6	300
382KN	Round mixing bowl	0.75	800	25	80	650
421KN	Pasta jar	3.00	1	3	5	220
444GH	Wall hook	0.75	200	86	60	3000
472GH	Dustbin + lid	9.00	300	3	10	180
506BR	Soap holder	1.20	10	9	20	400

* The reference number uses the following codes for ranges:

BB = Babycare BQ = Barbecue BR = Bathroom GD = Garden GH = General household KN = Kitchen

** Moulding rate is for the product as described (e.g. includes lids, or pack quantities).

Because of current high demand for many products, the backlog of work for planned stock replenishment currently averages two weeks and so all factory orders must be planned at least that far in advance. The re-order quantities (see Table 12.8) had always been established by the Estimating Department at the time when each new product was designed and the manufacturing costs were established, based on Marketing's estimates of likely demand. Recently, however, to minimize the total cost of set-ups and to maximize capacity utilization, all products are planned for a *minimum* production run of 20 hours. The individual re-order levels have not been reviewed for several years, but were originally based on two weeks' average sales at that time. About 20 per cent of the products are very seasonal (e.g. garden range), with peak demand from April to August. Storage bins sell particularly well from October to December. The European Marketing Manager summarized the current position: '*Our coverage of the market has never been so comprehensive; we are able to offer a full range of household plastics, which appeals to most European tastes. But we will not retain our newly developed markets unless we can give distributors confidence that we will supply all their orders within one week. Unfortunately, at the moment, many receive several deliveries for each order, spread over many weeks. This certainly increases their administrative and handling costs and our haulage costs. And sometimes the shortfall is only some small, low-value items like clothes pegs*.'

The factory operates on three seven-hour shifts, Monday to Friday: 105 hours per week, for 50 weeks per year. Regular overtime, typically 15 hours on a Saturday, has been worked most of the last year. Sunday is never used for production, allowing access to machines for routine and major overhauls. Machines are laid out in groups so that each operator can be kept highly utilized, attending to at least four machines. Any product can be made on any machine. Pierre Dumas, the production manager, was concerned about storage space:

'*At the moment our warehouse is full, with products stacked on the floor in every available corner, which makes it vulnerable to damage from passing forklifts and from double-handling. We have finally agreed to approve an extension (costing over €1 million) to be constructed in June to September this year, which will replace contract warehousing and associated transport which is costing us about 5 per cent of the manufacturing costs of the stored items. The return on investment for this project is well above our current 8 per cent cost of capital. There is no viable alternative because if we run out of space, production will have to stop for a time. Some of our products occupy very large volumes of rack space. However, in the meantime we have decided to review all the re-order quantities. They seem either to result in excessive stock or too little stock to provide the service required. Large items such as the baby bath (Item 143BB) could be looked at first. This is a good starting point because the product has stable and non-seasonal demand. We estimate that it costs us around 20 per cent of the manufacturing variable costs to store such items for one year*.'

Questions

1. Why is TEP unable to deliver all its products reliably within the target of one week and what effects might that have on the distributors?
2. Applying the EBQ model, what batch size would you recommend for this product? How long will each batch take to produce and how many batches per year will be made? Should this model be applied to calculate the re-order quantity for all the products, and if not, why?
3. How would the EBQ change if the set-up costs were reduced by 50 per cent, and the holding costs were re-assessed at 40 per cent, taking account of the opportunity costs of capital at TEP?
4. What internal problems result from the current planning and control policies? In particular, analyze stock turns and availability (e.g. high and low levels).
5. Using Pareto analysis, categorize the products into Classes A, B, C, based on usage value. Would this approach be useful for categorizing and controlling stock levels of all the products at TEP?
6. What overall recommendations would you make to Francis Lamouche about the proposed investment in the warehouse extension?

Other short cases and worked answers are included in the Companion Website to this book – **www.pearsoned.co.uk/slack**

Problems

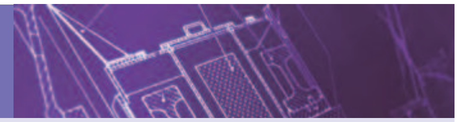

1 An electronics circuit supplier buys micro chips from a large manufacturer. Last year the company supplied 2000 specialist D/35 chips to customers. The cost of placing an order is $50 and the annual holding cost is estimated to be $2.4 per chip. How much should the company order at a time and what is the total cost of carrying inventory of this product?

2 Supermedicosupplies.com is an internet supplier of medical equipment. One of its most profitable lines is the 'thunderer' stethoscope. Demand for this product is 15,000 per year, the cost of holding the product is estimated to be €25 per year and the cost of placing an order €75. How many stethoscopes should the company order at a time?

3 Supermedicosupplies.com works a 44-week year. If the lead time between placing an order for stethoscopes and receiving them is two weeks, what is the re-order point for the thunderer stethoscopes?

4 The Super Pea Canning Company produces canned peas. It uses 10,000 litres of green dye per month. Because of the hazardous nature of this product it needs special transport; therefore the cost of placing an order is €2000. If the storage costs of holding the dye are €5 per litre per month, how much dye should be ordered at a time?

5 In the example above, if the storage costs of keeping the dye reduce to €3 per litre per month, how much will inventory costs reduce?

6 The dressings division of Strital produces sterile wound dressings in clean-room conditions. Because of the risk of cross infection the entire process must be cleaned and sterilized between batches of each product. This cleaning process takes 30 minutes. After manufacture, the dressings are stored in low-temperature sterile conditions in order to maintain their effectiveness. The cost to the company of running the storage area is £8000 per week. The demand for each type of dressing is currently 5000 per week. The manufacturing process is capable of producing at a rate of 175 dressings per hour irrespective of which product is being produced. The total operating cost (excluding materials etc.) for the production process is £1000 per hour irrespective of whether anything is being made or not. Assuming that the manufacturing process operates for 40 hours a week, (**a**) what is the economic batch quantity for the company's products, and (**b**) if demand increases to 6000 dressings per week for each product, how will the economic batch quantity change?

Study activities

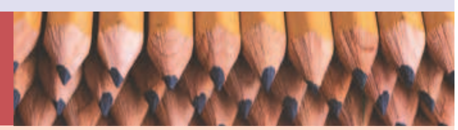

Some study activities can be answered by reading the chapter. Others will require some general knowledge of business activity and some might require an element of investigation. All have hints on how they can be answered on the Companion Website for this book that also contains more discussion questions – **www.pearsoned.co.uk/slack**

1 Estimate the annual usage value and average inventory level (or value) and space occupied by 20 representative items of food used within your household or that of your family. Using Pareto analysis, categorize this into usage-value groups (e.g. A, B, C), and calculate the average stock turn for each group. Does this analysis indicate a sensible use of capital and space, and if not, what changes might you make to the household's shopping strategy?

2 Obtain the last few years' annual report and accounts (you can usually download these from companies' websites) for two materials-processing operations (as opposed to customer- or information-processing operations) within one industrial sector. Calculate each operation's stock–turnover ratio and the proportion of inventory to current assets over the last few years. Try to explain what you think are the reasons for any differences and trends you can identify and discuss the likely advantages and disadvantages for the organizations concerned.

3 Make an appointment to visit a large petrol (gas) filling station to meet the manager. Discuss and analyze the inventory planning and control system used for fuel and other items in the shop, such as confectionery and lubricants. You should then obtain data to show how the system is working (for example, re-order points and quantities, use of forecasts to predict patterns of demand) and if possible, prepare graphs showing fluctuations in inventory levels for the selected products.

4 Using product information obtained from web searches, compare three inventory management systems (or software packages) which could be purchased by the general manager of a large state-run hospital who wishes to gain control of inventory throughout the organization. What are the claimed benefits of each system and how do they align to the theories presented in this chapter? What disadvantages might be experienced in using these approaches to inventory management and what resistance might be presented by the hospital's staff, and why?

Notes on chapter

1 Source: NBS website and discussions with NBS staff.
2 The 'stock to sales' ratio is a good indicator of the value of inventory in different businesses.
3 With special thanks to John Mathews, Howard Smith Paper Group.
4 Thanks to Manor Bakeries.

Selected further reading

Flores, B.E. and Whybark, D.C. (1987) 'Implementing Multiple Criteria ABC Analysis', *Journal of Operations Management*, Vol. 7, No. 1. An academic paper but one that gives some useful hints on the practicalities of ABC analysis.

Viale, J.D. (1997) *The Basics of Inventory Management*, Crisp Publications. Very much 'the basics', but that is exactly what most people need.

Waters, D. (2003) *Inventory Control and Management*, John Wiley and Sons Ltd. Conventional but useful coverage of the topic.

Wild, T. (2002) *Best Practice in Inventory Management*, Butterworth-Heinemann. A straightforward and readable practice-based approach to the subject.

Useful websites

http://www.inventoryops.com/dictionary.htm A great source for information on inventory management and warehouse operations.

http://www.mapnp.org/libary/opsmgnt/opsmgnt.htm General 'private' site on operations management, but with some good content.

http://www.apics.org Site of APICS, a US 'educational society for resource managers'.

http://www.inventorymanagement.com Site of the Centre for Inventory Management. Cases and links.

www.opsman.org Definitions, links and opinion on operations management.